UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND PHYSICAL SCIENCE
Electronics and Computer Science

Behavioural Biases and Agent-Based Asset Price Modelling

by

Radu T. Pruna

Thesis for the degree of Doctor of Philosophy

Supervisors: Prof. Nicholas R. Jennings, Dr. Maria Polukarov and Dr. Enrico Gerding
Examiners: Dr. Giacomo Livan and Prof. Frank McGroarty

November 16, 2018

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND PHYSICAL SCIENCE
Electronics and Computer Science

Thesis for the degree of Doctor of Philosophy

BEHAVIOURAL BIASES AND AGENT-BASED ASSET PRICE MODELLING

by Radu T. Pruna

The price, return and volume series of virtually all traded financial assets share a set of commonly observed statistical characteristics known as the stylized facts of financial data. In the last two decades, a body of literature has developed, attempting to explain these stylized facts as emerging properties from the interaction of a large number of heterogeneous market participants. The present thesis contributes to the literature on heterogeneous agent-based asset pricing models, that is, the computational study of financial markets as evolving systems of interacting agents.

Taking a prominent agent-based model (Franke and Westerhoff (2012)) as an example, we observe that its price series violates one of the core properties of real financial time series – its non-stationarity. We overcome this problem by extending the original model and drastically reduce the non-stationarity of the price series generated. Next, we estimate the model's parameters and evaluate the new setting, showing it is able to match a very rich set of stylized facts observed in real financial markets.

Now, a well defined agents-based asset pricing model able to match the widely observed properties of financial time series is valuable for testing the implications of various biases associated with investors' behaviour. In this context, we present two new behavioural asset pricing models. First, we define a setting where agents suffer from the disposition effect and test the implications of this behavioural bias on investors' interactions and price settings. We demonstrate that it has a direct impact on the returns series produced by the model, altering important stylized facts such as its heavy tails and volatility clustering. Moreover, we show that the horizon over which investors compute their wealth has no effect on the dynamics produced by the model.

Second, we present a new behavioural model of asset pricing where the agents are loss averse, and evaluate its implications. On the one hand, measuring how close the simulated time series are to its empirical counterparts, we show that the model with loss aversion better matches and explains the properties of real-world financial data, compared with the base model without the behavioural bias. On the other hand, we assess the impact of different levels of loss aversion not only on the agents' switching mechanism, but also on the properties of the time series generated by the model. We demonstrate how for different levels of the loss aversion parameter, the biased agents tend to be driven out of the market at different points in time. Since even the simplest strategies have been shown to survive the competition in an agent-based setting, we can link our findings with the behavioural finance literature, which states that investors' systematic biases lead to unexpected market behaviour, instabilities and errors.

Finally, we define a further behavioural heterogeneous agent-based asset pricing model with regret and analyse the implications of this behavioural bias on the model's dynamics. We study the coexistence of locally stable attractors of the corresponding nonlinear deterministic system, one of the most common and generic mechanisms for generating important properties observed in real financial markets. By incorporating regret in agents' expectations, we demonstrate that it can destabilise the price series and change a low volatility market regime into a highly volatile one. Consequently, we show that a change in investors' psychology contributes to the emergence of interesting new properties and that regret has the potential to explain key aspects of financial markets.

Contents

List of Figures

List of Tables

Nomenclature

p_t	Asset price at time period t
p_t^f	Fundamental price at time period t
y_t	Divident process of a risky asset t
z_t^h	Number of shares of the risky asset held by an investor type h for one time period starting at t
r	Constant risk free rate per annum
R	Gross return of the risk free asset t
K	Trading frequency t
R_t	Excess capital gain per share at time period t
a_1	Fundamentalists' risk aversion coefficient
a_2	Chartists' risk aversion coefficient
D_t^f	Fundamentalists' demand at time period t
D_t^c	Chartists' demand at time period t
D_t	Total demand at time period t
ϕ	Relative impact of fundamentalists' demand
α	Fundamentalists' price adjustments speed
χ	Relative impact of chartists' demand
γ	Chartists' extrapolation rate
ϵ_t^f	Noise in fundamentalists' demand at time period t
ϵ_t^c	Noise in chartists' demand at time period t
ϵ_t	Noise in total excess demand at time period t
σ_t^f	Noise variance in fundamentalists' demand at time period t

σ_t^c	Noise variance in chartists' demand time period t
σ_t	Noise variance in total excess demand
N	Half of population size
n_t^f	Varying fraction of fundamentalists in the market at time period t
n_t^c	Varying fraction of chartists in the market at time period t
n_1	Fixed fraction of fundamentalists in the market at time period t
n_2	Fixed fraction of chartists in the market at time period t
$q_{1,t}$	Total fraction of fundamentalists in the market at time period t
$q_{2,t}$	Total fraction of chartists in the market at time period t
u_t	Mean of chartists' learning process at time t
v_t	Variance of chartists' learning process at time t
x_t	Majority index of fundamentalists in the market at time period t
μ	Price adjustment of market impact factor
β	Intensity of choice
μ_p	Fundamental value percentage drift
σ_p	Fundamental value percentage volatility
a_t	Attractiveness level at time period t
w_t^f	Fundamentalists' wealth at time period t
w_t^c	Chartists' wealth at time period t
π_t^f	Fundamentalists' excess realised return at time period t
π_t^c	Chartists' excess realised return at time period t
λ_{LA}	Impact of loss aversion
λ_{DE}	Impact of the disposition effect
λ_{Rgrt}	Impact of regret
α_0	Predisposition parameter in the attractiveness level
α_n	Herding impact parameter in the attractiveness level
α_p	Price misalignment impact parameter in the attractiveness level

α_w	Wealth impact parameter in the attractiveness level
r_t	Asset return at time period t
m^{sim}	Vector of simulated moments
m^{emp}	Vector of empirical moments
C_h	Fixed cost associated with strategy h
J	Distance function between simulated and empirical moments
W	Weighting matrix in distance function between simulated and empirical moments
$\hat{\Sigma}$	Estimated variance-covariance matrix of the moments

Chapter 1

Introduction

A financial market is a broad term describing a place where buyers and sellers participate in the trade of various financial assets at prices that reflect supply and demand. One of the must influential ideas in modern financial markets is the Efficient Market Hypothesis (EMH), independently proposed by Paul Samuelson (1965) and Eugene Fama (1965; 1970), from two rather different research perspectives. It is well summarised by the title of Samuelson's article, "Proof that Property Anticipated Prices Fluctuate Randomly", and by Fama's well-known epithet, "prices fully reflect all available information ". It means that prices should adjust instantly and correctly to reflect new information. Accordingly, all currently available relevant information regarding a particular asset is already incorporated in the market price.

Depending on the information set available, there are different forms of the EMH. To date, most focus has been on its weak form, where the information set includes only the history of prices and returns. This form suggests that no abnormal profits can be made from analysis of historical stock prices or volume and is closely linked to the random walk hypothesis, suggesting that returns are serially independent. Thus, only new information can change an asset price and it would be an immediate reaction of the market triggered by its arrival. Moreover, prices fully reflect all available information and knowledge of past events never helps predict future values when all investors have rational expectations (Lucas Jr, 1978).

Over the years, many theoretical and empirical implications of the EMH have been tested. In particular, a range of studies raised several questions regarding the basic doctrines of the efficient markets model, especially on its view for asset price dynamics[1]. Furthermore, many critics of the EMH argue that investors are often, if not always, irrational, exhibiting predictable and financially inefficient behaviour (Lo, 2004). This comes as an alternative to the rational representative agent approach of the EMH.

[1]See the work of Grossman and Stiglitz (1980), Brock et al. (1992), Mehra and Prescott (1988), Engle (1982) and Hansen and Singleton (1983). For a review see Lo (2007).

However, most of the findings and debates regarding market efficiency and rationality are still unresolved, suggesting that markets may have non-trivial internal dynamics. Rationality implies not only that each and every agent (i.e. market participants or buyers and sellers) is rational and makes optimal decisions all the time, but it also requires full knowledge about the beliefs of all the other participants in the market. This it impossible in a heterogeneous framework, where some agents may be able to solve larger problems more quickly than others. On the other hand, heterogeneity obviously complicates the market models, making their analytical tractability almost impossible (Hommes, 2006). A computational approach is therefore better suited for investigating the heterogeneous world, ensuring tractability and offering a strong motivation for developing agent-based models of financial markets.

To this end, Agent-Based Computational Finance (ACF) is the computational study of financial markets as evolving systems of interacting agents (Tesfatsion, 2002). Generally speaking, agent-based models are systems in which a number of heterogeneous agents interact with each other and their environment (Wooldridge and Jennings, 1995). In particular, agent-based simulations dealing with economic environments are known as agent-based computational finance models, which include agent-based artificial markets (ABM) (Arthur, 1994). In this context, the financial markets can be viewed as an example of evolutionary systems populated by agents with different trading strategies.

In the last two decades a rich literature on agent-based financial models has developed, considering the financial markets as populations of different groups of agents with bounded rationality (Hommes, 2006). In this context, the term “bounded rationality”, first introduced by Simon (1955), suggests that individuals are not capable of the optimisation levels required by full rationality. Instead, because of their limitations, market participants make choices that are satisfactory, not optimal. In particular, agents behave in a manner that is as optimal with respect to their goals and resources. This means that agents usually follow strategies that have performed well in the recent past, according to different measures of accumulated wealth or profits. Interestingly, even the simplest trading strategies may survive the competition because of the co-evolution of prices and beliefs in the heterogeneous world. Thus, the new approaches shift from considering fully rational players that can be modelled and analysed mathematically, towards behavioural computer based simulations, where the numerical analysis becomes an important tool.

Nowadays, researchers do not only try to illustrate the basic mechanism of the models, but also quantitatively recreate the statistical properties of financial markets (Chen et al., 2012). Specifically, Heterogeneous Agent-Based Models (HAM) that use simple trading strategies have successfully generated a rich set of key properties that match the dynamics of asset prices in real financial markets. HAMs can provide useful insight on the behaviour of individual agents and also on the effects that emerge from their interaction, making the financial markets a very appealing application for agent-based modelling. This has led to the development of various models that aim to understand the links between empirical regularities of the markets and the complexities of the entire economic system[2].

To this end, independent studies have revealed a set of statistical properties common across different financial instruments, markets and time periods, which are difficult to explain using the asset pricing models of traditional finance. These characteristics are so robust across various financial markets, that they are termed as "stylized facts" in the econometrics literature (Cont, 2001). They target a wide range of aspects in the market, from price series and returns, to trading volume and volatility. HAMs attempt to explain these statistical properties of financial time series endogenously, by considering the interaction of market participants. Although some models have been able to provide possible explanations for various properties of financial markets, no single one has produced or explained all the important empirical features of trading data, including volume, duration, price and especially asset returns (Lux, 2008; Tesfatsion, 2002; Chen et al., 2012).

A similar line of research critiquing the EMH revolves around the preferences and behaviour of market participants. The idea is that investors behave according to their beliefs, which are updated based on new information. Hence, there are a variety of decisions in the behaviour of individuals and organisations that seem to lower their performance (De Bondt and Thaler, 1994). In order to explore the mechanisms of real markets, we need to better understand the updating processes of human beliefs. In this context, behavioural finance studies the nature of financial judgements and choices made by individuals, and examines their consequences on the markets (De Bondt et al., 2008). In contrast to other theoretical frameworks, behavioural finance models focus on the combination of investors' psychology and limits to rational arbitrage (see, e.g. Hirshleifer (2001) and Barberis and Thaler (2003)).

In recent years, a considerable focus has shifted from the econometric analysis of financial data towards developing modes of human psychology as it relates to financial markets (Shiller, 2003). The market participants only have limited information and restricted processing capabilities. Therefore, there is a strong possibility that investors have biases and deviate from rationality and this will have a large impact on the prices. These biases, also known as mental heuristics (Tversky and Kahneman, 1974), may generate predictable errors in individuals' judgement and choices (Kahneman, 2011).

However, due to the complexities of real financial markets, it is hard to assess the implications of various behavioural biases or economic policies on the interaction between participants and the overall market dynamics. For this reason, we believe that agent-based modelling and behavioural finance complement each other and can be used

[2]For surveys of this field of research see the work of Chen et al. (2012), Chiarella et al. (2008), Hommes (2006), LeBaron (2006) and Westerhoff (2010).

together. Not only can the agent-based framework serve as a useful theoretical tool for verifying the findings from behavioural finance, but it can also provide insights and which are otherwise hard to observe in the real financial markets. Models which are able to produce outcomes close to the empirical stylized facts are ideal for studying the effects of regulations, trading protocols or clearing mechanisms. In this context, the financial markets have become one of the most active research areas in agent-based modelling.

In the following section we will describe one of the widely used paradigms in agent-based modelling. We will see that numerical simulations and analysis demonstrate how most of these models are capable of generating important properties of real financial markets. Moreover, we will demonstrate how we can combine them with the findings of behavioural finance and use well-defined models to test the implications of various behavioural heuristics. Accordingly, in the next sections we will discuss some of the research challenges of agent-based financial modelling, continuing with a clear statement of the principal objectives and our contribution to the field.

1.1 Fundamentalists vs. chartists interaction

The first attempts to use agent-based financial models to explain some empirical properties observed in financial data began in the middle and late 1990s (Lux (1995, 1998); Brock and Hommes (1998); LeBaron et al. (1999)). Following the classification of Chen et al. (2012), there are, by and large, two main design paradigms of financial agents. The first one is referred to as the N-type design, considering a relatively small number of agents with simple interactions. The second one is referred to as the autonomous-agent design, where complex interactions may lead to a large number of agents. In this work we will mainly focus on the N-type design. The most familiar example of this is the fundamentalists versus chartists model (defined shortly), where there are only two types of agents. Highly motivated and backed by empirical work, we will show how this type of financial model successfully meets our objectives, discussed in Section 1.2.

Empirical evidence shows that, by and large, there are two kinds of forecasting behaviour in the market (Menkhoff and Taylor, 2007). The first class of agents, fundamentalists, act as a stabilizing force in the market. Fundamentalists base their trading strategies and expectations of future prices on economic and market fundamentals such as earnings, growth and dividends. They believe in the existence of a fundamental price and invest relatively to its value. Therefore, a fundamentalist will most likely buy an asset that is undervalued, that is, one whose price is below the believed fundamental value, and sell an asset that is overvalued, that is, one whose price is above the fundamental value. In their opinion, the price series will eventually move or converge towards a long term equilibrium or fundamental value.

On the other hand, we have chartists, also called technical analysts (Hommes, 2006). These investors forecast the future prices entirely by modelling historical data. They do not take into consideration the market fundamentals and base their decisions entirely on observed historical patterns in past prices. Technical analysts clearly believe that price developments display recurring patterns. Therefore, while the fundamentalists take into consideration only the fundamental price, the chartists base their decisions and trading strategies on historical prices.

One of the first surveys that brought to the attention of academics the use of technical analysis was conducted by Frankel and Froot (1987, 1990a, 1990b). However, their findings were viewed with scepticism because of the efficient market hypothesis, which, even in its weak form, states that it is impossible to earn excess return from forecasts based on historical price movements, once the risk-adjustments are made. Nevertheless, starting with the work of Allen and Taylor (1990, 1992), researchers began to observe the use of technical analysis in foreign exchange and describe it as an important tool in the decision making process in the financial markets. Many different questionnaire studies, summarised by Menkhoff and Taylor (2007), unanimously confirm that professional market participants indeed use technical and fundamental analysis.

In reality, one can argue that financial traders differ in many dimensions such as beliefs, trading strategies, risk attitudes, wealth, information exposure or the need for liquidity (Chen et al., 2012). Thinking of this high-dimensional heterogeneity, we should ask ourselves how much of it can be reflected in our artificial model? Ultimately, we have to decide how finely we want our agents to differ from each other. The design of the fundamentalists vs. chartists models is highly encouraged by observations made on how real financial traders behave. Notice that in such a setting the two agent types form two clusters of similar trading strategies.

Even though there are a large number of traders in the world, these two types cover virtually all the possible trading behaviours (Hommes, 2006). One can either be a fundamentalist or a technical analyst, regardless of the actual strategy used. For example, two fundamentalists may have different beliefs about the fundamental value of a particular asset, but their differences are negligible compared to those of chartists. The same reasoning applies in the case of two chartists using different methods for modelling historical data. Against this background, traders with small differences can be considered as being the same type and can be clustered together. For two traders to be of different types, their strategies have to be profoundly different.

Moreover, empirical evidence reveals that at the micro-structure level, the proportion of fundamentalists and chartists, also called market fraction, is constantly changing over time, leading to the market fraction hypothesis (Chen et al., 2012). As expected, the agents are continuously adapting to the currently observed market conditions, changing their strategies if needed. Specifically, the two classes of agents are governed by two fundamentally different beliefs regarding the price dynamics. Fundamentalists tend to believe that the current mispriced situation will soon be corrected, while the technical analysts think that in the short run it will continue. Given the current observed state of the market, each agent follows one of the two strategies. The reasoning behind this behaviour is believed to be based on a series of factors that can be included in the agent-based models. In general, traders, whether they are fundamentalists or chartists, can never be certain about the profitability of their strategies. Because of this uncertainty, they continuously adapt their behaviour by reviewing and revising their beliefs.

Throughout this thesis we are ultimately interested in constructing models that mimic the most important properties of financial data and which are analytically tractable at the same time. Since we encounter a variety of trading strategies, motives and irrationality in the real markets, such an approach is required.

More specifically, this work focuses on simple fundamentalists vs. chartist models, following some of the well-known approaches in the literature. In particular, in Chapters 3 and 4, we consider the structural stochastic volatility model (FW) introduced by Franke and Westerhoff (2011, 2012). It is acknowledged to be one of the most successful models in capturing the empirically observed traders' behaviour (Barde, 2015) . This model uses some of the well-established mechanisms for computing the agents' demands and price movements discussed in Section 2.1, with little modifications. Moreover, the market participants interact stochastically, their switching between strategies being set either by a transition probability (see Section 2.3.1) or a discrete choice approach (see Section 2.3.2).

We build on the FW model and incorporate two well-known heuristics in agents' behaviour, namely loss aversion and the disposition effect (see Section 2.4). In these new settings, we re-evaluate their parameters and show how the newly introduced behavioural biases affect the interactions between market participants and the time series generated by our models.

Furthermore, in Chapter 5 we depart from analysing the properties of the stochastic time series matched by the agent-based models and rather focus on the mechanism responsible for generating important stylized facts. Building on the work of Gaunersdorfer et al. (2008) and He et al. (2016), we consider the widely used adaptive belief system of Brock and Hommes (1998) to model regret, another important psychological model, in agents' expectations, and examine its implications on the overall market.

1.2 Research aims and objectives

In this thesis, we mainly focus on two issues in relation to heterogeneous agent-based models. First, designing the model's mechanisms and offering an explanation for the market behaviour. Second, recreating the statistical properties observed in real financial markets. In general, the aim of a well-defined agent-based model is to explain the characteristics of real financial markets as emergent properties of the interaction between participants in the market. In this vein, a key challenge is to mimic the behaviour of real markets while proposing an alternative to their apparent randomness. Specifically, we are interested in simple structures that can reproduce the empirical findings to a high degree and which are quantitatively close to the real ones. From a scientific point of view, stylized facts highlight general underlying financial mechanisms that are market-independent, with their faithful emulation being an active topic of research (see Section 2.5 for more detail).

In more detail, financial markets are complex structures with prices, trading volumes and strategies co-evolving over time, thus making it hard to find the origins of critical events or identify the structures responsible for particularly interesting dynamics. This is probably because of the important role of psychological and irrational behaviour in financial markets, well documented by the field of behavioural finance (see Section 2.4 for more details). Even in the simple fundamentalists vs. chartists models, all the interactions are continuously changing and evolving. Focusing on how agents behave (e.g. fundamentalists or chartists), we can observe the minimum conditions required to replicate the stylized facts in terms of both heterogeneity and complexity of the rules.

The models are constrained by the initial conditions imposed at the microlevel, but all the dynamics are governed by agents' interactions. In this way, we increase the transparency and the clarity at the system level, making it easier to analyse the results. In this context, finding the mechanisms responsible for generating the stylized facts is an important aim of our work. Specifically, noise enters at a micro level (pairwise agent interactions) and makes it difficult to assess the main causes for the observed properties at the macrolevel. Moreover, we need to ensure that the simulated experiments reflect some of the fundamental aspects of the considered problem.

Nevertheless, the economic estimation is an ambitious task. Just as for any other model, the parameters of the ABMs must be calibrated to real financial data in order to perform realistic emulation. This part of calibration is together with architecture design one of the most technical and crucial aspects of our work. We can identify and estimate all the ingredients of the financial model, but it is hard to be assured that these models are significant from an economical point of view, or that the estimations are stable. However, given the reduced simplicity of the agent-based models, it is easier to identify the underlying dependencies between its mechanisms and possible explanations can be tested.

A further aspect we consider is the modelling of the agents themselves. A vital question is how to model the behaviour of the agents that populate our system. One way would be to make the agents follow a logical paradigm as in traditional artificial intelligence (see

Brooks (1991) for an introduction), while another would consider a more psychological approach, as in evolutionary finance (see Evstigneev et al. (2008) for a survey on research and applications of evolutionary finance). Most of the models have unpretentious learning algorithms, in the form of simple updating equations with fixed parameters. No matter if we consider real or simulated financial markets, evidence strongly suggests that there is no single algorithm or strategy that performs best in all the situations (Chen et al., 2012).

Despite this, there have been attempts to model agents with complex learning strategies, the most famous example being the Santa Fe Institute artificial stock market. However, such extensions do not provide additional explanatory power. On the contrary, adding complexity usually implies less tractability. Specifically, complex interactions decrease the transparency and the clarity of the modelled process. This setting may be useful for finding long term equilibria, but it does not serve well the purpose of understanding the macrolevel effects of the micro interactions, which is one of our aims. Thus, on one hand side we want a model that can recreate all the important stylized facts, but on the other hand we are interested in the simplest possible one that can be used to estimate the heterogeneity.

In the young field of agent-based computational finance modelling, the research has been mostly computationally and theoretically oriented. There are a few attempts to estimate a model on economic data (see for example Boswijk et al. (2007), Franke and Westerhoff (2011) or Alfarano et al. (2007)) but little has been done on implementing economical and behavioural policies (Westerhoff and Franke, 2013). Heterogeneity is important in explaining the data generated, but much work is needed to investigate the robustness of empirical findings (Tesfatsion, 2002), this being one of our central objectives. It is important because we believe that for a better representation of how investors behave, agent-based financial models should capture the behavioural factors observed in real markets (LeBaron, 2006). Specifically, as the agents try to model real-life traders and should represent their behavioural biases to some extent.

To this end, a further objective of this thesis is to use well-defined heterogeneous agent-based models as a tool for incorporating psychological attributes in agents' behaviour and testing their implications on both the agents' interactions and the time series generated by the model. This implies that the realism of economic agents can be greatly increased by considering numerous cognitive and behavioural biases emulating those of human investors. In turn, this can lead to models that better match the properties of real financial markets.

1.3 Research contributions

Agent-based models that rely on simple trading strategies have proven themselves very efficient in generating important dynamics of real financial markets. In their original design, Franke and Westerhoff show that their model is able to successfully match a series of stylized facts including heavy tails, volatility clustering, long memory in absolute returns and absence of autocorrelation in raw returns.

However, the price series generated by the FW model violates one of the core properties of financial time series – its non-stationarity (see Section 2.5.1). In Chapter 3 we overcome this problem by extending the original model and changing the motion of the fundamental value over time. As a result, we drastically reduce the non-stationarity of the price series generated and ensure a better representation of the properties of real financial data, one of our research aims. Furthermore, after estimating the model's parameters, we show that it is able to match a very rich set of stylized facts of real financial markets. Thus, our model can be used for studying the effects of different behavioural factors and preferences.

For this reason, in Chapter 4 we address a further objective of our research (see Section 1.2) and modify the model to capture two of the most discussed empirical findings of behavioural finance, namely loss aversion and the disposition effect (see Section 2.4 for more details). They are both derived from prospect theory and refer to the asymmetry in how an agent perceives gains and losses. In order to accommodate for these preference, we use the way agents switch between strategies and define the models with loss aversion (WHP-LA) and the disposition effect (HPM-DE) in Sections 4.1.3 and 4.1.2, respectively. Moreover, we test the implications of these behavioural heuristics on both agents' interactions and the time series generated by our models (see Sections 4.3 for more details).

In addition, in Chapter 5 we address a further aim of our research and focus on the implications of regret, another well-known psychological model (see Section 2.4.3), on the mechanism responsible for generating important stylized facts. In this context, we define a new agent-based model that builds on the well-known adaptive belief system (see Section 2.3.2) and for the first time incorporates regret in agents' expectations (see Section 5.1.2). In doing so, we tackle another objective of our research, that agent based models can be used to test the implications of various behavioural factors.

Specifically, we make the following contributions:

- Checking for stationarity in the price series generated by the FW model, we observe that the unit root test is rejected in more than one in every four simulations (see Section 3.1.3 for more details). This strongly contradicts the well-known fact that financial price series are non-stationary, one of their most well-known stylized facts. Therefore, we build on the original model by making a novel change in the motion of the fundamental value over time. Accordingly, we assume it follows a geometric Brownian motion (GBM), the most widely used model of describing stock price behaviour (Hull, 2006). As a consequence, the price series generated by our modified model (FW+) rejects the unit root test for stationarity in less than 5% of the simulations, thus offering a better representation of the properties of financial data (see Section 3.1.3).

- One of the main objectives of agent-based financial modelling is to recreate the most important stylized facts of financial data. In this context, in Section 3.3 we show that our FW+ model is able to generate and match a rich set of properties usually observed in real financial markets including: lack of predictability and non-stationarity, absence of autocorrelation in raw returns, fat tails, volatility clustering and long memory in absolute returns, volume-volatility relations, aggregate Gaussianity, price impact and extreme price events. To date, this is the only model that is reported to match such a rich set of the stylized facts of real financial markets. Therefore, we illustrate one of the essential aims of agent-based financial modelling. Namely, that many of the stylized facts arise and can be explained from the interaction of market participants.

- While in the FW+ model the agents' proportions continuously change without a total domination of one type, in line with the market fraction hypothesis, in the WHP-LA setting we observe major differences in agents' interactions. Specifically, we show that for certain degrees of loss aversion, the agents are driven out of the market, leading to their complete extinction at different points in time, for the first time observed in an agent-based setting (see Section 4.3.2). This comes as a direct disruption caused by the behavioural bias since in a classical agent-based setting the traders' market fractions are continuously changing and even zero-intelligence investors have been shown to survive the competition.

- We further test the implications of loss aversion on the model's time series by measuring how close they are to their empirical counterparts. In Section 4.4 we show that our model with loss aversion better matches the data's stylized facts, compared to the same model without loss aversion.

- Similarly, we test the implications of the disposition effect on the both agents' interactions and price dynamics in the HPM-DE setting. For the first time in an agent-based model we find that while the time horizon over which investors consider their wealth has no impact on the overall setting, the level of the behavioural bias directly impacts the returns series produced by the model, altering its stylized facts and leading to disruptive behaviour (see Section 4.3.1 for more details).

- Finally, we depart from analysing the stochastic time series of our models and focus on the actual mechanism responsible for generating them. In Section 5.3 we show that an increase in agents' regret can destabilise the market and transform a low volatility regime characterized by a stable steady state into a highly volatile one characterized by a stable limit circle. Since the interplay between these coexisting attractors leads to important properties observed in financial markets, we conclude that regret can explain some key aspects of financial markets, marking our final research contribution.

1.4 Thesis outline

The remainder of the thesis is organised as follows:

- Chapter 2 provides a background on Agent-Based Financial Models, considering the most important developments in the field. We focus on fundamentalists vs. chartists models, presenting their evolution over time. Next, we present some of the latest advancements in the field of behavioural finance and their relation to agent-based models. Finally, we provide a formal definition of the commonly observed stylized facts in financial markets, their empirical confirmation and the existing models that matched them.

- Chapter 3 begins with a formal definition of the asset pricing model introduced by Franke and Westerhoff. We continue with the estimation and validation of the model. This is followed by our novel modification regarding the motion of the fundamental value and an in-depth analysis of the results. We will demonstrate that the model is able to recreate a rich set of the stylized facts of financial markets.

- Chapter 4 presents two new behavioural asset pricing models, one with loss aversion and another with the disposition effect in agents' interactions. We discuss their key dynamics including how prices are generated and how agents interact with each other. Next, we look at the differences made by the behavioural biases on both the agents' interactions and the price changes. We provide an in-depth analysis of the time series generated by the model with loss aversion and show that it is able to match a rich set of stylized facts.

- Chapter 5 presents a different behavioural model, this time incorporating regret in agents' expectations. We start by discussing the model's dynamics and focus on the stability conditions of its steady state. Finally, we explore the destabilising effect of regret.

- Chapter 6 summarises the outcomes of the research described in this work. We also provide a detailed description of the planned future research directions.

Chapter 2

Literature Review

In this chapter we present some of the most prominent work that has been carried out in the area of Agent-Based Financial Modelling, in relation to our research aims and objectives outlined in the previous chapter. We begin by discussing the structure and dynamics of some of the most relevant existing models, from the simple deterministic approaches discussed in Section 2.1 to more complex models that incorporate simple interactions in Section 2.2. Moreover, in Section 2.3 we present some of the more recent and advanced models, in which agents interact stochastically and that are capable of recreating some of the real life market dynamics. Next, since one of our main objectives is to test the implications of various behavioural policies in an agent-based setting, in Section 2.4 we present some of the latest advancements in the field of behavioural finance and their relation to agent-based models. Finally, we focus on another key aim of this work, that is recreating the properties of real financial time series, and discuss the most relevant stylized facts of financial data in Section 2.5. Here, not only do we offer an empirical background for the most important statistical properties of financial markets, but we also analyse how some of them are matched by existing agent-based models.

2.1 Agent-based models of asset pricing

Over the years, the fields of finance and economics have developed different approaches to model financial asset prices' dynamics. Historically, we distinguish between three important research directions. The first and most widely used are econometric statistical models calibrated to fit empirical time series such as past prices or returns (Engle, 2001). These can be useful in understanding or forecasting some of the statistics of the data, as long as the calibrated parameters stay more or less the same. The second class of models, known as Dynamics Stochastic General Equilibrium (DSGE) models, explicitly model agent-based micro foundations (Sbordone et al., 2010), usually by considering an utility maximising representative agent. Recent developments in the field try to make the models more realistic by accounting for agent heterogeneity, imperfect learning or incomplete information. In general, they try to replace the rational expectations hypothesis and recreate real life financial market conditions (Massaro, 2013). These two classes of models have shown various promising results over the years. However, since they adopt a top-down approach to system inference, we can say that they are rough approximations of reality (Farmer and Foley, 2009) and will not explain the diversity of market micro-structure usually observed in the real world.

This leads to a third class of models, called agent-based models or multi-agent systems, which try to recreate markets from a pure bottom-up approach and consider them as complex systems (LeBaron, 2002). As discussed in Chapter 1, the heterogeneous agents approach was developed as an alternative to the traditional, rational agent of efficient market hypothesis. However, many ideas in the agent-based approach have a long history in economics, even longer than the rational expectations and efficient market hypothesis. In more detail, Keynes (1937) argued that an investor's psychology plays an important role in financial markets. According to his theory, it is hard to compute the true fundamental value of a stock and for an investor it may be easier and less risky to estimate its value by a rule of thumb, similar to Simon's view that agents have limited knowledge about the markets and that they use simple but satisfactory trading strategies (for more details see Chapter 1).

Now, one of the first heterogeneous agent-based models for the stock market was developed by Zeeman (1974). Despite the lack of any micro foundations, the model incorporates some of the most important behavioural elements seen in more recent agent models and is able to recreate temporary bull and bear markets. He emphasized that in a market populated by fundamentalists and chartists (see Section 1.1), their behavioural assumptions and interaction explain the switching between bull and bear markets.

A few years later, Beja and Goldman (1980) introduce one of the first models that incorporates trend followers (chartists) and value investors (fundamentalists). With a behavioural model, they show that the price equilibrium is unstable by assuming the agents have linear trading rules. In the classical general equilibrium setting (Walrasian) the excess market demands (or, conversely, excess market supplies) must add up to zero. In this context, by excess demand and supply we mean their precise positive and negative values, respectively. Accordingly, the demands of fundamentalists and chartists would have to sum to zero. This assumption contradicts the usual setting of real financial markets, which may experience excess demand and supply (Keynesian setting).

In order to account for the Keynesian setting, the authors propose an agent-based model with a market maker that adjusts the prices according to cumulative excess demand. They provide an alternative to the perfect Walrasian market by proposing a price formation process that admits finite adjustment speed. Their framework permits transactions at disequilibrium prices. Furthermore, they assume that the market maker is risk neutral, setting the price in response to received orders, without worrying about accumulated inventory.

Mathematically, the price movements assume a finite adjustment speed in the direction of asset demand:

$$\frac{dp}{dt} = D_t^f + D_t^c, \tag{2.1}$$

where D_t^f and D_t^c represent the excess demand of fundamentalists and chartists, respectively. Let p_t^f denote the exogenously generated fundamental price and p_t the endogenously determined trading price. The fundamental excess demand is then formalized as a linear function of the form:

$$D_t^f = a(p_t^f - p_t), \tag{2.2}$$

where the coefficient $a > 0$ measures the relative impact of fundamental demand upon price movements. This format of fundamentalists' excess demand, standard in the literature, will be used in our fundamentalists vs. chartists models presented in Chapters 3 4 and 5.

With regards to expressing chartists' excess demand, there are more variations considered in the literature. Specifically, Beja and Goldman (1980) assume that it depends on the price change, formulated as the difference between their assessment of the current price trend, ψ_t, and the exogenously given return on an alternative security, g_t. Similarly to fundamentalists', chartists' excess demand is a linear function of the price change, that is:

$$D_t^c = b(\psi_t - g_t), \tag{2.3}$$

where the coefficient $b > 0$ measures the relative price impact of chartist demand upon price movements. The trend estimator ψ is adjusted using an adaptive system with a constant adaptation speed. In the agent-based models we will discuss in the next chapters, we make use of a widely used particular case of this general form and model chartists as simple trend followers that observe only the latest price changes. According to Equations 2.1 - 2.3, the price adjustment equation becomes:

$$\frac{dp}{dt} = a(p_t^f - p_t) + b(\psi_t - g_t) + e_t, \tag{2.4}$$

where e_t denotes an additional noise term. This general form describing the price dynamics in a two type model has been successfully preserved over the years and it is very similar to the price equations we will consider in the following chapters.

Chiarella (1992) extends the Beja and Goldman (1980) model in continuous time by considering non-linear trading strategies for chartists, however accompanied by the limiting assumptions of constant fundamental price $p_t^f = p_f$ and constant return on alternative investments $g_t = g$. They show that the price equilibrium is stable as long as the fraction of chartists is sufficiently low. When the chartists' market fraction exceeds a threshold, the equilibrium becomes unstable and is replaced by a limit cycle. That is, the excess demand of each trader type oscillates, causing sustained deviations from the equilibrium price. In particular, the prices fluctuate along the limit cycle without converging to the fundamental value. In addition, the interplay between different attractors has been proven to be a key mechanism for generating volatility clustering, an important property observed in real-life financial markets (He et al., 2016). This addresses one of the key aims of agent-based models (see Section 1.2), which we will discuss in more detail in Chapter 5.

A further influential discrete time model, introduced by Day and Huang (1990), considers two types of investors, one that base their decisions on a combination of economic fundamentals (similar to fundamentalists) and another that believe that the investment rules can be extrapolated from past deviations from fundamental value (similar to chartists). As in Beja and Goldman (1980), the presence of a market maker comes as a justification of the price adjustment rule represented by Equation 2.1. The market maker mediates the transactions between investors and provides liquidity. He sets the price by supplying stock out of its inventory and raising the price if there is excess demand, while accumulating stock and lowering the price when there is excess supply. The market maker should therefore be viewed as a stock exchange. Their model is one of the first to exhibit complex, chaotic price fluctuations around a fundamental price, qualitatively similar to real stock market movements. The authors also demonstrate that a large market fraction of chartists does not guarantee that the price will converge to its fundamental equilibrium. Instead, it creates bubble episodes of over- and undervaluation.

A similar price setting rule called the market impact function has been developed by Farmer and Joshi (2002), and since then used by many influential agent-based models, including the ones presented in the next chapters. Their original setting consists of N directional traders who place market orders always filled by a market maker. As before, the prices are increased (decreased) when there is excess demand (supply). However, there is a difference in the algorithm used by the market maker to set the prices. More specifically, the authors assume that the relative increase of the price from time period t to $t+1$ is an increasing function ϕ of the current demand:

$$\frac{p_{t+1}}{p_t} = \phi(D_t), \tag{2.5}$$

where $\phi' > 0$ and $\phi(0) = 1$. Taking logarithms and expanding in a Taylor series gives:

$$\log(p_{t+1}) - \log(p_t) \approx \frac{D_t}{\lambda}, \tag{2.6}$$

where $\lambda := 1/\phi'(0)$ normalizes the order size, called market depth or liquidity. Note that the price updating rule is essentially the same as the price adjustment of Equation 2.1 used in Beja and Goldman (1980), Day and Huang (1990) and Chiarella (1992), except that now it holds the logarithm of the price instead of the price itself. Their model is able to reproduce a series of stylized facts observed in financial markets, such as noise amplification, excess and clustered volatility.

The basic elements introduced in this section are still used in more recent models, with little modifications, as we will see in the following sections. In our work, presented in detail in the following chapters, we consider behavioural agent-based models that are similar in many ways to the ones introduced more than 30 years ago. Specifically, we will consider fundamentalists vs. chartists models where the agents' demand functions are similar to the ones of Equations 2.2 and 2.3. Moreover, because of its popularity and mathematical tractability, the (log) pricing mechanism considered in Chapters 3 and 4 will be established by a market maker, taking into consideration a market impact factor, similar to Equation 2.4.

2.2 Interacting agents

All the models presented so far are deterministic, meaning that the agents do not interact with each other. One may think that, due to law of large numbers, stochastic interactions have little to no effect on the aggregate variables. However, this is not the case. Simple local interactions may generate sophisticated behaviour at the macro level, as we will demonstrate in the next chapters.

In this vein, Kirman (1991) proposes an exchange rate model based on the interaction between fundamentalists and chartists. More specifically, the switching mechanism between agents is based on a herding mechanism (see Bikhchandani and Sharma (2000) for a review of herding behaviour in financial markets) inspired from the communication behaviour of ants (Kirman, 1993). When ants face two identical food sources in the vicinity of their nest, a majority of the population will concentrate on one of the food sources at any point in time. This is because of the chemical information transmission via pheromones by which the first ants that visited the food source recruit followers guiding them to the same site. However, after enough time their proportions switch randomly. Hence, ants facing symmetric situation behave asymmetrically.

In financial markets, some people often have better information than others so it makes sense for investors to consider the decisions of other market participants, especially if they are thought to be better informed. Hence, one way of overcoming informational problems is to simply copy the behaviour of others. Herding on its own can be either rational or irrational and in many situations it is linked to bubbles and market inefficiencies (Shiller, 2015). Agents usually follow others because they are exposed to similar information and share the same mental frames, or simply because they don't know what else to do. This was later recognised as an essential ingredient of any agent-based financial model (Alfi et al., 2009) and will be also included in our agent-based models discussed in the next chapters.

In one of the first models with a herding mechanism, Kirman (1993) proposed an elegant framework based on ants' interactions. In a two-type agent-based model, an investor's choices and expectations are affected by random meetings with other market participants or spontaneous changes. In this context, Kirman investigates the equilibrium distribution of this model and shows that it depends on the magnitudes of the probabilities of being converted and of self-conversion. More specifically, he finds that investors usually maintain their strategy, with only occasional shifts between them. When the market is dominated by fundamentalists, the exchange rate is stable and close to the fundamental value. In contrast, if chartists dominate the market, the exchange rate is either unstable or stable but far from the unit root process. Moreover, when chartists (fundamentalists) dominate the market, the volatility of the exchange rate fluctuations is high (low). Therefore, with regular switches between the two strategies, the phenomenon of clustered volatility can be observed. In a later paper, Kirman and Teyssière (2002) discuss the stylized facts generated by the model in detail, demonstrating that it also produces long memory in volatility.

The ant model marked the beginning of new horizons in the literature of Agent-Based Computational Finance Modelling. The researchers began to use econometric techniques as tools for estimating and validating the models. For this reason, the new generation of models were starting to successfully generate important properties observed in real-life financial markets. This has become one of the key objectives of agent-based financial models (see Section 1.2), which we will discuss in more detail in Section 2.5.

Taking a step further, the emerging literature was not only concerned with creating models that can match the stylized facts, but also with the estimation of their parameters. One way of estimating the parameters is by using a direct method that applies statistical techniques to the aggregate equation in which all the model's parameters are encapsulated. However, the agent-based models are hard to estimate because of their complexity induced by the large number of degrees of freedom and the absence of such criterion functions, leading to one of the main challenges in the field (see Section 1.2 for an outline of our research challenges and objectives). This inspired the development of new procedures, referred to as the simulation-based econometric methods (Chen et al., 2012).

The key idea in the simulation-based methods is to form a vector of parameters which will be calibrated such that the properties generated by the model match those of real data. Since we are mainly interested in the emerging properties of ABM, one of the most widely used technique is the Method of Simulated Moments (MSM) (McFadden, 1989)[1]. In this, we first choose the vector of parameters used by the agent model to generate the time series (i.e. asset price series). We then compare a set of statistics

[1]Other estimation methods used in the literature are the maximum likelihood method, nonlinear least squares, ordinal least squares and interactive evolutionary computation. See Chen et al. (2012) for a detailed survey of these methods and models implementing them.

(moments) generated by the simulated time series (simulated moments), with those of real data (sample moments). The objective is to minimise the distance function between these two sets by searching over the entire parameter space.

Formally, let X be the set of real statistics derived from the real data, $\mathbf{X} = (X_1, X_2, ..., X_n)$, and Y the set of moments from the simulated data, $\mathbf{Y} = (Y_1, Y_2, ..., Y_n)$. Let $\mathcal{L}$ be the distance function between $\mathbf{X}$ and $\mathbf{Y}$. Then, the MSM estimation involves solving for the set of parameters θ:

$$\theta^* = \arg\min_{\theta\in\Theta} \mathcal{L}(\mathbf{X}, \mathbf{Y}; \Theta). \tag{2.7}$$

The first estimation was conducted by Winker et al. (2001), who applied the method of simulated moments to estimate Kirman's ant model (1991, 1993) for the U.S.$ and German Deutsche Mark Rate (DEM) exchange. They tried to simulate the two parameters of the model, namely the self-conversion rate and the probability of being converted. Same authors (Gilli and Winker, 2003) considered a three-parameter ant model by including a noise related parameter. They found significant changes in the market fraction between fundamentalists and chartists, each agent type having periods of dominating the market. In the next chapters (see Sections 3.2 and 4.2), we will estimate our models' parameters using a similar method of simulated moments technique, first developed by (Franke and Westerhoff, 2016).

The simple agent interactions proposed by Kirman (1993) provided the basis for further improvements. Because of their key role in replicating real-life properties, more complex, stochastic interactions have been developed throughout the years (see Section 2.3). This is a huge step forward towards a better representation of how real financial traders behave, which is one of the aims in agent-based financial modelling (see Section 1.2).

2.3 Social interactions

In this section we discuss a more advanced class of models in the chartist-fundamentalist literature. While Kirman allows pair-wise interactions only, these models use a mean-field approach (Hommes, 2006). Specifically, this implies that all agents influence all other agents in the market with the same intensity. Social interaction assumes that the individuals will be influenced by the overall mood of the market and that the strategy of each agent depends directly on the choices of the others. The frameworks discussed in this section bring us a step closer to one of our research aims, that is a better representation of how real life investors behave. Moreover, the introduction of social interactions in agent-based models has led to structures that are capable of recreating the stylized facts of financial markets, another important objective of our work.

Recent laboratory experiments in Hommes (2011); Anufriev and Hommes (2012); Hommes et al. (2017) also show that agents using simple rule of thumb trading strategies are able to coordinate on a common prediction rule. The authors emphasise that heterogeneity in expectations and adaptive behaviour are crucial to describe individual forecasting and aggregate price behaviour. Such an influence gives us the possibility of studying the aggregate behaviour more easily in an agent-based framework and also means that small changes at the individual level cause large changes at the macro level. In this context, in Chapters 4 and 5 we test the effects of important behavioural biases by studying the macrodynamic changes brought by mirolevel agents' interactions (see Sections 4.3 and 5.3 for a more detailed discussion).

In these settings, agents can change their beliefs and switch to using different strategies. In particular, the agents' switching mechanism has been dominated by two different frameworks. The first one, the transition probability approach (TPA) was introduced by Lux (1995), while the second, known as the discrete choice approach (DCA) was introduced by Brock and Hommes (1997). In the models presented Chapters 4 and 5, we consider stochastic interactions between market participants and model specific behavioural heuristics using both mechanisms. In the following sections we discuss the two frameworks in detail.

2.3.1 The transition probability approach

One particular widely used framework was first introduced by Lux (1995) and further developed by Lux (1998) and Lux and Marchesi (1999, 2000). They describe the market as being populated by $2N$ speculative agents (chartists), divided in two groups, optimists and pessimists. At time t, there are n_+ optimists and n_- pessimists, i.e. buyers and sellers, such that $n_+ + n_- = 2N$. The overall average of speculators is captured by the opinion index x:

$$x = \frac{n_+ - n_-}{2N} \tag{2.8}$$

Obviously, $x = 0$ corresponds to a balanced situation where the number of optimists equals the number of pessimists, whereas $x = 1$ ($x = -1$) means that there are only optimists (pessimists) in the market.

In Lux's model, the switching mechanism between agents is captured by two determinants: a herding mechanism and the profit differential. Being a social interactions model, it assumes that all other individuals have the same influence on any one. The feedback received can lead to migrations of agents between the two groups. Formally, these transition probabilities can be modelled by Poisson distributions in continuous time. Let p_{+-} and p_{-+} be the probabilities of a pessimistic trader becoming optimistic and vice-versa. Then, we have:

$$p_{+-} = v\, exp(\alpha x), \quad p_{-+} = v\, exp(-\alpha x), \tag{2.9}$$

where α captures the strength of the herding effect and v measures the speed of the infection. Consequently, we expect a fraction n_-p_{+-} to switch from n_- to n_+ and similarly the other way around. A simplifying assumption is that every individual may change his opinion only once at each iteration.

In the next step, the mechanism of contagion is linked to the price formation process in order to provide a satisfactory description of the stock market dynamics. The approach is similar to the one of Beja and Goldman (1980). The optimistic individuals are assumed to buy additional units of the asset, while the pessimists will sell some of the asset, entering the supply side. Let t_N be the number of stocks each individual can buy or sell. This leads to the aggregate demand, D_N, of the speculators being defined as:

$$D_N = n_+t_N - n_-t_N = xT_N, \quad T_N := 2Nt_N, \tag{2.10}$$

where T_N denotes the trading volume of speculators.

Lux introduces a third class of traders, fundamentalists, whose excess demand depends on the difference between the fundamental value p_f and the actual price p. Hence, the aggregate demand, D_F, becomes

$$D_F = T_F(p_f - p), \tag{2.11}$$

where $T_F > 0$ captures the aggregate trading volume of fundamentalists. A market-maker is assumed to provide liquidity and adjusts the prices proportional to the aggregate net asset demands of both traders:

$$\dot{p} = \frac{dp}{dt} = \mu(D_N + D_F) = \mu[xT_N + T_F(p_f - p_t)], \tag{2.12}$$

where μ is the speed of price adjustment. In order to encapsulate both speculative behaviour and fundamentalism, the transition probabilities are updated as follows:

$$p_{+-} = v\, exp(\alpha_1\dot{p}/v + \alpha_2 x), \quad p_{-+} = v\, exp(-\alpha_1\dot{p}/v - \alpha_2 x), \tag{2.13}$$

where v is the speed of opinion changes, as before. But α has now been split between α_1 which is a weight factor describing how much information the traders try to draw from the prices and α_2 which determines the strength of contagion from other trades.

Lux (1995) demonstrates that the trading price p has a unique equilibrium when $x = 0$ and $p = p_f$ for $\alpha_2 \leq 1$, with two additional bubble equilibria at $p \neq p_f$. The bubble equilibria become further displaced from the fundamental price, as we observe a larger trading volume of chartists, T_N, relative to the fundamentalists, T_F, and an increasing infection parameter α_2. We should note that even when the equilibrium is unique, it can be either stable or unstable. If the values of T_N, v, and α_1 are high and T_F is low with respect to T_N then the favour is towards instability. Furthermore, if the equilibrium is unique and unstable, there exists at least one stable limit cycle such that all trajectories of the system converge to a periodic orbit. In such cases, the model explains periodic switches between overvaluation and undervaluation (bull and bear markets) generated by the endogenous processes of mimetic contagion and trend following. Lux and Marchesi (1999 and 2000) showed that the model is able to match a series of important stylized facts of real financial markets including uncorrelated returns with clustered, long range dependent volatility and heavy tails.

Alfarano, Lux and Wagner propose a simple two-type version of the Lux model including fundamentalists and noise traders, who are subject to irrational fads and moods, originally introduced by De Long et al. (1990). The model, restricted to having only the herding mechanism in the transition probabilities, is estimated on the Germany market (Alfarano et al., 2005), Australian market (Alfarano et al., 2006) and Japanese market (Alfarano et al., 2007). They try to characterize the markets by a dominance of either fundamentalists' or noise traders' activity through estimating the magnitude of the parameters in the switching mechanism. With the exception of the Australian stock market, the noise traders dominate almost every individual stock in Japan and Germany. They are also able to match some of the stylized facts of financial data including heavy tails and volatility clustering as emergent properties of the interaction among traders. Since one of our work's objectives is to test the effects of various behavioural biases in a well-defined agent-based setting (see Section 1.2), we will use the transition probability approach defined in this section, that is capable of recreating the dynamics of real life financial markets, to model our agent's interactions and heuristics (see Section 4.1.2 for more details).

2.3.2 The adaptive beliefs system

In this section we discuss the adaptive belief system (ABS) introduced by Brock and Hommes (1997, 1998), another widely influential framework of social interactions. An ABS can be described as a standard discounted value asset pricing model derived from mean-variance maximization and extended to the case of heterogeneous beliefs (Hommes, 2006). As we have seen so far, the agents have bounded rationality and select their strategy based on recent performance.

In the original setting, agents can either invest in a risky asset or a risk free one. The risk free one is always supplied and pays a fixed return R, while the risky one (e.g. a stock) pays an uncertain dividend. In more detail, let p_t be the stochastic price per share of the risky asset at time t, and let y_t be its stochastic dividend process. Given this, denote by $R_{t+1} = p_{t+1} + y_{t+1} - Rp_t$ the excess capital gain per share made on the risky asset at $t+1$. Note that while R represents the fixed return of the risk free asset, R_t is the excess gain per share made by an agent that invested in the risky one.

It follows that the wealth of investor of type h at $t+1$ is given by:

$$W_{h,t+1} = RW_{h,t} + (p_{t+1} + y_{t+1} - Rp_t)z_{h,t}, \tag{2.14}$$

where $W_{h,t}$ is the investor's wealth at time t and $z_{h,t}$ represents the number of shares of the risky asset (demand) held by investor type h, one time period from t to $t+1$. Note that by type h we refer to a subclass of investors all following the same particular strategy. It is assumed that the investors are myopic mean variance maximizers, so that the optimal demand $z_{h,t}$ of each trader of type h for the risky asset solves:

$$z_{h,t} = \max_{z_{h,t}} \left\{ E_{h,t}[R_{t+1}] - \frac{a_h}{2} V_{h,t}[R_{t+1}] \right\}, \tag{2.15}$$

i.e.

$$z_{h,t} = \frac{E_{h,t}[R_{t+1}]}{a_h V_{h,t}[R_{t+1}]}, \tag{2.16}$$

where $E_{h,t}$ and $V_{h,t}$ denote the subjective beliefs of trader h about the expectation and variance of excess return and a is the risk aversion parameter. Note that $z_{h,t}$ is not the demand of trader h but rather its target holdings, derived from maximization of absolute utility. It will coincide with the actual demand only if the trader has no prior holdings. Moreover, this mean variance optimisation is equivalent to maximising a constant absolute risk aversion (CARA) utility wealth function $U_h(W) = -\exp(-a_h W)$, where a_h represents a trader's type h risk aversion coefficient.

If we have H different trader types, the aggregate demand becomes:

$$z_t = \sum_{h=1}^{H} n_{h,t} \frac{E_{h,t}[R_{t+1}]}{a_h V_{h,t}[R_{t+1}]}, \tag{2.17}$$

where $n_{h,t}$ denotes the fraction of trader type h in the market at time t.

Brock and Hommes introduce a key endogenous selection mechanism for updating the market fractions of agent type h using a multinomial logit model of discrete choice. Hence the name of the discrete choice approach (DCA). Each strategy gets a fitness function $A_{h,t}$ measuring the performance of strategy h at time t. An example of performance can be the realized profits. The fractions $n_{h,t}$ are then updated using the following formula:

$$n_{h,t} = \frac{exp(\beta A_{h,t-1})}{Z_{t-1}}, \quad Z_{t-1} = \sum_{h=1}^{H} exp(\beta A_{h,t-1}), \tag{2.18}$$

where β is called the intensity of choice and Z_t denotes a normalization factor.

Note that the ABS can be applied to a finite number of trader types, with the fundamentalist vs chartist model being just a particular example of an ABS with 2 types of traders. Extensions to this framework were made by considering an infinite number of agent types or strategies. To this end, Brock et al. (2005) introduced the large type limit (LTL) theory describing the dynamic behaviour of heterogeneous markets with many trader types, whereas Diks and Van Der Weide (2005) proposed a continuous belief system (CBS). Both theories are constructed on the idea of distribution of beliefs (Chen et al., 2012). This can be viewed as a distribution over a belief space, from which the observed beliefs are sampled. In the finite example, we are dealing with a sample of beliefs, with the switching mechanism operating on this fixed sample. On the other hand, in the infinite case we consider the population of beliefs, with the discrete choice problem being expanded to continuous choice.

As is usually the case for early original models, their formal estimations were performed afterwards. Boswijk et al. (2007) estimated an ABS model of fundamentalists vs. chartists using yearly Standard & Poor's 500 stock market index (S&P 500) data from 1871-2003. Their findings show a coexistence of both beliefs, with the market fraction switching from periods of fundamental domination to chartism domination as in the empirical results of Winker et al. (2001) and Gilli and Winker (2003). Their findings provided the foundation of the Market Fraction Hypothesis (MFH) which says that the financial markets can be described by agents switching between different types or beliefs (Chen et al., 2012). This is one of the fundamental findings of agent-based financial models. In addition, Boswijk et al. (2007) found that fundamentalists and chartists coexist and co-evolve in the market, with their market fractions changing constantly. However, in Chapter 4, we will see that this is not always the case. In addressing one of our main research aims, that is, testing the implications of various behavioural biases on the model's dynamics (see Section 1.2 for more details), we demonstrate how different heuristics in agents' behaviour can lead to their complete disappearance from the market (see Section 4.3.2).

Further estimations of the two- or three-type ABS models are made by Amilon (2008). In the three-type model, the author uses a framework similar to the one used by Lux, with two types of chartists. He incorporates a wide range of parameters, leading to a 9-parameter two-type ABS model and a 14-parameter three-type ABS model. The two models are associated with different structures and estimated using different methods. While in the two-type model the maximum likelihood method is used, the three-type model is estimated using the method of simulated moments. The interesting finding is that the inclusion of a second type of chartists makes is hard for fundamentalists to survive, meaning that a fundamentalists-chartists model is sufficient for agent-based modelling purposes.

The agent-based models incorporating stochastic interactions between market participants offer a very good representation of how real financial traders behave. The agents are allowed to switch strategies according to a series of factors such as herding mechanisms or past performance. In our systems, presented in the next three chapters, we model stochastic interactions between market participants and base their switching between strategies on both the discrete choice approach and the transition probability approach. In the next section we discuss further factors that may influence agents' behaviours and strategies, based on observations from real financial markets. This way, we strive to improve the degree to which agents represent investors' behaviour, while also testing the implications of this behaviour on the overall market dynamics (see Section 1.2 for a more detailed discussion on our research aims and objectives).

2.4 Behavioural biases

In the previous sections we have seen how agent-based representations of financial markets have developed from simple deterministic models incorporating no interactions between agents, to more complex settings where agents are allowed to interact stochastically and switch strategies. In this section we discuss in more detail some of the factors that may influence agents' actions given the observed state of the market. By incorporating such factors in our agent-based models defined in next chapters, we address two different aims of our research (see Section 1.2). In this context, not only do we offer a better representation of investors' behaviour but we also test the implications of different behavioural biases on the overall model setting.

More specifically, in this section we summarize some of the developments of an expanding line of research, known as behavioural finance. This comes as an alternative to most of the conventional financial theories assuming that new information is immediately reflected in the asset prices and it can be handled by rational decision-making agents. However, a few major implications, such as optimisation problems too hard to solve and biased beliefs, make traditional economics models simply too difficult to implement. In reality, market participants only have limited information and restricted processing capabilities. Kahneman and Tversky (1979) introduced prospect theory which suggests incorporating psychological factors when modelling the decision-making of economic agents. There is a good possibility that investors have biases and deviate from rationality, which, in turn, will have a large impact on the prices. Moreover, these biases may generate predictable errors in individuals' judgement and choices (Shefrin, 2008; Kahneman, 2011).

The work of Kahneman and Tversky (1979); Shiller (2003) and Thaler (2000) among others, takes forward the arguments of rational economic agents of neoclassical economics to quasi-rational human beings. In more detail, behavioural finance focuses on identifying various biases and heuristics that can potentially affect financial decision-making. There are a wide range of factors considered, from situational or personality-related (Holden, 2010) to entirely non-economic ones, associated to weather or other external events (Hirshleifer and Shumway, 2003). Human decisions are the product of emotion, habit, reason and instinct and how these factors interact with each other depends on time, place and circumstances (Brennan and Lo, 2011). Behavioural finance has been rapidly expanding and many researchers have studied the investors' biases in their information processing techniques. The main findings of behavioural finance are described in the work of Barberis and Thaler (2003), De Bondt (2005), Thaler (2005), and Shefrin (2002). More recently, significant focus has shifted from the econometric analysis of financial data to developing models of human psychology and its relations to financial markets (Shiller, 2015).

Notably, psychological experiments show that people's beliefs are often predictable. In many cases, the source of the problem is cognitive and therefore a function of how people think. An extensive list of cognitive biases and their description can be found in De Bondt et al. (2008). Some of the most compelling behavioural concepts include representativeness (the tendency to assign to certain information more importance than to another, without any basis), overconfidence (the agents' belief that they know more than they actually do), excessive optimism (exaggeration of one's own abilities), anchoring (tendency to stick to a reference point), availability heuristic (tendency to form estimates of judgements based on ease of instances coming to mind rapidly), self-attribution error (tendency to find information that confirms an existing belief) and overreaction bias (tendency to react sharply). All of these behavioural elements play a role in the decision making of economic agents in general and stock market participants in particular, thus affecting the asset pricing mechanism (Chandra and Thenmozhi, 2017).

A further key element of financial models is represented by people's preferences. Some of the most well-known features of behavioural preferences include loss aversion (tendency to value losses more than gains of equal amount), mental accounting (tendency to think of value in relative rather than absolute terms) and regret aversion (motion of pain and anger felt by agents when they observe that they took a bad decision in the past and could have taken one with better outcomes).

We believe that agent-based modelling and behavioural finance complement each other and can be used together as the agent-based framework can serve as a useful theoretical tool for verifying the findings from behavioural finance. LeBaron (2006) also argues that agent-based models are well suited for testing behavioural theories and anticipates that the connection between them will become more intertwined as both fields progress.

In this context, in the following sections we discuss in more detail two particularly interesting theories of decision-making under uncertainty, namely prospect theory and regret theory. Both of the approaches have been widely used in the literature and demonstrated to have major implications for understanding investors' preferences. Moreover, they are capable of explaining empirical observations that cannot be predicted by traditional theory. In Sections 2.4.1 and 2.4.2 we discuss prospect theory and its related biases, namely loss aversion and the disposition effect, followed by a review of regret theory in Section 2.4.3. In Chapters 4 and 5, we will see how these behavioural biases affect the market dynamics in agent-based settings.

2.4.1 Loss aversion

We now turn our attention to the trade-off between risk and return, a crucial aspect of financial decisions. To this end, there are several behaviourally based preference frameworks. In particular, prospect theory, developed by Tversky and Kahneman (1974; 1992), attempts to describe how people think and the way they systematically violate the axioms of expected utility. Importantly, prospect theory and its successors state that people are risk averse, risk neutral and risk seeking, depending on their personality and the situation in which they find themselves (De Bondt et al., 2013).

In general, people want safety and they worry about the downside risk, but others are more focused on the upside potential of risky opportunities. Nevertheless, the target outcome or the reference points play a key role in decision-making. Therefore, people are worried about losses or below target results and at the same time tend to avoid danger if they can achieve their goals. This explains the risk seeking by otherwise cautious individuals. According to prospect theory, agents are risk-averse in the domain of gains but risk-seeking when the changes in wealth are perceived as losses. This is demonstrated in the following experiment, where a subject has to choose between two pairs (A vs. B and C vs. D):

- A = sure gain of \$24,000
- B = 25% chance to gain \$100,000 and a 75% chance to gain nothing.
- C = sure loss of \$75,000
- D = 75% chance to lose \$100,000 and 25% chance to lose nothing.

Kahneman and Tversky found that more people chose A than B and more people chose D than C. While the choice of A over B is consistent with risk aversion, the choice of D over C is considered risk seeking. Note that the \$25,000 expected gain of B (25% of a \$100,000 gain), is greater than the sure \$24,000 gain of A. Hence, the common choice of A over B is consistent with risk aversion. On the other hand, the expected loss of \$75,000 of D (75% of a \$100,000 loss) is equal to the sure \$75,000 loss of C, but riskier since it can impose a \$100,000 loss. Kahneman and Tversky call the choice of D over C as aversion to a sure loss, since C imposes a sure loss and D doesn't. Therefore, someone who is not satisfied with his losses is likely to accept gambles that would be unacceptable otherwise (Han and Hsu, 2004).

In general, an individual is loss averse if she or he dislikes symmetric 50-50 bets and, furthermore, the preference increases with the size of the stakes. As a result, there is an asymmetry in how an agent perceives gains and losses of the same amount. In more detail, the loss aversion refers to investors' reluctance to realise losses, i.e. losses loom larger than gains (Kahneman and Tversky, 1984; Tversky and Kahneman, 1991). This concept applies when one is trying to avoid a loss even if it means accepting a higher risk.

Tversky and Kahneman (1992) argue that people weight losses twice as much as gains of a similar magnitude. Therefore, investors may prefer to avoid risk to protect their existing wealth, but also take risks in order to avoid certain losses. In general, the form of a value function $v(x)$ is different depending on whether the change from the reference point x is in the range of profits or losses. The function is concave for gains, convex for losses and has a gradient in the range of loss twice as steep as the one in the range of profit. Specifically, Tversky and Kahneman (1992) define the value function as a two part power function:

$$v(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x)^{\beta} & \text{if } x < 0, \end{cases} \tag{2.19}$$

where the median exponents α and β were found to be 0.88 for both gains and losses, and the median λ was 2.25, indicating loss aversion.

Bracke and Tenreyro (2016) find evidence that supports the loss-aversion hypothesis in real world housing data by showing that sellers demand a higher price if they bought the house when prices were high. This effect is also uniform across those with or without mortgages. Specifically, Tversky and Kahneman studied loss aversion in experiments on human subjects and their hypothesis has been supported by real world data.

However, even though these biases have been clearly observed in real life investors' behaviour or identified in laboratory experiments, examining their implications on the overall market is still a challenging topic. Because of the high complexity of financial markets, it is almost impossible to assess the macro effects of a particularly interesting micro dynamic. For this reason, it is hard to directly quantify the impact of various behavioural heuristics on the interaction between market participants or on the observed time series. Some famous counterexamples include the work of Barberis et al. (1999) and Benartzi and Thaler (1993), who proposed equilibrium behavioural models based on prospect theory and showed that the loss aversion feature was successful in explaining some of the real life financial properties such as the premium puzzle and the volatility puzzle.

Despite their promise, these behavioural financial models have been able to address only a few statistical features among financial stylized facts. In this context, Shimokawa et al.

(2007) proposes a prospect theoretical equilibrium model that incorporates loss-averse investors. In a market populated by smart fundamentalists and loss averse traders, the biased agents forecast the value of an asset by indirectly observing the others' private signals. However, the information they can utilise is only past price sequence and it is assumed that they place excessive importance on expected losses. In order to account for loss aversion, the authors consider that the agents' optimal demand is given by a modified utility function such that:

$$\begin{aligned} z_{h,t} &= \max_{z_{h,t}} E[U_h(W_{h,t+1})], \\ U_h(W_{h,t+1}) &= -\exp[-a_h B_{h,t} W_{h,t+1}], \end{aligned} \tag{2.20}$$

where $B_{h,t}$ is a parameter related to the loss aversion such as:

$$\begin{cases} B_{h,t} & = 4 \quad if E[v - p_{t-1}^R] < 0, \\ B_{h,t} & = 1 \quad if E[v - p_{t-1}^R] \geq 0, \end{cases} \tag{2.21}$$

where v is the settlement price or fundamental value of the risky asset and p_{t-1}^R is a reference point determined by averaging the last R periods' prices. This means that investors decide their demands by considering their expectations about future gains and losses. If the prefiguration related to the asset price in the current period is below the reference point, suggesting a loss, they make the loss-averse parameter increase.

With this modification of the utility function, in an agent-based modelling approach, Shimokawa et al. (2007) show that incorporating loss aversion in agents' preferences helps to explain a good number of financial stylized facts including excess kurtosis, asymmetry of return distribution, auto-correlation of return volatility and cross-correlation between return volatility and trading volume.

In a different agent-based setting, in order to account for traders' heterogeneity and irrationality, Rekik et al. (2014) consider a market with three types of traders, fundamentalists, non-fundamentalists and loss averse investors. Using the same stylized utility function as Shimokawa et al. (2007) to model loss aversion and employing agents' learning through artificial neural networks, Rekik et al. (2014) demonstrate that heterogeneous agents, that are affected by behavioural biases, help explain the dynamics of the market prices. Specifically, the authors show that prices agree the most with their underlying fundamental values in a market consisting of a combination of fundamentalists and irrational traders with various behavioural biases such as loss aversion, herding and anchoring.

More recently, Polach and Kukacka (2017) use the same general idea as Shimokawa et al. (2007) to model the most important features of prospect theory into the popular adaptive beliefs system of Brock and Hommes (see Section 2.3.2). Accounting for loss aversion with reference point dependence and distorted treatment of gains and losses, Polach and Kukacka (2017) demonstrate that the original model can be consistently extended with the most relevant features of prospect theory, while its intrinsic stylized structure may remain intact. Similar to our findings discussed in the next chapters, the occurrence of fundamentalists is more extreme and the overall stability of the model is increased when behavioural biases are incorporated in agents' strategies. The presence of the loss aversion feature increases the chances of fundamental traders surviving in the market compared to the benchmark simulation.

Selim et al. (2015) provide further evidence that incorporating behavioural biases in agents' beliefs improves the stylized facts generated by an artificial market model. In a fundamentalists vs. chartists setting similar to the one we will present in Chapter 4, Selim et al. model chartism loss aversion by considering a value function that recognises losses more than gains. In more detail, the agents' adaptive behaviour is modelled via a discrete choice approach based on a fitness measure similar to wealth. However, this is modified such that chartists weigh losses twice as much as gains of a similar magnitude, causing a more rapid switch to other groups in the domain of losses. Monte Carlo simulations reveal that the model with loss aversion is capable of explaining a number of important stylized facts of stock markets, such as random walk price behaviour, bubbles and crashes, fat-tailed return distributions and volatility clustering. By modelling loss aversion in a similar way (see Section 4.1.3 for more details), we take a step forward and show that incorporating loss aversion in agents' interactions leads to a model that clearly improves the degree to which it can match the financial stylized facts, compared to the model without loss aversion, thus addressing our key research aims of a better representation of investors' behaviour and an improvement of the recreation of real life financial properties in an agent-based setting (see Section 1.2).

2.4.2 The disposition effect

A second main component of prospect theory research in finance is to understand how people trade financial assets over time. One particularly interesting empirical finding is the phenomenon that investors appear reluctant to realise losses. That is, investors have a tendency to sell winning stocks too early and hold losing stocks for too long. This pattern, labelled as the disposition effect, was first documented by Shefrin and Statman (1985) and explained with a combination of mental accounting and loss aversion biases coming from prospect theory.

For example, assume that an investor bought a stock at \$50 and is now trading at \$40. Suppose that, in the next month, the price could go either up or down \$10, with equal probability. The investor has to choose between selling the stock now and realising a \$10 loss, or keeping the stock and facing a 50-50 chance of either losing \$20 or breaking even. A risk averse investor will sell the stock. However, according to prospect theory, a risk seeking investor in the domain of losses, who sets the purchase price as the reference point to compute gains and losses, will not sell the stock. The prospect theory investor prefers the risk of breaking even to the certain pain of realising a loss. Conversely, suppose that the investor purchased one share at $50 and the trading price is now $60, with a 50-50 chance of going up or down $10. In this case, a prospect theory investor finds himself in the domain of gains and will sell the stock, preferring the immediate realisation of the $10 gain.

Over the years, with transaction data becoming readily accessible, the disposition effect has become a widely documented empirical regularity. In one of the first such studies, Odean (1998) analyses the trading records of a large discount brokerage house between 1987 and 1993 and finds a strong preference of investors for realising winning trades rather than losers. Moreover, it is shown that some of the most obvious potential explanations, based on informed trading, rebalancing or transaction costs, fail to capture important features of the data. Motivated by these findings, various other studies report the existence of the disposition effect in mutual funds data (Frazzini, 2006), the Finnish stock market (Grinblatt and Keloharju, 2001), the Taiwanese stock market (Barber et al., 2007), the Chinese stock market (Chen et al., 2007), the company stock options market (Heath et al., 1999), and the housing market (Genesove and Mayer, 2001).

Although evidence supports the idea that greater investor sophistication is associated with less susceptibility to the disposition effect (Frazzini, 2006), professional traders still suffer from this behavioural bias. To this end, Locke and Mann (2000) analyse the patterns of professional futures traders and find that while all traders tend to hold losers longer than winners, some of the least successful ones hold losers the longest. On the other hand, the most successful traders hold losers the shortest time. Moreover, Shapira and Venezia (2001) find evidence of the disposition effect among professional traders in Israel, while Jordan and Diltz (2004) examine day trader transactions for evidence of the disposition effect.

In addition, this behaviour is particularly puzzling because stock returns usually exhibit momentum patterns. That is, stocks that have recently done well continue to outperform, on average, while those that have done poorly continue to do so (Frazzini, 2006). For this reason, investors should sell the stocks with poor past performance, but instead they do the opposite. This apparent unwillingness to realise losses comes naturally from prospect theory and the convexity of the value function in the region of losses (Shefrin and Statman, 1985).

A number of papers have tried to formalise this intuition, but, as often happens with the quantification of behavioural biases, the task turns out to be harder than expected. In particular, the argument is that, for the intuition to work, the investors' value function needs to be much more convex over losses than the experimental evidence suggests that it actually is. While this issue continues to be debated (Barberis and Xiong, 2009), other authors suggest that the disposition effect can be better understood as a consequence of the realisation utility. That is, the idea that people derive utility directly from selling at a gain relative to the purchase price, and disutility from selling at a loss. In other words, people think that selling at a gain is a good recipe for long term wealth accumulation or, conversely, selling at a loss is a poor recipe for wealth accumulation. Note that even though this explanation for the disposition effect is different than the one based on the convexity of the prospect theory value function, it is still rooted in the idea that investors derive utility from gains and losses rather than absolute wealth levels.

In one of the few agent-based examples implementing the disposition effect, Li et al. (2014) use a similar model to the one we will present in Chapter 4 to account for this behavioural bias. In a fundamentalists vs. chartist setting, the authors build on the model of Westerhoff (2008) to account for the disposition effect. Specifically, the investors' disposition effect is introduced into the model changing the chartists' trend following demand function such that:

$$z_t^c = \begin{cases} 0 \quad if \quad p_t - p_{t-1} < f, \\ b(p_t - p_{t-1}) + \beta \quad otherwise, \end{cases} \tag{2.22}$$

where $f < 0$ is a fixed threshold, p_t represents the price at time t, $b > 0$ is a chartists reaction coefficient to the price trend and $\beta_t \sim N(0, \sigma^\beta)$ captures the noise or demand heterogeneity. Accordingly, if the price falls more than a fixed value $|f|$, chartists continue to hold their assets and have a demand equal to 0, thus exhibiting the disposition effect. Otherwise, chartists go with their usual strategy and follow the latest trend. With this modification in the chartists' demand functions, Li et al. (2014) show that the disposition effect can reproduce the widely observed property of asymmetric volatility. That is, the impact of bad news is greater than the impact of good news. Furthermore, the authors argue that investors disposition behaviour slows the release rate of bad news, which affects the asset price fluctuations and reduces the deviations between the price and its fundamental value.

Evidence suggests that investors' unwillingness to cut their losses is closely related to the notion of regret, which may be accentuated by having to admit their mistakes to other people (Kahneman and Tversky, 1982). Put differently, investors may not realise a loss because they are trying to avoid the regret associated with their initial (wrong) judgement.

2.4.3 Regret

Regret is a cognitive emotion of pain and anger felt by agents when they observe that they took a bad decision in the past and could have taken one with better outcomes.

In financial markets, regret is experienced when an investor's decision yields, ex-post, a worse performance than an alternative one he could have chosen (Bleichrodt and Wakker, 2015). In general, there are two components leading to regret, one associated with the comparison between alternative outcomes and the other with the feeling of self-blame for having made a poor choice.

Consider a general choice situation between two actions A and B. Regret theory generalises expected utility by assuming that the utility experienced under choice A is affected by what would have happened, had a different choice B been chosen instead of A, and vice versa. Agents feel regret about A if the result of the alternative choice B had been better.

In more detail, an economic theory of regret was independently formulated by Bell (1982) and Loomes and Sugden (1982). This theory of choice under uncertainty explains some of the observed violations of the traditional expected utility approach by assuming that economic agents are rational and not only base their decisions on future expected payoffs but also on expected regret. Therefore, investors are assumed to make decisions by maximising the expected value of a modified utility function that takes into account both traditional risk and regret.

The causes and effects of regret have been extensively discussed in psychology literature, which shows that regret is more intense when the unfavourable outcomes are the result of action rather than inaction (Kahneman and Tversky, 1982). Moreover, the effects of this behavioural bias are amplified by the visibility of the outcomes of alternative options. In this context, the degree of regret seems to be proportional with the closeness to unchosen options (Kahneman and Miller, 1986). Ritov and Baron (1995) show that anticipated regret was greater when subjects knew they would have complete information about outcomes, rather than the situations when information was available just for the chosen outcome. In other words, the anticipated pain of regret can be reduced or entirely eliminated if people do not know the outcomes of the forgone alternatives. However, in the case of financial markets, investors will always have information about alternative choices.

In addition, regret is such a strong emotion that its negative implications can influence an agent's decision making under uncertainty (Connolly and Zeelenberg, 2002). The prospect of future regret may determine investors making sub-optimal, non-rational decisions relative to their utility models. Therefore, since the anticipation of future regret may change the agents' decisions, their satisfaction with investments is not simply a function of outcome (Fogel and Berry, 2006). In general, it is assumed that people maximise a slightly modified utility function that builds on traditional utility theory in a natural way. That is because an individual who experiences regret is expected to try to anticipate the feeling and take it into consideration when making a decision under uncertainty (Loomes and Sugden, 1982).

Originally, regret was introduced as part of a modified utility function such that:

$$u(x, y) = v(x) + f(v(x) - v(y)), \tag{2.23}$$

where $u(x, y)$ is the utility of achieving x, knowing that y could have been achieved and $v(x)$ is the traditional utility function. The difference $v(x) - v(y)$ is the loss/gain value of having chosen x instead of y and $f(v(x) - v(y))$ indicates the regret associated with having chosen x instead of y. Although this function has been initially defined for pair-wise choices, it can be easily extended to more general choice sets. Consider an agent that can choose among various investments i, with outcomes x_i. A modified utility of choosing i is given by:

$$u(x_i) = v(x_i) + f(v(x_i) - v(\max[x_i])), \tag{2.24}$$

where $\max[x_i]$ is the best ex-post outcome among the feasible alternatives. Note that investors now choose the optimal investment portfolio by maximizing their expected modified utility of all possible investment choices.

Intuitively, while traditional utility theory of choice is defined over the agent's portfolio, this modified utility includes a comparison with other portfolios that could have been chosen but are not owned by the agent. Contrary to the prospect theory of Kahneman and Tversky (1979) and the disappointment theory of Gul (1991), regret theory does not introduce an increased risk aversion but rather two dimensions of risk. The first one addresses the deviation of the portfolio return from its expected value, while the second one is linked to deviations from best past alternatives.

Now, although the experimental psychology literature supports the assumption that regret influences decision-making under uncertainty, it has largely been ignored in the finance literature. Some counterexamples include Barberis et al. (2006) using regret theory to explain the stock market participation puzzle, Braun and Muermann (2004) applying it to derive the optimal demand for insurance, Muermann et al. (2006) using it for asset allocation in pension schemes, Wong (2015) using it to model optimal capital structure, Dodonova and Khoroshilov (2005) including a simplified form of regret in asset pricing and showing it can explain the high volatility of stock returns, Michenaud and Solnik (2008) deriving closed form solutions for optimal investment choices in currency exposure decision using a utility function with regret, and more recently Simões et al. (2018) minimising regret as a constraint in portfolio theory. Nevertheless, none of this work models regret in an agent-based setting where the focus is on the implications of particularly interesting microlevel factors on the macro dynamics. In contrast, the focus is on deriving comparative statics and quantifying the influence of different levels of regret aversion. In many cases this limitation is brought about by the technical difficulty of optimising an expected utility function with two attributes: value and regret.

Apart from cognition and emotion, the behavioural literature also studies social psychology. As we have seen in Section 2.3, social interactions are an important factor in financial markets and researchers try to capture them in the agent-based models. For a better representation of how investors and traders behave, which is one of our research objectives (see Section 1.2), the agent-based financial models should capture the behavioural factors observed in real markets. The agents try to model real life traders and therefore must represent their behavioural biases to some extent. Unfortunately, there are only few models in the literature that address this issue.

In our models, presented in Chapters 4 and 5, we address this major objective of agent-based modelling by building on existing fundamentalist vs. chartist models to test the implications of some of the most discussed behavioural preferences of real market participants, loss aversion and the disposition effect. While still achieving one of the main goals in the field, i.e. matching a rich set of stylized facts (as discussed in Chapter 1), the new models will display new interesting dynamics. Moreover, in Chapter 5 we define a new agent-based model that accounts for regret. Accordingly, we model investors' decisions such that they are influenced by the comparison with other market participants, as this is a more realistic portrait of the real world. With a focus on the mechanism responsible for generating important stylized facts (one of our research aim outlined in Section 1.2), we will discuss how regret affects the inner workings of the model and indirectly impact that interesting mechanism (see Section 5.3).

2.5 Stylized facts

In this section we discuss in more detail one of the main goals of this work, that is matching a rich set of the properties of real financial markets (see Section 1.2), each of which will be presented in this section. For a long time researchers have been interested in the nature of economic time series and in the 1990s some of the first surveys appeared on the observed statistical properties of financial data (De Vries and Leuven, 1994; Pagan, 1996; Granger and Ding, 1995).

One may think that, since different assets are not influenced by the same events or information, their price series will exhibit different properties. Nevertheless, empirical studies have found that financial time series are characterised by a set of statistical properties that seem to be common to a wide variety of financial markets, instruments and even periods (Cont, 2001). These properties, better known as stylized facts, have puzzled researchers for decades. Even though the statistical properties are well known, most of the attempts at their explanation have remained unsuccessful. In particular, it is still unclear what are the economical or financial mechanisms that cause the appearance of these properties. Interestingly, the stylized facts are very similar to the scaling laws (Stanley et al., 1996) in natural sciences which can be viewed as the results of microscopic interactions in complex systems. This offers a strong motivation for modelling financial markets using multi-agent interactions (Lux, 2008). As discussed in Chapter 1, agent-based models try to explain the mechanisms that generate the observed behaviour of market prices using simple, interacting, heterogeneous agents.

Below we discuss some of the commonly observed stylized facts of financial time series including the martingale property and lack of predictability, absence of autocorrelations, heavy tails of returns, volatility clustering, long-memory and long range dependencies, return-volume relations, aggregate gaussianity and the price impact. These properties will be reproduced by our agent-based models presented in Chapters 3 and 4 (see Sections 3.3 and 4.5).

2.5.1 Lack of predictability and non-stationarity

The martingale property is one of the fundamental properties of price series, being at the heart of the Efficient Market Hypothesis. It is a fundamental ingredient of theoretical models and a wide range of literature tries to explain it. The martingale property states that knowledge of past events never helps predicting the future movements. As a consequence, the realized price changes are random variables driven by news arrival processes.

The lack of predictability of price series is well known and documented, with a well-established body of literature offering plausible generic explanations of this stylized fact (Lux, 2008). It is explained by traditional finance and the EMH as a consequence of the informational efficiency, that is, all currently available relevant information is already embedded in the market price. Hence, the price can be changed only by the arrival of new information. The predictability of price changes has been questioned in many papers, artificial intelligence algorithms and data-mining techniques being used to search for hidden patterns in price series (see Taylor (2011) for a review on artificial intelligence algorithms in financial markets).

A similar, very common property of financial time series data is its non-stationarity. In particular, a non-stationary process is a stochastic process whose joint probability distribution changes when shifted in time. Consequently, its mean and variance change over time. Usually, stock prices are defined as examples of random walks, a non-stationary process. As a rule, non-stationary data is unpredictable and cannot be forecasted. Indeed, the lack of predictability of the future price is in accordance with informational efficiency of the market.

In more detail, the non-stationarity of financial data has been discussed and studied for a long time and, similarly to the martingale property, arises from the theory of efficient markets (Pagan, 1996). A property associated with non-stationarity is the unit root, which states that one is not able to reject the hypothesis that the time series follows a random walk or a martingale process. Alternatively, the series is said to be integrated of order one. A huge body of literature has been spawned on the question of unit roots, resulting in a wide range of statistical tests that can be applied to financial data. An example of such tests is Augmented Dickey-Fuller (ADF) (Dickey and Fuller, 1979), which will be used in Chapter 3 to show that we can improve the non-stationarity of our model's price series by imposing a change in the motion of the asset's fundamental value (Section 3.1.3 for more details).

2.5.2 Absence of autocorrelations

It is a well-known fact that price movements do not exhibit any significant autocorrelation (Cont, 2001). The autocorrelation function (ACF) defined as:

$$C(\tau) = corr(r(t, \Delta t), r(t + \tau, \Delta t)), \tag{2.25}$$

where *corr* denotes the sample correlation, rapidly decays and even for small interval times is close to zero.

In early literature, Working (1934) cited King et al. (1930) as one of the first to note the fact that stock price changes appear to be uncorrelated. Furthermore, Kendall and Hill (1953) provide one of the first in-depth analyses of stock price series. They find very small (usually positive) autocorrelations in the weekly return series of 19 British stock indices between 1928-1938. The highest measured autocorrelation coefficient is below 0.24, implying a predictability of less than 6%. Similar findings were reported by Fama (1965), who studied the behaviour of daily and weekly returns for individual stock between 1957-1962. He found small, usually positive autocorrelations (below 0.1) for daily returns and even smaller, predominantly negative autocorrelations for weekly returns (usually above -0.05).

Many other studies have found a rapid decline of autocorrelations after the first lag, therefore confirming the absence of (linear) autocorrelation in returns at all horizons and making it a well accepted stylized fact. Absence of autocorrelations in returns is often cited as support for the efficient market hypothesis. Intuitively, if the price changes exhibit any significant correlation, one may create strategies to exploit it. These strategies, also known as statistical arbitrage (Bondarenko, 2003), will therefore reduce the autocorrelations except for very short time scales representing the time the market needs to react to new information.

More recently, Bouchaud and Potters (2003) report that between 1991-1995, for GBP/USD currency negative autocorrelations persisted for timescales up to 20-30 minutes but this is no longer the case. Moreover, Chakraborti et al. (2011) found that autocorrelations do not persist in timescales over 1 minute in the Euronext market between 2007-2008, while

Cont et al. (2014) report no significant autocorrelation for timescales over 20 seconds in the NYSE during 2010.

2.5.3 Excess kurtosis and heavy tails

Kurtosis is a descriptor of the shape of the probability distribution of a random variable, measuring its peakedness. Returns of stock prices, like returns of many other financial assets, are bell shaped similar to the normal distribution, but contain more mass in the peak and the tail than the Gaussian distribution. Such distributions have excess positive kurtosis, being called leptokurtic. The kurtosis coefficient, $k(X)$ of a random variable X is defined as:

$$k(X) = \frac{E[X - E(X)]^4}{\{E[X - E(X)]^2\}^2}, \tag{2.26}$$

where $E[X]$ denotes the expected value of X. The normal distribution has a benchmark kurtosis of 3, hence we will use the term excessive kurtosis defined as $k-3$ for leptokurtic random variables. An elegant way of visualizing the leptokurtic behaviour is represented by an elongated S in the so called PP-plots (see e.g. Figure 3.6), which compare the empirical cumulative distribution function with the cumulative distribution function of a normal distribution, having the same mean and variance as the empirical distribution. This is a quite useful way of qualitatively comparing the normality of the data without any further assumptions. If the data was drawn from a normal distribution, we would expect a linear plot along the 45 degree line. Any deviations from this are evidence of non-normality. In the case of heavy tails, the central part of the distribution is aligned with the normal distribution, while the greatest deviations occur at the tails in each case.

Some of the first to discover the leptokurtic behaviour of returns were Fama (1965), who reported it for each constituent of the Dow Jones Industrial Average stock index, and Mandelbrot et al. (1963), who referenced it for other financial time series.

The positive kurtosis implies a peakiness of the distribution bigger than normal and a slow asymptotic decay of the probability distribution function. This non-normal decay is often termed as a fat tail. Fat tails are defined as tails of the distribution that have a higher density than what is predicted under normality assumptions (LeBaron and Samanta, 2005). For example, a distribution with exponential decay (as in the normal) is considered thin tailed, while a power decay of the density function is considered a fat tail distribution.

The tail index is a key concept of extreme value theory allowing for a classification of the extreme values (maxima and minima) of iid random variables with continuous distributions. In more detail, let $M_n = \max\{x_1, x_2, ..., x_n\}$ be the maximum of n sample observations of the iid random variables $\{X_1, X_2, ..., X_n\}$. For such settings, Fisher and Tippett (1928) showed that, after an appropriate change of scale and location, the limiting distribution of M belongs to one of only three classes of distribution functions.

These distribution can be mapped onto different domains of attraction:

1. The tail can decay exponentially; these are the standard cases of normal, log normal or gamma distributions.
2. It can decay by a power, as in the cases of Pareto or Cauchy distributions. These are no longer integrable when weighted by the tail probabilities, leading to fat tails.
3. The tail can decay with a finite tail index; these will be thin tailed.

Formally, the tail probabilities $\bar{F} = P(X > x)$ of a random variable X whose maxima and minima are described by one of the three types, are related to its distribution $G(x)$ through the relation:

$$\bar{F} = -\ln G(x), \quad \ln G(x) > -1 \tag{2.27}$$

This implies the following tail probabilities for the random variables X:

1. Type I: Medium-tailed:
$$\bar{F} = exp(-x)\mathbb{I}_{x\geq 0}, \tag{2.28}$$
2. Type II: Fat-tailed:
$$\bar{F} = x^{-\alpha}\mathbb{I}_{x\geq 1}, \tag{2.29}$$
3. Type III: Thin-tailed:
$$\bar{F} = (-x)^{\alpha}\mathbb{I}_{-1\leq x\leq 0}. \tag{2.30}$$

Therefore, extreme value theory tells us that even without knowing the distribution of a random variable X, we can still derive certain limiting results of the distribution of its maxima. The extremes may be characterised by exponentially or hyperbolically decays for type I and II respectively, whereas distributions with extremal behaviour of type III have finite endpoints.

The tail probability of the Generalised Pareto Distribution (GDP) provides a unifying representation of 2.28-2.30:

$$\bar{F}\xi = (1 + \xi x)^{-1/\xi}, \tag{2.31}$$

where the sign of ξ determines the distribution into type I ($\xi \to 0$), type II ($\xi > 0$) and type III ($\xi < 0$). Moreover, the tail index is given by $\alpha = 1/|\xi|$.

Therefore, in order to discriminate between the types of limiting distributions we need to estimate the tail index. After all, measuring the tail index of a distribution provides a quantitative measure of the heaviness of the tail. One of the widely used methods, introduced by Hill et al. (1975), provides an estimator of the tail index ξ using the maximum likelihood method:

$$\hat{\xi} = \frac{1}{k}\sum_{i=1}^{k}\{\ln x_{n-i+1} - \ln x_{n-k}\}, \tag{2.32}$$

where x_i is the i'th order statistic of the sample and k denotes the number of the n sample observations for which the asymptotic behaviour is assumed to be valid. Intuitively, the Hill index $\hat{\xi}$ uses the step size between extreme observations to generalise the behaviour of the tails into the broader part of the distribution (LeBaron and Samanta, 2005). Hill's estimate has been shown to be a consistent estimate for the fat tailed distributions by Mason (1982). Although the Hill index is commonly used for approximating the tail index, alternative estimations have been proposed (Dekkers and De Haan, 1989).

There are many empirical studies that address this stylized fact. For example, Jansen and De Vries (1991) were some of the first to apply extreme value theory to financial data. Using the Hill estimate, they obtain a tail index in the range 3.2-5.2 for daily returns of various stock indices in 1962-1986. Loretan and Phillips (1994) report a tail index in range 3.1-3.8 for daily returns for the S&P index in the period 1962-1987 and between 2.5-3.2 for monthly stock index returns from 1834-1987. A dataset of daily US stock returns from 1985-1990 was studied by Longin (1996) at various frequencies, finding a tail index in the range 3-4. Furthermore, Lux (1996) applies the Hill estimate to daily returns of the German stock index DAX and its constituents in 1988-1989 obtaining a tail index of 2.3-3.8.

A more recent extensive study on tail indices was conducted by Jondeau and Rockinger (2003), using a database consisting of daily stock-market returns for 20 countries. They found the tail indices of emerging market share prices hover in the same range as those of major industrialised countries. Fat tails of returns were found in the American Stock Exchange by Plerou and Stanley (2008), Paris Bourse by Bouchaud et al. (2008), Euronext by Chakraborti et al. (2011), NYSE and NASDAQ by Gopikrishnan et al. (1998). The precise form of the distribution changes with the time-scales of returns. In particular, Gu et al. (2008) found that in the Chinese stock market the tail of the distribution can be approximated by a power law with exponent $\alpha \approx 3$. Notably, most data sets studied reveal a tail index higher than two and less than five (Cont, 2001).

2.5.4 Heteroscedasticity and volatility clustering

The absence of autocorrelation discussed in Section 2.5.2 does not rule out the possibility of nonlinear dependencies in returns. It is well known that the absence of serial correlation does not imply independence. Studies found that tests of independence usually reject the null hypothesis of independence in financial returns (Pagan, 1996). Even simple visual representations of return series reveal heteroscedasticity as a violation of the assumption of independently and identically distributed returns. That is, volatility measured as squared or absolute returns is not constant in time. Fluctuations of volatility were first noted by Mandelbrot et al. (1963) in daily returns of cotton prices. The author reported that periods of high volatility alternate with periods of low volatility. Further, early examples of heteroscedasticity were found in Fielitz (1971), Wichern et al. (1976) and Hsu (1977), to name a few.

Furthermore, nonlinear representations of returns, such as absolute, squared or various powers of returns, exhibit a much higher positive autocorrelation that persists over time. The powers of absolute returns can be seen as measures of volatility, indicating a high degree of predictability of volatility, although without providing any information on the direction of price movements. This phenomenon is stable across different financial instruments and time periods, being a quantitative signature of the well known volatility clustering: large price variations are more likely to be followed by large price variations (Cont, 2001).

A widely used measure of volatility clustering is the autocorrelation function (ACF) of squared returns:

$$C_2(\tau) = corr(|r(t+\tau,\Delta t)|^2, |r(t,\Delta t)|^2), \tag{2.33}$$

although any power of absolute returns can be used:

$$C_\alpha(\tau) = corr(|r(t+\tau,\Delta t)|^\alpha, |r(t,\Delta t)|^\alpha). \tag{2.34}$$

This behaviour has been first noted by Mandelbrot et al. (1963), with the first models covering it being introduced by Engle (1982) in its famous AutoRegressive Conditional Heteroscedasticity (ARCH) framework. The slow decay of a positive autocorrelation is sometime called the ARCH effect. Interestingly, Ding et al. (1993) studied the effects of different powers α and found that, for a given lag τ, the correlation is highest for absolute returns with $\alpha \approx 1$, meaning that the absolute returns are the most predictable. An extensive study on volatility clustering in financial markets can be found in Cont (2005).

2.5.5 Long memory and long range dependencies

A property closely related to the volatility clustering effect is the decay of the autocorrelation function. The ARCH models give rise to exponential decay in autocorrelations of squared or absolute returns. However, visual inspections of the ACF raise doubts over such a fast decay (see e.g. Figure 3.8). Empirical studies have found that the decay of the function $C_\alpha(\tau)$ can be reproduced by a power law:

$$C_\alpha(\tau) \sim \frac{A}{\tau^\beta}, \tag{2.35}$$

with the coefficient $\beta \leq 0.5$, usually between 0.2 and 0.4 (Liu et al. (1997), Cont et al. (1997), Cont (2005)). This slow decay (hyperbolic) is a defining property of long memory or long range dependence. Following the definition provided by Cont (2005), a stationary process Y_t (with finite variance) is said to have a long range dependence if its autocorrelation function $C_\alpha(\tau) = corr(Y_t, Y_{t+\tau})$ decays as a power of the lag τ:

$$C_\alpha(\tau) = corr(Y_t, Y_{t+\tau}) \underset{\tau\to\infty}{\sim} \frac{L(\tau)}{\tau^{1-2d}} \qquad 0 < d < 1/2, \tag{2.36}$$

where L is slowly varying at infinity, i.e. verifies $\forall a > 0$, $\frac{L(at)}{L(t)} \to 1$ as $t \to \infty$.

One of the widely used quantitative measures of the long range dependencies in time series is the Hurst exponent (Hurst, 1951), which falls in the range $1/2 < H < 1$ for a long memory process. This is the traditional statistic used for detecting the long memory of a process elegantly formalized by Lillo and Farmer (2004). By letting $\mathbf{X} = X(t_1), X(t_2), ..., X(t_k)$ be a real valued time series, we can characterise the long memory using the diffusion properties of the integrated series $\mathbf{Y}$ as $Y(l) = \sum_{i=1}^{l} X(t_i)$. Furthermore, let $V(l) = Var(Y(t_i+1), Y(t_i+2), ..., Y(t_i+l))$ for some $i \in \{0, 1, ..., kl\}$. Then, in the limit $l \to \infty$, if $\mathbf{X}$ is a short-memory process, then $V(l)$ scales as $O(l)$, whereas if $\mathbf{X}$ is a long-memory process, then $V(l)$ scales as $O(l^{2H})$, for some $H \in (0.5, 1)$. The exponent H is called the Hurst exponent. According to Cont (2001), most studies on financial time series have reported a Hurst exponent between 0.55 and 0.6.

Numerous empirical tests have been designed aiming to offer quantitative explanations for the long memory effect. Specifically, Mandelbrot (1971) was the first one to suggest this stylized fact and observed it in many empirical studies (Mandelbrot and Taqqu, 1979). Long range dependencies have been found across different markets and periods[2]. Moreover, similar to ARCH models, there are agent-based models specifically designed for generating long-memory. For example, Muzy et al. (2000) consider the correlation function 2.34 with $\alpha = 1$ and show that it exhibits a slow decay, represented in a logarithmic form. Many others have tried to show that specific functions fit the autocorrelation of volatility without finding common grounds.

2.5.6 Volume-volatility relations

The relationship between volatility and volume traded is important for understanding how information is transmitted and embedded in markets. This is at the heart of

[2] See e.g. Liu et al. (1997), Liu et al. (1999), Cont (2005), Chakraborti et al. (2011) and Cont et al. (1997)

determining the underlying mechanisms of crashes and rapid changes in the stock prices. Furthermore, an intensive analysis on volume can help investors identify periods in which informational shocks occur, providing valuable information about future price changes.

To this end, the dependence between volatility and volume traded was first observed by Osborne (1959), who noted that it changes very rapidly. The positive correlation between the two key market descriptors was first detected by Ying (1966). He compared the S&P prices and NYSE trading volume, concluding that small (large) trading volumes were usually accompanied by a decrease (increase) of prices. Moreover, a large volume increase was followed by either a large rise or fall in prices. This empirical finding has been confirmed by later studies[3]. In line with earlier studies, a positive volatility/volume dependence was found by Marsh and Wagner (2004) in an extensive study on large international equity markets including Holland, France, Germany, Hong Kong, the U.K., and Japan. In addition, the causalities of volume on price changes were studied in Chuang et al. (2009) using data from different stock indices such as S&P 500, FTSE 100, and NYSE, concluding that it exhibits a V-shape relation so that the diffusion of volatility distribution increases with volume. It is important to note that all measures of volatility are positively correlated with traded volume (Cont, 2001). Models related to this topic include the sequential information arriving model (Copeland (1976); Jennings et al. (1981)) and mixture of distributions model (Clark (1973); Tauchen and Pitts (1983)).

2.5.7 Aggregate gaussianity

The distribution of returns is probably the most studied stylized fact in the quantitative finance and econometrics literature. Apart from its heavy tails discussed in Section 2.5.3 and volatility clustering in Section 2.5.4, another important statistical property can be observed in most of the studies. As we have seen, the distribution of returns exhibits excessive kurtosis, having a higher peak than the normal distribution and a fat tail. However, as we increase the time scale over which returns are calculated, their distribution becomes more and more like the normal distribution (Cont, 2001). In particular, the shape of the return distribution does not remain the same if we change the time-scale. This can easily be seen in different markets and time periods and is referred to as aggregate gaussianity.

2.5.8 Price impact

As we have seen in Section 2.5.6, there is a positive correlation between volatility and volume. In addition, researchers have been interested in the relationship between actual prices and volume. More specifically, the focus has been on how trades impact price

[3]For a early survey see Karpoff (1987)

changes. These changes in prices that occur as a result of traders' actions are known as the price impact.

One of the first to study this relationship was Kyle (1985), who suggested that the price impact should be linear. This means that the price changes are proportional to the volume of the underlying orders. However, in the last two decades, there has been a visible shift toward a non-linear relationship between impact and volume. In particular, there is an increasing evidence of a square-root relationship between the two (Moro et al., 2009). Moreover, the impact of a single order has been widely found to be a concave function of volume.

Empirically, Gopikrishnan et al. (2000) originally observed that the volume of trades in the NYSE are asymptotically distributed. Building on their work, Gabaix et al. (2003) found that while the price change has a power law behaviour with a fat tail index ≈ 3 (in line with previous findings presented in Section 2.5.3), the volume distribution can be described by a power law with a tail ≈ 1.5. Letting Δp and V denote the price change and volume, respectively, the same authors found a relationship of the form:

$$\Delta p \simeq kV^{1/2}, \tag{2.37}$$

for some constant k, thus confirming a square-root behaviour of the price impact and its concave behaviour.

2.6 Summary

In this chapter, we have reviewed the literature related to Agent-Based Financial Modelling, exploring some of the most relevant and promising approaches in the field. The research related to this work can be split into two broad categories: the modelling itself and one of its primary objectives, the stylized facts we are trying to reproduce. These two areas are the main focus of this work (as outlined in Chapter 1) and therefore required the detailed discussion in this chapter.

In the first part of the chapter, we discussed the evolution of agent-based models from simple deterministic models with no interaction between its participants to more complex ones, based on social interactions that incorporate various non-trivial mechanisms. Thus, models of financial markets are more and more able to integrate realistic mechanisms and interactions between market participants, one of the aims of our research outlined in Section 1.2. The models discussed here serve as a base for the ones we present in the next chapters later. In particular, the ways of computing demands and the price dynamics have not encountered major modifications over the years, being similar to one ones presented in Section 2.1. Furthermore, since allowing for social interaction has improved the degree to which models are capable of matching real-life dynamics, in the settings presented in the next chapters we will allow for such market behaviour, using both the discrete choice and the transition probability approach as switching mechanisms between strategies (as per Section 2.3).

However, while in the original settings presented in this chapter agents change their beliefs according to simple utility functions, we will see how other empirically observed behavioural processes may influence traders' decisions. Indeed, since one of the key aims of our work is to test the implications of various behavioural and economic factors (see Section 1.2), we discuss some of the main findings of behavioural finance in the second part of this chapter, presented in Section 2.4. These behavioural factors will be modelled and incorporated in our agent-based settings presented in the next chapters, where we will discuss in detail their implications on the overall market dynamics.

A further research aim of this work is to develop agent-based models that can recreate the most important stylized facts observed in real financial markets (as per Section 1.2). In the final part of this chapter we provided an in-depth analysis of all the important stylized facts of real financial markets. To this end, we continue in the next two chapters by considering and building on a structural stochastic model of asset pricing, with a focus on the rich set of empirical statistical properties it is able to match and the implications of loss aversion and the disposition effect on the models' dynamics. Finally, in Chapter 5 we model regret in our agents' behaviour and focus on the mechanism responsible for generating interesting stylized facts, another important objective of our research (as discussed in Section 1.2).

Chapter 3

An Agent-Based Model of Asset Pricing and its Stylized Facts

In this chapter we extend the fundamentalists vs chartists asset pricing model of Franke and Westerhoff (the FW model) to capture new properties usually observed in real financial markets. In doing so, we address two of the main aims of our research, focusing on the structure and dynamics of the model itself and the properties it is able to recreate (see Section 1.2 for more details).

This chapter is organised as follows. In Section 3.1, we formally define the fundamentalists vs. chartists asset pricing model and its key mechanisms. Motivated by the violation of a central property of real financial price series, their non-stationarity, Section 3.1.3 summarises our first extension to the original model, resulting in an improved setting (termed the FW+ model) and representing our first research contribution (as per Section 1.3). This is followed by the estimation and validation of the model's parameters in Section 3.2, using the method of simulated moments.

Next, in Section 3.3, we present our results. We give a thorough analysis of the price series generated by the FW+ model, integrating a wide range of tests that reinforce the presence of heavy tails, volatility clustering, long memory in absolute returns, as well as absence of autocorrelation in raw returns. Moreover, we demonstrate that our FW+ setting matches a set of new properties usually observed in real financial markets including volatility-volume correlations, aggregate Gaussianity, price impact and extreme price events. To date, this is the only model that is reported to match such a rich set of the stylized facts of real financial markets, marking our second contribution. In doing so, we illustrate one of the essential aims of the agent-based financial modelling: that many of the stylized facts arise and can be explained just from the interaction of market participants (as discussed in Section 1.2).

3.1 Model definition

Heterogeneous agent-based models that rely on simple trading strategies have proven themselves very efficient in generating important dynamics of real financial markets. In this respect, the two type model of Franke and Westerhoff (2012) if one of the most successful in capturing the empirically observed traders' behaviour (Barde, 2015). For simplicity, we will refer to the original Franke and Westerhoff model as the FW model. In this simple framework, introduced in Section 1.1, the market is populated by two different types of agents. On one hand, we have the fundamentalists. These agents act as a stabilizing force in the market and base their trading strategies on the presence of a fundamental value. On the other hand, we have chartists or technical analysts. This second class of investors base their strategies on historical price movements.

In the following sections we define the structure of the FW model, starting with its basic price dynamics presented in Section 3.1.1, and its agents' interactions in Section 3.1.2. Next, motivated by the violation of a central property of real financial price series, their non-stationarity (see Section 2.5.1 for more details), we propose a novel change in the system. In Section 3.1.3 we present our first research contribution (as outlined in Section 1.3) and change the way the fundamental value is computed, modelling it as a geometric Brownian motion.

3.1.1 Price dynamics

Following some of the most prominent examples in the literature (see Section 2.1), the asset price changes are determined by excess demand. The specific demand of each trader type is kept as simple as possible; it is presented in the form of the demand per average trader. For fundamentalists, the demand is inversely related to the mispricing of the current price relatively to the fundamental value. That is, at time t, their core demand D_t^f is proportional to the gap $(p_t^f - p_t)$, where p_t is the log price of the asset at time t, while the p_t^f is the fundamental log value at time t. In the model proposed by Franke and Westerhoff, the fundamental value is kept constant. However, the unrealistic assumption of a constant fundamental value leads to stationary price series, thus violating one of the main properties of financial data. As a solution, in our extended setting, we allow the fundamental value p_t^f to change over time, following a geometric Brownian motion (see section 3.1.3 for more details). Similarly, the core demand of the chartists' group, D_t^c, is proportional to the price changes they have just observed, $(p_t - p_{t-1})$, where p_t and p_{t-1} are the log prices at time t and $t-1$, respectively.

The wide variety of within-group specifications are captured by noise terms added to each of the core demands. The noise terms are sampled at every time step for each of the trading groups. Hence, they encapsulate the within-group heterogeneity and scale with the current size of the group. Specifically, the noise is represented by two normally distributed random variables, ϵ_t^f and ϵ_t^c, for fundamentalists and chartists, respectively. In other words, one can think of the noise variables as a convenient way of capturing the heterogeneity of markets populated by hundreds or thousands of different agents. Each of the two noise terms are added to the deterministic part of the demand, leading to the core demand per agent within the corresponding group.

To summarise, combining the deterministic and stochastic elements, the net demand of each group for the asset in period t is given by:

$$D_t^f = \phi(p_t^f - p_t) + \epsilon_t^f \qquad \epsilon_t^f \sim \mathcal{N}(0, \sigma_f^2) \qquad \phi > 0,\ \epsilon_f > 0, \tag{3.1}$$

$$D_t^c = \chi(p_t - p_{t-1}) + \epsilon_t^c \qquad \epsilon_t^c \sim \mathcal{N}(0, \sigma_c^2) \qquad \chi > 0,\ \epsilon_c > 0, \tag{3.2}$$

where ϕ and χ are constants denoting the aggressiveness of the traders' demand and σ_f^2 and σ_c^2 are noise variances. Note that different choices of parameters can lead to dominations of either the deterministic part or the stochastic noise. In particular, a higher signal to noise ratio in Equation 3.1 implies a stronger mean reversion of the price towards the fundamental value, leading to a negative autocorrelation in the raw returns. On the other hand, a higher signal to noise ratio in Equation 3.2 will lead to more bubbles in the price series, implying a positive autocorrelation of returns. The exact values of the parameters are estimated in Section 3.2 using the Method of Simulated Moments (MSM) (see Section 2.3 for other models estimated using this method).

The agents are allowed to switch strategies at each iteration, so their market fraction may fluctuate over time. For simplicity, the agents' population size is fixed at $2N$. Let n_t^f and n_t^c be the proportions of fundamentalists and chartists in the market at time t, respectively. The fundamentalists' majority index is defined as:

$$x_t = (n_t^f - n_t^c)/2N \tag{3.3}$$

Therefore, $x_t \in [-1, 1]$, a value of $x_t = -1$ ($x = +1$) corresponds to a market where all traders are chartists (fundamentalists). Moreover, the market fractions can be expressed in terms of the majority index as:

$$n_t^f/2N = (1 + x_t)/2, \qquad n_t^c/2N = (1 - x_t)/2. \tag{3.4}$$

The (normalised) total demand D_t, given by the equation:

$$D_t = n_t^f D_t^f + n_t^c D_t^c = (1 + x_t)D_t^f/2 + (1 - x_t)D_t^c/2 \tag{3.5}$$

will not generally add up to zero and we will have an excess of either supply or demand. Following the early examples in the literature (see Section 2.1), a market maker is assumed to absorb the excess supply and provide any excess demand. The market maker mediates the transactions between investors and provides liquidity. He sets the price by supplying stock out of its inventory and raising the price if there is excess demand, while accumulating stock and lowering the price when there is excess supply. Specifically, the market maker reacts to the imbalance between demand and supply by proportionally adjusting the price with a constant factor $\mu > 0$, similar to the market impact function presented in Section 2.1.

Accordingly, the equation determining the price for the next period $t + 1$ results from Equations 3.1-3.5 as:

$$p_{t+1} = p_t + \frac{\mu}{2}[(1 + x_t)\phi(p_t^f - p_t) + (1 - x_t)\chi(p_t - p_{t-1}) + \epsilon_t], \tag{3.6}$$

$$\epsilon_t \sim \mathcal{N}(0, \sigma_t^2), \qquad \sigma_t^2 = [(1 + x_t)^2\sigma_f^2 + (1 - x_t)\sigma_c^2]/2. \tag{3.7}$$

The latter equation is derived as the sum of the two normal distributions from Equations 3.1 and 3.2, multiplied by the market fractions $(1 \pm x_t)/2$. The result is a new normal distribution with mean zero and variance equal to the sum of the two single variances. The combined variance σ_t depends on the variations of the market fractions of the fundamentalists and chartists. This random time-varying variance is a key feature of the model, being termed the Structural Stochastic Volatility (SSV) of returns (i.e. the log differences in prices) in the Franke and Westerhoff paper.

3.1.2 Social interactions

To complete the model, it is necessary to set up the motions of the market fractions n_t^f and n_t^c. They are predetermined in each period and change only from one period to the next one. Following the discrete choice approach (DCA) introduced by Brock and Hommes (see Section 2.3.2), the two market shares n_{t+1}^s $(s = f, c)$ can be determined using the formula $n_{t+1}^s = exp(\beta u_t^s)/[exp(\beta u_t^f) + exp(\beta u_t^c)]$, where u is the agents' utility and β is the intensity of choice. Dividing both numerator and denominator by $exp(\beta u_t^f)$, we can write the market fraction of fundamentalists as $n_{t+1}^f = 1/\{1 + exp[-\beta(u_t^f - u_t^c)]\}$.

In this basic setting, we consider the utilities u_t^c and u_t^f, usually derived from the past gains of the two groups. Thus, $(u_t^f - u_t^c)$ represents the difference in any utility variables and can be viewed as a measure of relative attractiveness of fundamentalist trading. Hence, notation may be changed to make use of the attractiveness level, a_t, defined in the FW model as the difference $(u_t^f - u_t^c)$. The discrete choice approach becomes:

$$n_{t+1}^f = \frac{1}{1 + exp(-\beta a_t)}, \quad n_{t+1}^c = 1 - n_{t+1}^f. \tag{3.8}$$

The demand terms in the price rule summarized in Equation 3.6 are normalised using a market impact $\mu = 0.01$. Furthermore, the intensity of choice is fixed at $\beta = 1$, which is common to all the DCA versions. Of course, setting these values is just a matter of scaling the market impact on prices and the relative attractiveness level a_t of fundamentalism, respectively.

Note that the market fractions are directly influenced by the attractiveness level. An increase in the index a_t leads to an increase in the market share of the fundamentalists. For this reason, it is extremely important to define the exact mechanism of the attractiveness level and all of its components.

Following the work of Franke and Westerhoff, we consider three principles that may influence the way in which the agents choose one of the two strategies. The first principle is a herding mechanism, meaning that the more agents are in a group the more attractive that group becomes. This idea has been long used in the literature (see, for example, Kirman (1993) and Lux (1995)) and can be easily represented by a term proportional to the most recent difference in the market fractions $(n_t^f - n_t^c)$.

The second principle is based on a certain predisposition towards one of the two trading strategies. This can be directly captured by a constant α_0, which is positive (negative) if the agents have a priori preference for fundamentalism (chartism). Finally, the third component encapsulates the idea of price misalignment. This is empirically backed by Menkhoff et al. (2009), who observes that when the price is further away from its fundamental value, professionals tend to anticipate its mean reversion toward equilibrium and trade more aggressively. In a setting where the current price tends to move far away from the fundamental value, chartism becomes riskier and the fundamentalism is more attractive. Hence, it is convenient to make a_t rise in proportion to the squared deviations of the price p_t from the fundamental value p_t^f. Combining the three components of the attractiveness level, we have:

$$a_t = \alpha_0 + \alpha_n(n_t^f - n_t^c) + \alpha_p(p_t - p_t^f)^2, \tag{3.9}$$

where α_0 is the predisposition parameter, $\alpha_n > 0$ captures the herding parameter and $\alpha_p > 0$ measures the influence of price misalignment.

In summary, the model is governed by two central dynamic equations. Firstly, the price adjustments are made according to Equations 3.6 and 3.7, with the structural stochastic component σ_t. Secondly, the changes of the market fractions n_t^f and n_t^c are described by Equations 3.8 and 3.9, encapsulating a herding mechanism corrected by a strong price misalignment. The whole system has a recursive structure that is easily forward iterated. Next, motivated by the violation of a crucial property of financial data, we present the our first extension of the original model, resulting in an improved setting (the FW+ model).

3.1.3 Motion of the fundamental value

Financial time series are characterised by the lack of predictability, mathematically known as the martingale property (see Section 2.5.1 for more details). It is a fundamental ingredient of theoretical models stating that knowledge of past events never helps predicting the future movements. The lack of predictability of price series is well known and documented, with a well-established body of literature offering plausible generic explanations of this stylized fact (Lux, 2008). It is explained by traditional finance as a consequence of the informational efficiency, that is, all currently available relevant information is already embedded in the market price. Hence, the price can be changed only by the arrival of new information.

A similar, very common property of financial time series data is its non-stationarity (see Section 2.5.1 for more details). In particular, a non-stationary process is a stochastic process whose probability distribution changes when shifted in time. Consequently, its mean and variance change over time. Usually, stock prices are defined as examples of random walks, a non-stationary process. It is commonly assumed that non-stationary data is unpredictable and cannot be forecasted.

In more detail, the non-stationarity of financial data has been discussed and studied for a long time and, similarly to the martingale property, arises from the theory of efficient markets (Pagan, 1996). A property associated with non-stationarity is the unit root, which states that one is not able to reject the hypothesis that the time series follows a random walk or a martingale process. A huge body of literature has been spawned on the question of unit roots, resulting in a wide range of statistical tests that can be applied to financial data (see Section 2.5.1).

One of the widely used tests of non-stationarity is the Augmented Dickey Fuller (ADF) unit root test. Following this approach, we apply the classical ADF test (allowing for a constant and trend order) to the FW model, where the fundamental value is kept constant, $p_t^f = 0$. In 10000 simulations, the unit root test was rejected 2736 times with the p-value being less than the critical value of 10%. Moreover, the test was rejected 1637 times with the p-value less than the critical value of 5%. The p-values are obtained using the updated tables from MacKinnon (2010). Therefore, we can say that the price series generated by the model using a constant fundamental value are stationary approximately 27% of the time, which strongly contradicts the behaviour of real financial time series.

The failure to pass the unit root test is mainly due to the unrealistic assumption of a constant fundamental value. As a solution, we extended the FW setting by allowing the fundamental value p_t^f to change over time. Specifically, we set p_t^f to follow a geometric Brownian motion. This is the most widely used model of describing stock price behaviour (Hull, 2006) and is usually applied in quantitative finance. Mathematically,

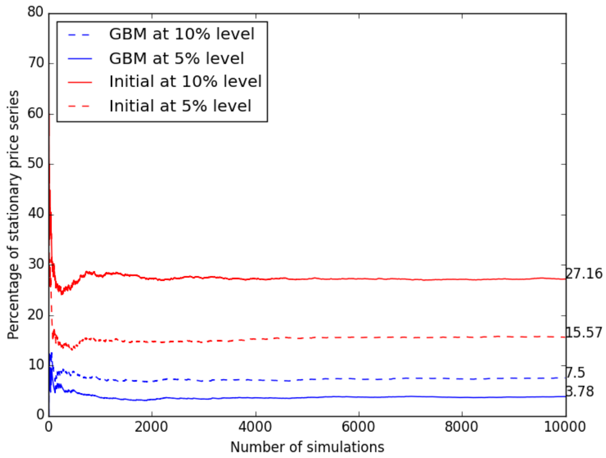

Figure 3.1: Augmented Dickey-Fuller stationarity test results for simulated FW (initial) and FW+ (GBM) data at the 5% and 10% critical values.

the fundamental price is given by

$$dp_t^f = \mu_p p_t^f dt + \sigma_p p_t^f dW_t, \tag{3.10}$$

where W_t is a Wiener process, μ_p is the percentage drift and σ_p is the percentage volatility. The exact values of the drift and volatility (μ_p, σ_p) will be estimated together with all the other model's parameters in Section 3.2.

With this change in the model, we run the Augmented Dickey Fuller test once again over 10000 different simulations. Now, the test is rejected only 823 times with the p-value less than the critical value of 10% and 418 times with the p-value less than the critical value of 5%. Thus, we can say that the price series generated by the new model (FW+) are non-stationary in more than 93% of the simulations (with 95% confidence interval). A visual representation of the number of non-stationary price series generated in both models can be seen in Figure 3.1. We plot the percentage of the non-stationary price series generated as we increase the number of simulations. The red and blue lines represent the initial and our modified model, respectively. The hard (dotted) lines are the computed ADF results at 10% (5%) critical values. We can observe how the number of stationary price series is drastically reduced, from 27% with the initial model (at 10% critical size) to 4% in our modified model (at 5% critical size).

Therefore, we can say that the new model produces a more realistic price series, since in real financial markets the prices follow martingales where knowledge of past events does not help predict the mean of the future changes. In this context, we address one of the main aims of our research (as per Section 1.2), that is, we improve the degree to which a model is able to recreate the statistics of real financial data. Unlike the original FW model, our improved FW+ setting better reflects some of the fundamental properties of real financial time series – their non-stationarity and unpredictability. This summarises our first contribution to the model. By building on the original FW model and making a novel change in the motion of the fundamental value over time, we are able to overcome the violation of a crucial property of real financial time series (see Section 2.5.1), which marks the first contribution of our work (as outlined in Section 1.3). In the next section we provide a formal estimation of the model parameters using a new method, for the first time applied to this kind of interaction between agents.

3.2 Estimation of the model's parameters

One of the biggest issues of agent-based modelling rises from its high complexity (as discussed in Section 1.2). Given this, we need to ensure that our simulations reflect some of the fundamental aspects of the real financial markets. Therefore, we have to identify and estimate all the model's parameters for which the system is both stable and significant from an economic point of view. A simple way to calibrate the model is by selecting a number of statistics for the price or return series and choosing the values of parameters such that the statistics generated by the model in a few runs are quantitatively similar to the ones of empirical data. However, a more rigorous way of estimating the numerical values of the model parameters is the Method of Simulated Moments (MSM). This technique, briefly discussed in Section 2.2, is based on an objective function that is optimised across the set of FW+ model parameters $(\phi, \chi, \sigma_f, \sigma_c, \alpha_0, \alpha_n, \alpha_p, \mu_p, \sigma_p)$. It refers to a set of statistics, also known as moments, arising from the simulations. The basic idea is that these moments should be close to their empirical counterparts, with the distance between them being captured by an objective function. For an overview of the method of simulated moments and its applications to agent-based models see for example the work of Gili and Winker (2001, 2003), Amilon (2008) and Franke (2009), to name a few.

An agent-based model is not supposed to mimic the exact economy or financial markets, but rather it should explain some of the most important stylized facts of financial markets, usually at daily frequency. In this work, the empirical data considered for model calibration consists of T=6866 daily observations of the S&P stock market from January 1980 to mid-March 2007. Following the work of Franke and Westerhoff (2011), the model was estimated such that it matches the four most discussed statistical properties of financial data including absence of autocorrelation in raw returns, heavy tails, volatility clustering and long memory. The returns, or (log) price changes, are expressed as percentage points, such as:

$$r_t = 100(p_t - p_{t-1}). \tag{3.11}$$

Moreover, the volatility of returns is defined as their absolute value $v_t = |r_t|$.

Table 3.1: The empirical moments S&P 500.

r AC-1	Hill	v	v AC					
			1	5	10	25	50	100
0.019	3.299	0.713	0.184	0.215	0.156	0.126	0.108	0.070

In order to conduct the quantitative analysis, the four stylized facts were measured using a number of summary statistics, or moments. The first moment considered is the first-order autocorrelation coefficient of raw returns. It needs to be close to zero, so that it agrees with the empirical findings on this matter described in Section 2.5.2. This will limit the chartists' price extrapolations of the most recent price changes. Furthermore, if the model generates coefficients that are insignificant, all the autocorrelations at longer time lags will vanish too. The remaining three moments are dealing with the volatility of returns. First of all, the model should suitably scale the overall volatility, thus limiting the general noise brought by the two variances σ_f and σ_c. The mean value of the absolute returns is considered. Next, the heavy tail is measured by the Hill tail index of the absolute returns. The tail is specified as the upper 5% in order to eliminate bias and consider a more accurate tail index.

The long memory effects are captured by the autocorrelation function (ACF) of the absolute returns up to a lag of 100 days. The ACF decays as we increase the lag, but it doesn't become insignificant. Hence, the entire profile has to be matched and is sufficiently well represented by the six different lag coefficients ($\tau = 1, 5, 10, 25, 50, 100$) that we consider.

Thus, the model is evaluated on the basis of nine moments (the Hill estimator, volatility, first order autocorrelation of the raw returns and the autocorrelation at lags $\tau = 1, 5, 10, 25, 50, 100$ for the absolute returns), summarised in a (column) vector $m = (m_1, ..., m_9)$. Applying the method of simulated moments, they have to come closer to their equivalent empirical moments, m^{emp}, calculated on the daily S&P 500 stock market index, whose values are presented in Table 3.1.

The distance between the two vectors m and m^{emp} is defined as a quadratic function with a suitable weighting matrix $W \in \mathbb{R}^{9x9}$ (defined shortly). Hence, the distance is given by:

$$J = J(m, m^{emp}; W) = (m - m^{emp})W(m - m^{emp})^T. \tag{3.12}$$

The weight matrix W accounts for the moments' sampling variability, its determination being crucial for the model's parameters. The idea is that the higher the sampling variability of a given moment i, the larger the differences between m_i and m_i^{emp} that can still be considered insignificant. This behaviour can be achieved by correspondingly small diagonal elements w_{ii}. Moreover, the matrix W should support possible correlations between single moments. One obvious choice for W is the inverse of an estimated

Algorithm 1 Bootstrap estimation of the covariance matrix

Set $I_0 = \{1, 2, \ldots, T\}$
for $b = 1 : B$ **do**
  Draw T random numbers with replacement from I_0
  Construct bootstrap sample $I_b = \{t_1^b, t_2^b, \ldots, t_T^b\}$
  Calculate vector of moments $m_b = (m_1^b, \ldots, m_9^b)$ from I_b
$\bar{m} = \frac{1}{B}\sum_b m_b$
$\hat{\Sigma} = \frac{1}{B}\sum_b (m_b - \bar{m})(m_b - \bar{m})'$
$W = \hat{\Sigma}^{-1}$

Algorithm 2 Parameter estimation

for $a = 1 : 1000$ **do**
  $\hat{\theta}^a =$ minimise(funJ, method=Nelder-Mead)
$\hat{\theta} = \hat{\theta}^{\tilde{a}}$, where $\tilde{a}$ such that $\hat{J}^{\tilde{a}}$= median of $\{\hat{J}^a\}_{a=1}^{1000}$ and $\hat{J}^a = J[m^a(\hat{\theta}; S), m_T^{emp}]$

def funJ(θ):
  Simulate model using vector of parameter θ
  Get simulated moments $m_{sim} = (m_1, \ldots, m_9)$
  Get $J = J(m^e mp, m_T^{emp}; W) = (m^{sim} - m_T^{emp})' W (m^{sim} - m_T^{emp})$
  Return J

variance-covariance matrix $\hat{\Sigma}$ of the moments (Franke, 2009),

$$W = \hat{\Sigma}^{-1}. \tag{3.13}$$

The covariances in $\hat{\Sigma}$ are estimated by a bootstrap procedure used to construct additional samples from the empirical observations. In the literature, this is often carried out by a block bootstrap (Winker et al., 2007; Franke and Westerhoff, 2011, 2012). However, the original long-range dependence in the return series is interrupted every time two non-adjacent blocks are pasted together. Hence, the independence of randomly selected blocks cannot reproduce the dependence structure of the original sample. This is known as the joint-point problem (Andrews, 2004). Since our estimation is concerned with summary statistics, we can overcome the joint-point problem by avoiding the block bootstrap.

Correspondingly, in order to estimate the variance-covariance matrix $\hat{\Sigma}$, we use a new bootstrap method first described by Franke and Westerhoff (2016). In their work, the authors apply this method to the FW model with the interaction between agents being modelled by the transition probability approach (see Section 2.3.1). Therefore, we are for the first time applying this new bootstrap framework to a model where the agents interact stochastically according to the discrete choice approach (see Section 2.3.2). Departing from the traditional block bootstrap, we sample single days and, associated with each of them, the history of the past few lags required to calculate the lagged autocorrelations. In more detail, the procedure is presented in Algorithm 1. We construct the

Table 3.2: Fw+ model parameters

ϕ	χ	σ_f	σ_c	α_0	α_n	α_p	μ_p	σ_p
0.121	1.555	0.592	1.917	-0.301	1.990	22.741	0.01	0.157

set of time indices, $I_0 = \{1, 2, \ldots, T\}$, and for each bootstrap sample b, we can sample directly from it. Accordingly, a bootstrap sample is constructed by T random draws with replacement from I_0. Repeating this B times, we have $b = 1, \ldots, B$ index sets,

$$I_b = \{t_1^b, t_2^b, \ldots, t_T^B\}, \tag{3.14}$$

from which the bootstrapped moments are obtained.

For a good representation, the bootstrap method is repeated 5000 times, obtaining a distribution for each of the moments. Formally, let $m^b = (m_1^b, ..., m_9^b)$ be the resulting vector of moments, and $\bar{m} = (1/B)\sum_b m^b$ be their mean values. Then, the estimate of the moment covariance matrix $\hat{\Sigma}$ becomes,

$$\hat{\Sigma} = \frac{1}{B}\sum_{b=1}^{B}(m^b - \bar{m})(m^b - \bar{m})^T. \tag{3.15}$$

Going back to the estimation problem, we are interested in the set of parameters that minimise the distance function J from Equation 3.12. In order to reduce the variability in the stochastic simulations, the time horizon is chosen to be longer than the empirical sample period T, commonly defined as $S = 10 \cdot T$. Repeated simulations over S periods (or days) are carried out, in search for the set of parameters that minimise the associated loss. To this end, let θ be the vector of parameters and $m = m(\theta; S)$ denote the moments to which vector θ gives rise.

Furthermore, the comparability of different trials of θ is determined by a random seed $a = 1, 2, \ldots,$ let us say. Moreover, let $m^a(\theta, S)$ denote the moment vector obtained by simulating the model with a parameter vector θ over S periods on the basis of a random seed a. The parameter estimates on a random seed a, denoted $\hat{\theta}^a$, are the solution of the following minimisation problem,

$$\hat{\theta}^a = \underset{\theta}{\operatorname{argmin}}\, J[m^a(\theta; S), m_T^{emp}], \tag{3.16}$$

where m_T^{emp} is the moment vector for the empirical S&P 500 data[1]

Although one may think that a simulation over $S = 68660$ days provides a large sample to base the moments on, the variability arising from such different samples is still considerable. Hence, it seems most appropriate to carry out a large number of estimations

[1] For the actual minimisation problem we use the Nelder-Mead simplex search algorithm (Nelder and Mead, 1965).

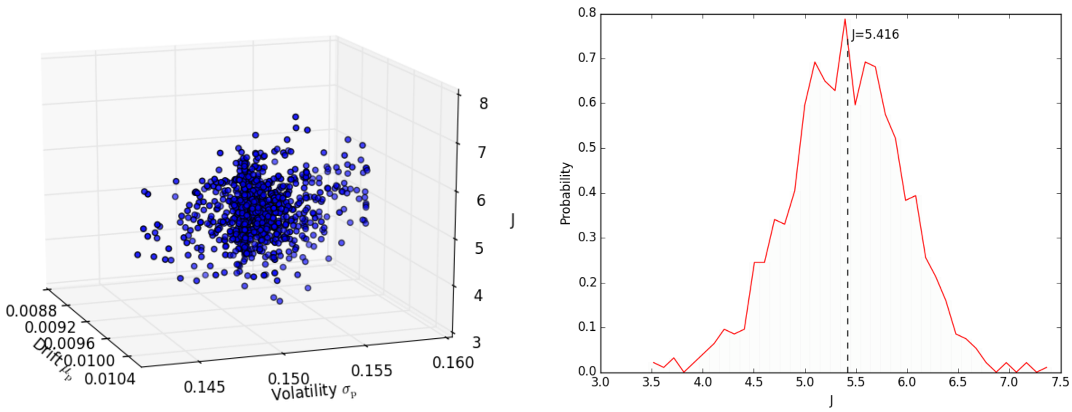

Figure 3.2: Distribution of objective function J.

($a = 1000$) and choose the one with the median loss. The parameter set $\hat{\theta}$ giving this associated loss will be our represeentative estimation (see Algorithm 2). Specifically, using Equation 3.16 we have,

$$\begin{aligned} &\hat{\theta} = \hat{\theta}^{\tilde{a}}, where\ \tilde{a}\ is\ such\ that\ \hat{J}^{\tilde{a}}\ is\ the\ median\ of\ \{\hat{J}^a\}_{a=1}^{1000},\ and \\ &\hat{J}^a = J[m^a(\hat{\theta}^a; S), m_T^{emp}], \quad a = 1, \ldots, 1000 \end{aligned} \tag{3.17}$$

The final parameter vector $\hat{\theta}$ resulting from the estimation has already been reported in Table 3.2. The corresponding minimized loss is,

$$\hat{J} = J[m^{\tilde{a}}(\hat{\theta}; S), m_T^{emp}] = 5.416 \tag{3.18}$$

A visual representation of the distribution of the objective function J can be observed in Figure 3.2. In the left panel, we plot the function as small changes of the last two parameters of the model, the drift and volatility of the fundamental price. We observe how it doesn't depart from small values, staying in the neighbourhood of the medians we chose as the optimal values. In the right panel, we plot the actual distribution of J.

To sum up, the estimation of the model's parameters is based on the minimisation of Equation 3.16, where the objective function J is defined by Equations 3.12, 3.13, 3.15 and the set of nine moments described earlier. The model was validated on the S&P 500 data, leading to the set of parameters presented in Table 3.2. These are the values of parameters we used to generate all the results discussed in the following section.

3.3 Stylized facts

In this section we address one of our main research aims (see Section 1.2) by exploring the statistical properties generated by the new FW+ agent-based model described in the previous sections. Specifically, we provide both quantitative and qualitative measures that demonstrate the existence of the most important stylized facts of financial markets. An in-depth analysis of the price, returns, volatility and volume series generated by the new model will be performed. We show that the model is able to match a rich set of properties including martingales, absence of autocorrelations in raw returns, heavy tails, volatility clustering and long memory in absolute returns, volume-volatility relations, aggregate gaussianity, a concave price impact and extreme price events.

As mentioned above, the model's numerical parameters are given in Table 3.2, with the time unit considered being a day. In Figure 3.3 we present a simple run of the model over 6867 days, covering the same time span as the empirical data. The top panel illustrates the (log) price series p_t generated by the model and the fundamental values p_t^f. We can observe the irregular swings in the prices, with a considerable amplitude, similar to the behaviour of real financial series, reflected by the empirical returns in Figure 3.3(c). The second panel shows the corresponding composition of the traders, or the fluctuations in the market fraction. The majority index x_t is moving continuously, changing from periods of fundamentalists domination (positive values) to periods of chartists domination (negative values). However, it shows that most of the time the market is dominated by fundamentalists ($x_t > 0.6$), with sudden swings to chartism ($x_t \approx -0.65$).

A quick visual inspection of simulated price series and majority index plotted in Figures 3.3 (a) and (c), respectively, reveals a high market fraction of fundamentalists in the presence of strong mispricing. On the other hand, the chartists increase in number when the price gets closer to the fundamental value. This behaviour is obvious from the model definition. The mispricing element in the attractiveness level of Equation 3.9 makes fundamentalism more appealing when the price drifts from the fundamental value. However, an increasing number of fundamentalists begin to dominate the market and their actions start the mean reverting process of the price towards the fundamental value. As the current price gets closer to the fundamental value, fundamentalism is no longer attractive or profitable, making the agents switch to chartism. The herding element takes over, leading to a majority fraction of chartism, whose actions will move the price away from the fundamental value.

The third panel in Figure 3.3 demonstrates the implications of these irregular switches in the agents' strategies on returns, specified as percentage points (as in Equation 3.11). Since the chartists have a greater variability in demand, $\sigma_c^2 > \sigma_f^2$, by comparing Figures 3.3 (b) and (c) we observe that the level of returns during a chartism domination exceeds its level in a fundamentalism regime. Therefore, it appears that normal sequences of

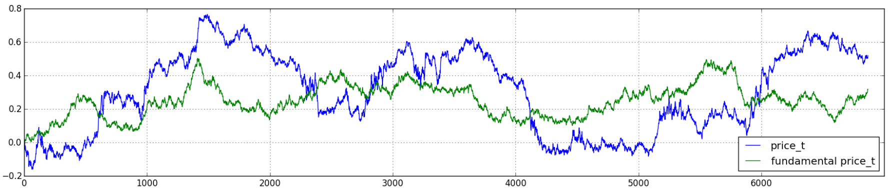

(a) Simulated log prices

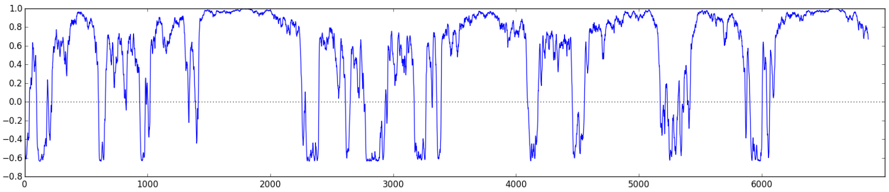

(b) Simulated majority index

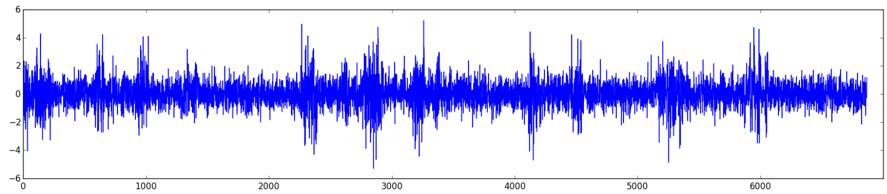

(c) Simulated returns

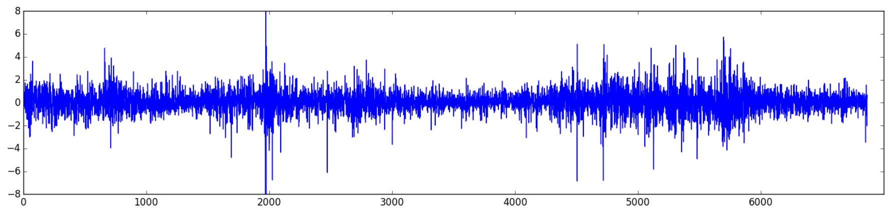

(d) Empirical returns

Figure 3.3: Simulated FW+ price series (p_t), fundamental price series (p_t^f), majority index (x_t) and returns (r_t) over 6867 time periods, together with empirical S&P 500 returns.

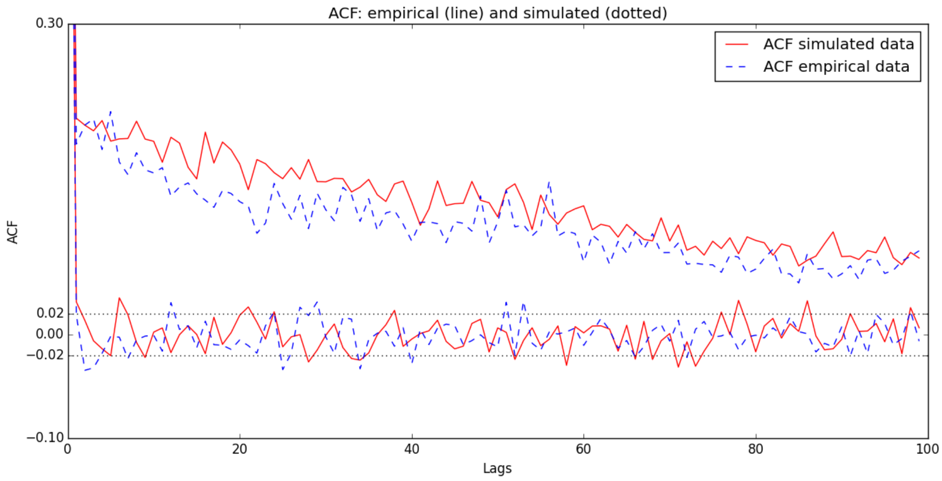

Figure 3.4: Autocorrelation function of FW+ returns. The red lines (blue dots) indicate empirical (simulated) returns. The two upper (lower) lines represent the ACF of absolute (raw) returns r_t at lags $\tau = 1, ..., 100$

returns are interrupted by outbursts of increased volatility, when the majority of agents are chartists.

In the next subsections we discuss all the stylized facts of real financial markets recreated by our model. We offer both qualitative and quantitative measures that demonstrate one of the main objectives of agent-based financial modelling. That is, our model is able to match a rich set of statistical properties observed in financial series.

3.3.1 Absence of autocorrelations

A well-known property of financial data states that price movements do not exhibit any significant autocorrelation. Many studies have found a rapid decline of autocorrelation function (ACF) after the first lag, therefore confirming the absence of (linear) autocorrelation in returns at all horizons and making it a well accepted stylized fact (see Section 2.5.2). The absence of autocorrelation in returns can easily be demonstrated using the autocorrelation function (ACF). This basic stylized fact is demonstrated in Figure 3.4, where the ACF of raw returns is computed at lags from 1 to 100. We compare the autocorrelation function of both simulated and empirical returns showing that it becomes close to zero after the first lag. This behaviour is in line with the empirical findings discussed in Section 2.5.2 and reinforces the idea that price changes are not correlated.

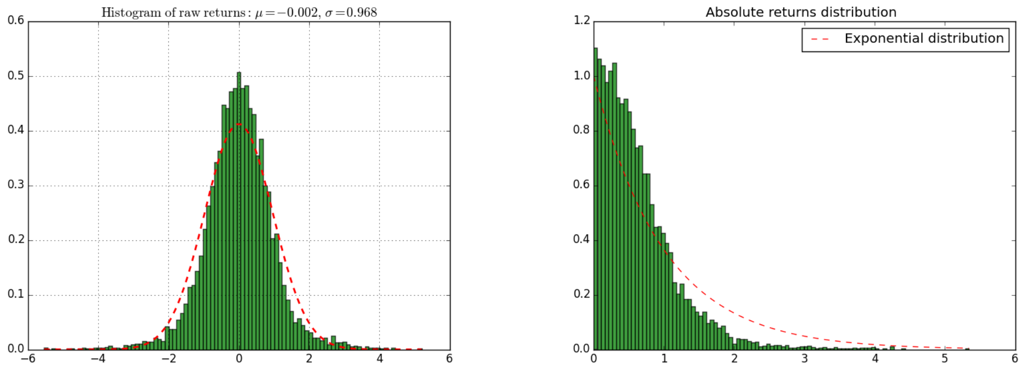

(a) Raw returns distribution. The red line represents a normal distribution superimposed on the return distribution.

(b) Absolute returns distribution. The red line represents an exponential distribution superimposed on the return distribution.

Figure 3.5: Distribution of FW+ (a) raw returns and (b) absolute returns.

3.3.2 Heavy tails

The distribution of returns is another challenging topic in the econometrics literature. We consider the distribution on raw returns, plotted in Figure 3.5(a). We can immediately observe that the distribution displays strong deviations from Gaussianity, which are typical of most financial series. We can also see that the distribution contains more mass in the peak than the normal distribution, having an excess kurtosis of 2.49 and a small negative skewness of -0.0055.

Furthermore, the deviations from Gaussianity can be noticed in the probability plots (PP-plot). Specifically, the PP-plot compares the empirical cumulative distribution function of the returns' distribution with the normal one, calculating a best-fit line. Therefore, we generate a probability plot of our sample data against the quantiles of the corresponding normal distribution. The leptokurtic behaviour can be observed as an elongated S in Figure 3.6(a).

It is clear that the distribution of price variations, or returns, is not a Gaussian. In detail, very large fluctuations are more likely in a stock market and dramatic crashes are observed more frequently. In order to characterise the probability of such events, we introduce the complementary cumulative distribution function (CCDF),

$$F(x) = 1 - Prob(X < x), \tag{3.19}$$

which describes the trail of the distribution $P(x)$ of returns.

The complementary cumulative distribution of returns is found to be approximately a power law $F(x) \propto \tilde{x}^{-\alpha}$ with an exponent in the range 2-4 (Bouchaud and Potters, 2003). Since its decay is slower than Gaussian, we say the returns have a heavy or fat tail. The right tail (positive returns) is usually characterised by a different exponent with respect

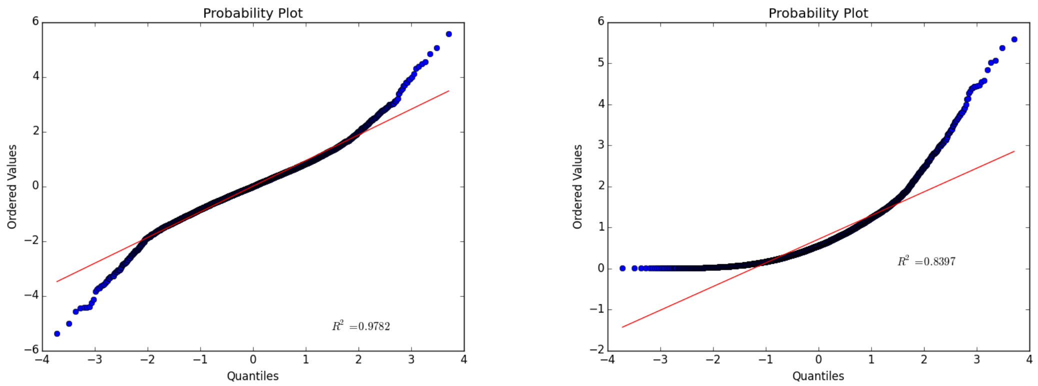

(a) Raw returns vs Normal distribution

(b) Absolute returns vs Normal distribution

Figure 3.6: PP-plot of FW+ (a) raw returns and (b) absolute returns versus Gaussian distribution. The elongated S in (a) indicates an excess kurtosis in the returns distribution, while (b) indicates a distribution of absolute returns with heavy tails.

to the left tail (negative returns). Accordingly, the distribution is asymmetric with respect to the mean and the left tail is heavier than the right one. In Figure 3.7 we present the complementary cumulative distribution function of returns compared to a power law decay with exponent $\alpha = 4$ and a Gaussian with the same mean and variance.

Next, we consider the distribution of absolute returns, or volatility. As we can see from Figure 3.5(b) the distribution of absolute returns has a bigger decay than the exponential distribution superimposed on it. This indicates that the distribution has a fat tail. However, the exact distribution is hard to find. As expected, when performing the Kolmogorov-Smirnov test or the Anderson-Darling test (Stephens, 1974), comparing the absolute returns distribution with various different distributions, we fail to find an exact one that would best describe the distribution of absolute returns. Furthermore, its heavy tail can be observed in the Figure 3.6(b), the tails of the distribution being far away from the corresponding normal ones.

The power law behaviour in the distribution of absolute returns can be approximated by fitting a distribution of the form,

$$p(x) \propto x^{-\alpha}, \tag{3.20}$$

where α is the power law exponent. When we fit such a power law distribution to the distribution of absolute returns, we find an exponent $\alpha = 4.034$, very similar to the one of empirical data $\alpha_{emp} = 4.031$. These values are computed using the end of the tail that minimizes the differences between the power law and the actual distribution. Note that the power law exponent is in line with most of the other empirical findings discussed in Section 2.5.3.

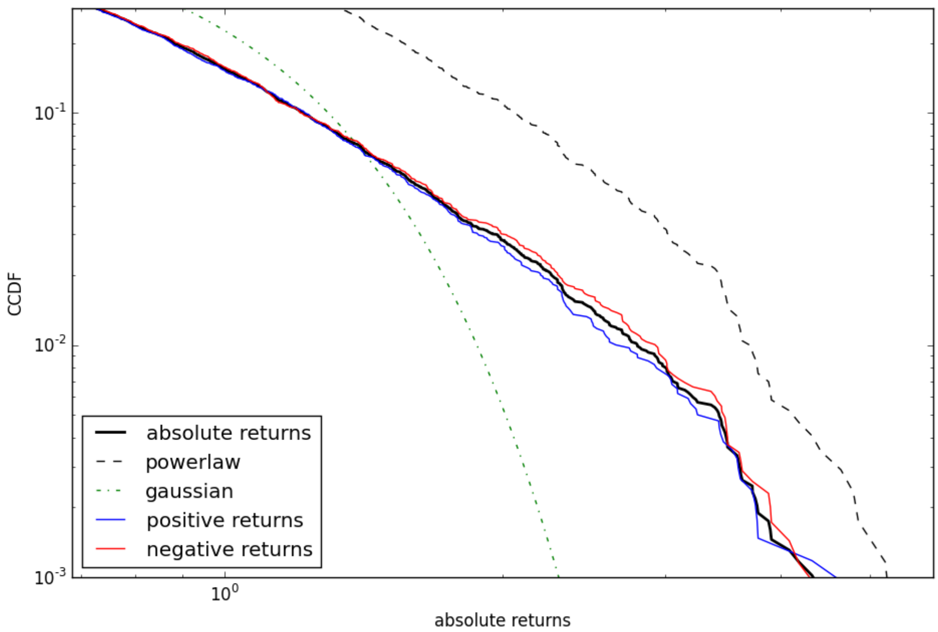

Figure 3.7: Complementary cumulative distribution function of FW+ absolute returns (solid black line). The green dashed line is the CCDF of a Gaussian with the same mean and variance of the return distribution. The dashed black line is a power law decay with exponent $\alpha = 4$. The blue and red lines are the CCDF for positive and negative returns respectively. We observe that the red curve has a slower decay with respect to the blue one.

Finally, we analyse the distribution of absolute returns by computing the well-known Hill tail index. In particular, we obtain a Hill index of 3.69 for the simulated absolute returns, which is not significantly different from the empirical Hill index of 3.30 for the S%P 500 returns. As before, we computed the well-known Hill estimator using the same end of the tail that minimised the differences between the power law and the actual distribution. Once again, this finding is in harmony with the behaviour of the index in empirical financial data (see Section 2.5.3).

3.3.3 Volatility clustering

A simple visual inspection of the volatility measured as absolute returns in Figure 3.5(b) reveals a violation of the assumption of the independently and identically distributed returns. Accordingly, we can say that volatility reveals the so called heteroscedasticity (see Section 2.5.4), with clear fluctuations in the distribution over time.

Furthermore, looking at the returns generated by the model (Figure 3.3(c)), we observe periods of relatively small returns interrupted by abrupt increases in returns. Moreover, these periods tend to be clustered together. Specifically, small price changes are followed

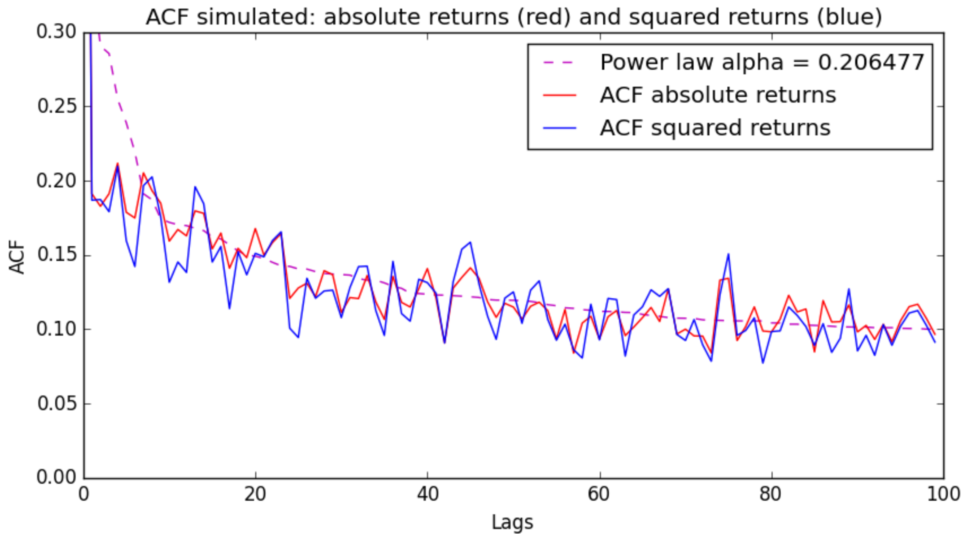

Figure 3.8: Autocorrelation function of FW+ volatility. The red (blue) lines represents the ACF of volatility measured as absolute (squared) returns at lags $\tau = 1, ..., 100$.

by small ones and large price changes by large ones. This is the signature of the well-known volatility clustering. A common way of confirming the presence of volatility clustering is by considering its autocorrelation function. Even though there are different ways of measuring volatility, the most commonly used ones are the absolute and squared returns. Therefore, we plot the ACF of volatility measured as both absolute and squared returns in Figure 3.8. The positive autocorrelation that persists over time, doubled by its slow decay, is a clear presence of volatility clustering. The same kind of behaviour can be seen in empirical data and most of the other empirical studies discussed in Section 2.5.4.

3.3.4 Long memory

The long memory effect specifically addresses the decay in the autocorrelation function of volatility. Usually, if the decay is slow, similar to a hyperbolic function, we can say that the corresponding process exhibits long memory. The autocorrelation of absolute returns remains positive and decays slowly as we increase the lags, but faster than the autocorrelation function of squared returns, in harmony with empirical findings.

One way of observing the decay in the ACF is by fitting a power law as per Equation 2.35. We notice that a close fit of the ACF can be obtained using a power law of the form $p(\tau) \sim A/\tau^{\alpha}$, with an exponent $\alpha \approx 0.21$ (Figure 3.8). A similar behaviour can be seen for the empirical data, as discussed in Section 2.5.5. Notably, these findings are in line with other empirical studies described in Section 2.5.5, where the exponent α is usually reported to fall in the range $[0.2, 0.4]$.

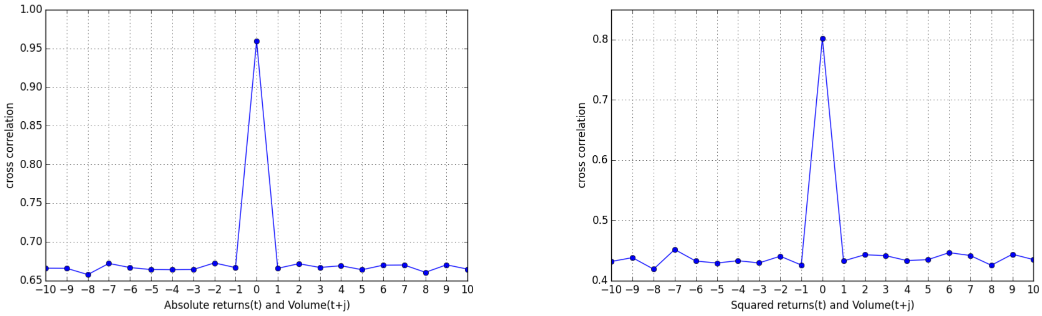

(a) Absolute returns vs Volume

(b) Squared returns vs Volume

Figure 3.9: Volume-Volatility relations: cross correlation between FW+ daily volatility measured as (a) absolute returns and (b) squared returns and volume.

Another widely used method of testing the long memory effect is by using the Hurst exponent. Computing the Hurst exponent of the simulated volatility measured as absolute (squared) returns, we obtain the values of 0.70 (0.71). Similarly, the Hurst exponent in empirical absolute returns is equal to 0.69. Therefore, the presence of a long memory process is clearly demonstrated. We should also note that most empirical studies report a Hurst exponent in the interval $[0.5, 1]$ for long memory processes.

A final formal test for long range dependence in volatility is the logperiodogram regression for estimation of the long memory parameter d, described by Geweke and Porter-Hudak (1983). Using the implementation presented in Pape (2007), we obtain significantly positive estimates of the long memory parameter d over the full samples. For the simulated volatility, we have $d \approx 0.4$, while its empirical correspondent is $d \approx 0.46$. These significantly positive estimates suggest a strong presence of long memory in the volatility of simulated returns.

3.3.5 Volume-volatility relations

The dependence between volatility and volume traded has been noticed and documented across different financial instruments at different time scales. Specifically, the relationship between the two market descriptors can be observed using the cross correlation function. In Figure 3.9 we can observe a significantly positive cross correlation between volatility and volume. We define the volume at time t as the total absolute demand at t of both fundamentalists and chartists. It is important to mention that all measures of volatility are positively correlated with traded volume.

The significantly positive cross correlation means that a small (large) trading volume is accompanied by a small (large) change in volatility. This behaviour is reflected in the way future prices are determined by the market maker, demands (and therefore volume) directly influencing volatility changes. It is important to note that the dependence

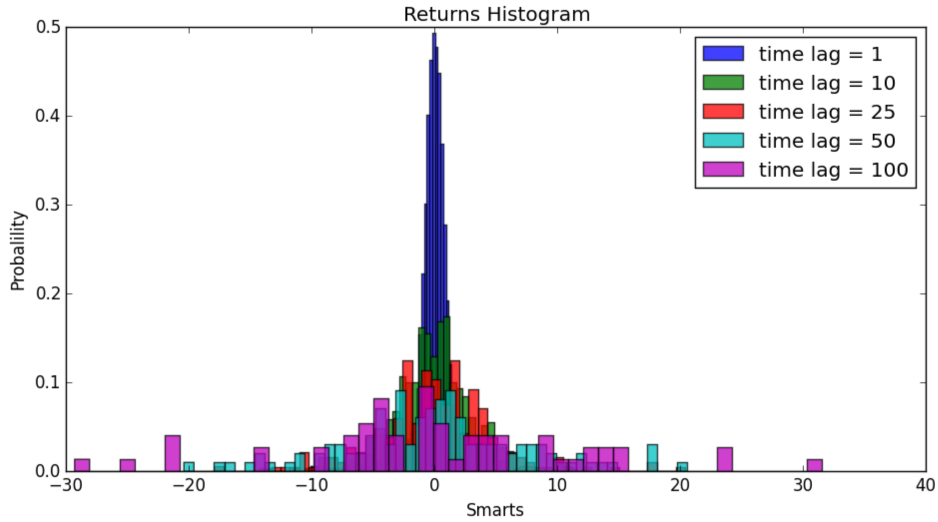

Figure 3.10: Aggregate Gaussianity: distribution of FW+ returns calculated at different time scales $\tau = 1, 10, 25, 50, 100$

.

between volume and volatility remains significant as we increase the time lag. From the investor's point of view, this seems to be the only dependency that can be exploited. As we have seen so far, the returns are not correlated (see Section 3.3.1) and even though the volatility displays positive correlation, it is less significant and slowly decays as we increase the lag, a signature of volatility clustering (see Section 3.3.3).

3.3.6 Aggregate gaussianity

A further easily observable property of returns is the aggregate Gaussianity. As we have seen, the distribution of returns has excessive kurtosis and heavy tails, with a higher peak than the normal distribution and power law tails. However, this behaviour changes as we increase the time scale over which the returns are calculated. More specifically, the distribution of returns changes its shape, becoming more like the normal distribution. This stylized fact can easily be observed in different markets and time periods.

In more detail, the aggregate Gaussianity of the simulated returns generated by our model is clearly visible in Figure 3.10. We plot the distribution of returns calculated for different time lags $\tau = 1, 10, 25, 50, 100$. We can observe how the distribution of returns becomes more like the normal distribution as we increase the time lags. The peaks of the distributions decrease and the tails become wider. Specifically, when returns are calculated at small time scales their distribution is extremely peaked in the middle with very heavy tails. However, as we increase the time scale, the return distributions become more flat, looking like the normal distribution. Moreover, the clear decrease of the excess kurtosis is presented in the following table:

Table 3.3: Excess kurtosis and skewness as we increase the time lags at which FW+ returns are computed.

time lag	1	10	25	50	100
kurtosis	2.5116	2.2022	1.6289	1.0661	0.4954
skewness	-0.0049	-0.0128	-0.0324	-0.0409	-0.0648

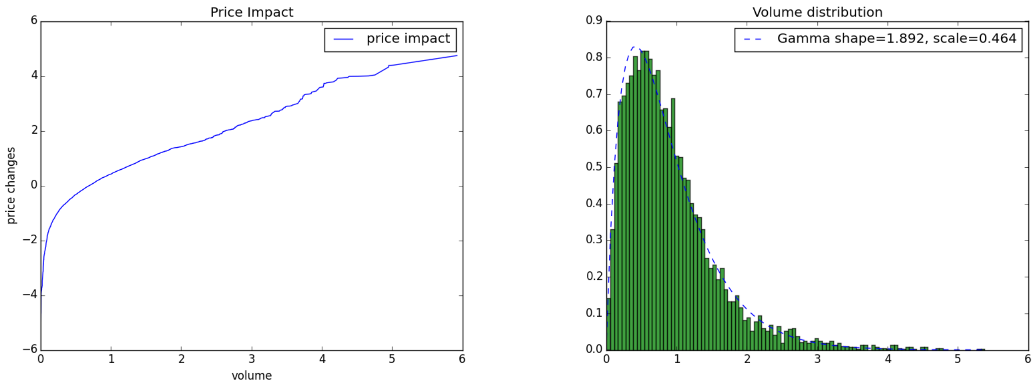

(a) Price impact (b) Volume distribution

Figure 3.11: FW+ price impact function and the FW+ volume distribution.

3.3.7 Price impact and extreme price events

So far, we have discussed the volume-volatility and return-volatility correlations. The price impact tackles the correlations between the volume and the price changes. More specifically, we are interested in the impact of the volume traded on the prices. There are a few different theories regarding this relation, from linear to non-linear dependencies (as discussed in Section 2.5.8). The most recent evidence shows a square-root relationship between the two market descriptors, the order impact being described as a concave function of volume (Mastromatteo et al., 2014).

Using our simulated data, we plot the changes in prices as a function of volume. As before, we refer to the volume in time period t as the absolute value of the total agents' demand in that specific time period. In Figure 3.11(a), we plot the price changes as a function of volume, clearly obtaining a concave function of price impact. We were expecting this kind of behaviour from the model definition. This is because the market maker computes the next price as a function of the agents' demand multiplied by their market fraction. Clearly, if the absolute demands increase, the next price will also increase, but not linearly.

Moreover, in Figure 3.11(b), we show the volume distribution. Although the exact distribution of volumes is hard to find, we notice that the Gamma distribution is the best fit of it. Furthermore, the tail of the volume distribution can be described by a power

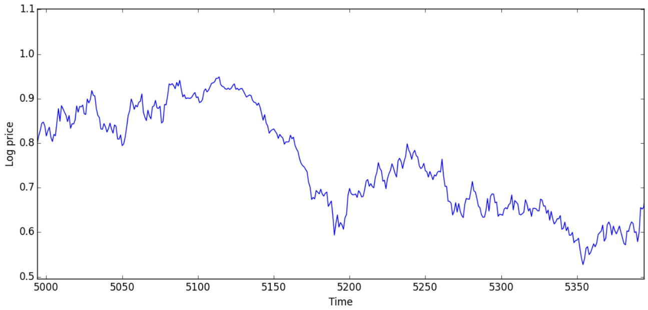

(a) Price series

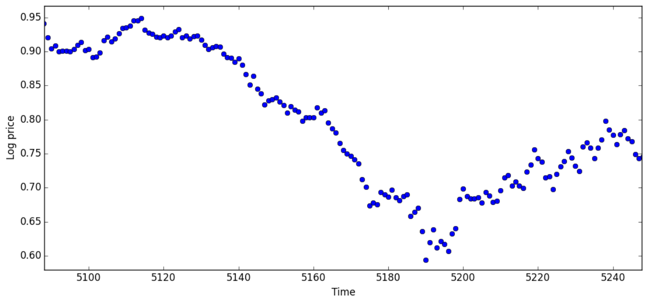

(b) Ticks in price series

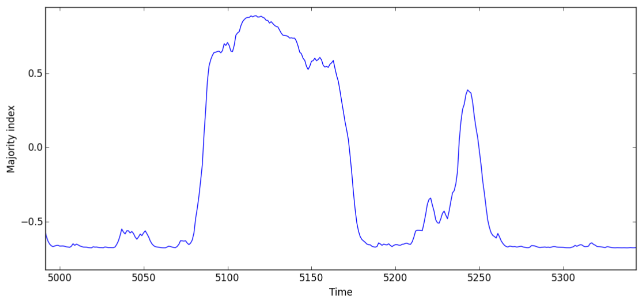

(c) Majority index

Figure 3.12: Extreme price event: a crash in the FW+ price series from $t = 5163$ to $t = 5175$.

law with $\alpha \approx 3.4$. This is similar to the empirical findings regarding the distribution of the number of trades outlined in Section 2.5.8.

Finally, we discuss one of the most recently observed phenomena in real financial markets, the so called extreme price events. Following the definition of Johnson et al. (2013), a spike (crash) is an occurrence of an asset price ticking up (down) at least ten times before ticking down (up) with a price change exceeding 0.8%. On average, we observe 12 spikes or crashes in each simulation (6866 time periods), while there are 9 extreme price events present in the empirical S&P 500 data. In Figure 3.12 we present an example of a crash in the simulated prices occurring between periods $t = 5163$ and $t = 5175$. We observe a drop of 17.19% from the initial price at time $t = 5163$. There is no surprise that the crash occurs in a period of switch from fundamentalism (majority index

> 0.5) to chartism domination (majority index < -0.5), when the herding mechanism suddenly pushes the price away from the fundamental value. The price becomes stable when the majority index changes its sign and the fundamentalists begin to increase in numbers and dominate the market. Moreover, we note that even if the changes in prices are significant, the extreme price events do not usually last much longer that 10 time periods.

3.4 Summary

In this chapter we described our improved structural stochastic volatility agent-based model of asset pricing and the financial markets' stylized facts it is able to recreate. Starting from the original fundamentalist vs. chartist model proposed by Franke and Westerhoff, we discuss its key elements including how the asset prices are generated and how the market participants interact with each other, influencing their strategies. In these first sections, we address one of our key objectives outlined in Section 1.2 and focus on the structure of the model itself.

Next, motivated by the violation of a central property of financial asset prices, their non-stationarity, we introduced our first research contribution. Based on the mathematical description of real financial price series, we changed the mechanism generating the exogenous fundamental value to a random walk. More specifically, the fundamental value was set to follow a geometric Brownian motion. In this way, the number of stationary price series generated was drastically reduced, matching the way real financial data behave. Furthermore, we showed how the model is estimated and validated on empirical S&P 500 index data using a modified method of simulated moments for the first time applied to this setting.

Finally, we tackle one of the main aims of our research, namely the stylized facts matched by an agent-based model. We provide a thorough analysis of the time series produced by the new FW+ model and show that we are able to recreate a rich set of statistical properties observed in real financial markets. Moreover, not only did we match the stylized facts reported by Franke and Westerhoff (2012), but also offered more statistical tests that reinforced their presence. In addition, we were able to recreate new stylized facts including volume-volatility relations, aggregate gaussianity, price impact and extreme price events. By so doing, we address our second contribution discussed in Section 1.3, that our FW+ model is the only one that matches such a rich set of the stylized facts of real financial markets.

Chapter 4

Behavioural Biases in Agent-Based Models of Asset Pricing

In this chapter we further extend the FW+ model presented in the previous chapter to test the implications of some of the most researched behavioural heuristics observed for real life investors. Indeed, one of the objectives of a well-defined ABM that is able to replicate the real life properties of financial data is to test various behavioural and economical theories (as discussed in Section 1.2). Thus, although the FW+ model already incorporates some crucial social psychological factors such as herding and predisposition towards a certain strategy, we are interested in observing how further additions change the interactions between agents and prices. Specifically, we will focus on two of the most discussed and interconnected behavioural preferences observed in financial markets, namely loss aversion (see Section 2.4.1 for more details) and the disposition effect (see Section 2.4.2 for more details).

In general, the field of behavioural finance studies the most relevant psychological aspects of investors' decision-making. Human-based experiments have shown that when it comes to financial decisions, people's beliefs are often predictable. In many cases, the source of the problem is cognitive and therefore a function of how people think. There are a wide range of biases, heuristics and preferences considered, ranging from situational or personality-related, to completely non-economical ones (for more details see Section 2.4).

Now, we are particularly interested in two of these factors, both derived from prospect theory (see Section 2.4). On one hand, the so-called loss aversion, which was shown to effectively explain some of the financial properties which cannot be justified by traditional finance theory, refers to investors' unwillingness to realise and cut his or her losses. This concept applies when one is trying to avoid a loss even if it means to accept a higher risk.

In more detail, the concept of loss aversion can be easily illustrated by an early example presented in Samuelson (1963), who once offered a colleague to flip a coin, heads you win \$200 and tails you lose \$100. He reported that his colleagues turned down this bet because he would feel the \$100 loss more than the \$200 gain. Intuitively, this is the sentiment behind the concept of loss aversion. In other words, there is an asymmetry in how people perceive gains and losses. Specifically, Tversky and Kahneman (1992) found that people are twice more affected by losses than by gains of equal amounts.

On the other hand, Shefrin and Statman (1985) use prospect theory to show that investors may hold losing trades longer than profitable ones – a now well-known phenomenon labelled as the disposition effect. Put differently, investors may not realise a loss because they are trying to avoid the regret associated with their initial (wrong) judgement.

However, although behavioural biases have been clearly observed in real-life settings, examining their implications on financial markets is still a challenging topic. Nevertheless, a well-defined agent-based model is an important tool for testing behavioural and economical theories (Westerhoff and Dieci, 2006) and understanding their influence on the interplay between agents and prices (as outlined in Section 1.2). For this reason, we present two new behavioural agent-based models that, as we demonstrate, capture the loss aversion and the disposition effect in traders' behaviour.

In the loss averse setting, we address one of the main objectives in agent-based financial modelling, namely how good is the model in explaining real world data (as per Section 1.2). We demonstrate that introducing loss aversion leads to a new behavioural model that better matches the properties of financial time series. Measuring how close a model is to its empirical counterparts, we show that the degree to which the data's summary statistics are matched is higher in the model with loss aversion compared to the same model without loss aversion. This represents one of our research contributions discussed in Section 1.3. Moreover, we show that for certain degrees of loss aversion, the model will produce some interesting results, leading to the disappearance of the loss averse traders from the market and altering the returns series generated by the model (see Section 1.3).

Similarly, we test the implications of the disposition effect on both agents' interactions and the asset price dynamics. We find that while the time horizon over which investors consider their wealth has no impact on the overall setting, the level of regret directly impacts the returns series produced by the model, altering its stylized facts and leading to disruptive behaviour (in line with our contributions discussed in Section 1.3). We can relate these findings with the behavioural finance literature by showing that investors' biases generate unexpected market behaviour, instabilities and systematic errors, thus offering an alternative to the rational expectations hypothesis.

This chapter is organised as follows. In Section 4.1, we formally define the two behavioural asset pricing models, having the same price dynamics as the FW+ model but with different factors influencing the agents' interactions. We then estimate the models' parameters values in Section 4.2. Next, in Section 4.3 we assess the implications of the newly introduced behavioural factors and the differences they make on both agents' interactions and the time series generated by the models. In Section 4.4, motivated by the implication of loss aversion, we show that this model better matches the empirical data's statistics, compared to the model without loss aversion. Finally, in Section 4.5 we discuss all the statistical properties generated by the new model, showing that it is able to match the same rich set of stylized facts.

4.1 Model definition

In this section, we further extend the fundamentalist vs. chartist model presented in Chapter 3 (FW+), to capture two of the most researched biases in investors behaviour. In Section 4.1.1 we specify the models' price dynamics, common for the two settings. Next, in Sections 4.1.2 and 4.1.3 we define the agents' interactions, building on the FW+ model to accommodate the introduction of the disposition effect and loss aversion, respectively.

4.1.1 Price dynamics

As we have seen in Chapter 3, heterogeneous agent-based models with simple trading strategies are good at generating the most important properties of real financial markets. For this reason, we will build on the FW+ fundamentalist vs. chartist model to test the effects of two highly discussed behavioural preferences. As before, the market is populated by two classes of agents, fundamentalists and chartists.

The new behavioural models defined in this chapter have many common characteristics with the previous FW+ one and a similar price formation mechanisms. In the new settings, prices are determined by a market maker as in Equations 3.6 and 3.7, with fundamentalists' and chartists' demand given by Equations 3.1 and 3.2, respectively. Similar to the FW+ model, the fundamental price is considered to follow a geometric Brownian motion, as per Equation 3.11.

So far, the settings of the new behavioural models are identical to the one presented in the previous chapter. However, the concepts of loss aversion and the disposition effect are both linked to a level of wealth, from which agents can derive the domains of gains and losses (Tversky and Kahneman, 1992). In this context, one of the principles that influence agents' preferences is based on differential profits, incorporating some inertia. In detail, with respect to strategy $s = f, c$, let g_t^s be the short term capital gains that an average agent in this group could realise at time t. This is obtained from the demand formulated at time $t-2$ and executed at the price p_{t-1} for the next day. In general, a utility v_t^s obtained from these gains is defined as $v_t^s = g_t^s + \eta v_{t-1}^s$, where η is a memory coefficient between 0 and 1. Following the agents' wealth definition of Franke and Westerhoff (2012), we consider a weighted average of g_t^s and v_t^s such that:

$$\begin{aligned} g_t^s &= [exp(p_t) - exp(p_{t-1})]d_{t-2}^s, \\ w_t^s &= \eta w_{t-1}^s + (1-\eta) g_t^s. \end{aligned} \tag{4.1}$$

Note that we replaced the utility symbol u with w to accommodate our interpretation. Furthermore, w_t^s represents the accumulated profits, discounted by the coefficient $\eta < 1$, that would have been earned by an agent who had consistently followed strategy s on a daily basis. In other words, w^s is the accumulated wealth attributed to strategy s.

Now, one of the key elements that differentiates the way we model loss aversion and the disposition effect is given by the motion of the market fractions, n_t^f and n_t^c. In the next two sections we discuss in more detail the way we incorporate these behavioural biases in our agents' behaviour.

4.1.2 Social interactions with the disposition effect

In order to incorporate the disposition effect into our agent-based setting, we model agents' interactions by making use of the transition probability approach (TPA), presented in Section 2.3.1. Its basic idea is that agents switch strategies (fundamentalism and chartism) with certain probabilities, relying on the exponential of a switching index, a_t, such that:

$$\begin{aligned} P_t^{cf} &= \nu \exp{(a_t)}, \\ P_t^{fc} &= \nu \exp{(-a_t)}, \end{aligned} \tag{4.2}$$

where P_t^{cf} and P_t^{fc} are the probabilities of shifting from chartism to fundamentalism at time t and vice-versa, respectively. The parameter a_t is determined by several factors including a predisposition towards one of the two agent types, a herding factor and a price misalignment between p_t and p_t^f, as in Equation 3.9.

To model and test the implications of the disposition effect in our agent-based model, we use the notion of wealth, formally defined by Equation 4.1. In this context, some agents are assumed to be biased towards retaining losing positions in order to avoid the regret of being wrong. As a result, whenever they find themselves in a losing strategy, they do not realise the loss and stick to that strategy hoping that their bad luck will soon reverse. Formally, the probability of the biased agents switching between strategies decreases when they are losing money:

$$\begin{aligned} \hat{P_t^{cf}} &= P_t^{cf} - \lambda_{DE} P_t^{cf}, \quad \text{if } w_t^c < 0, \\ \hat{P_t^{fc}} &= P_t^{fc} - \lambda_{DE} P_t^{fc}, \quad \text{if } w_t^f < 0, \end{aligned} \tag{4.3}$$

where the value of $\lambda_{DE} \in [0, 1]$ measures the agents' unwillingness to cut their loses, avoiding regret. We can view λ_{DE} as the percentage of biased investors within the population.

A further important concept in prospect theory is the horizon over which gains and losses are measured – in other words, how often an agent evaluates his investment performance. In our model, this is specified by the memory parameter $\eta \in [0, 1)$, which indicates how much weight an agent puts on his latest profits and how much on his past wealth, when computing his current, accumulated wealth. Simply put, η relates to the look-back window used by an agent when calculating his wealth, with the time horizon increasing with η. Since the disposition effect is associated with the investors' tendency for ignoring past mistakes and putting them into perspective, the memory parameter should play a role in how this bias affects the agents.

In summary, the new behavioural model with disposition effect, denoted by HPM-DE, is governed by two central dynamic equations. Firstly, the price adjustments are made by a market maker according to Equations 3.6 and 3.7. Secondly, the changes of the market fractions n_t^f and n_t^c are changing according to the transition probabilities approach described by Equation 4.3, based on an attractiveness level encapsulating a predisposition towards one of the two agent types, a herding mechanism and a price misalignment, as in Equation 3.9. In Section 4.3.1 we will discuss how this behavioural addition changes the interaction between agents and market.

4.1.3 Social interactions with loss aversion

In this section, we define a different behavioural agent-based model, incorporating loss aversion in agents' interactions. In order to account for this behavioural bias, we model agents' switching mechanism using the discrete choice approach (DCA) (see Section 2.3.2 for more details), since Franke and Westerhoff demonstrated that it better represents the empirical properties of financial data, compared to the TPA.

In the FW+ model, the switching mechanism was based on the discrete choice approach, using the notion of an attractiveness level, a_t, defined as function of some behavioural factors such as herding, price misalignment and the predisposition towards a particular strategy (see Equation 3.9). However, the concept of loss aversion is linked to a level of wealth, from which we can derive the domains of gains and losses (as discussed in Section 2.4.1).

Our implementation of loss aversion is slightly different from prospect theory in the sense that agents are considering a form of profits incorporating some inertia, formally defined in Equation 4.1, instead of just the latest changes. This difference is imposed by the setting of the model. Specifically, the agents choose their strategies (chartism vs. fundamentalism) based on how well they have worked over time, rather than just the latest observed performance. Therefore, loss averse agents consider a modified wealth function such as:

$$w_t^s = \begin{cases} w_t^s & \text{if } w_t^s \geq 0 \\ \lambda_{LA} w_t^s & \text{if } w_t^s < 0, \end{cases} \tag{4.4}$$

where λ_{LA} is the loss aversion coefficient, $s = f, c$. In their seminal paper, Tversky and Kahneman (1992) argue that people weigh losses twice as much as gains of a similar magnitude, fixing the loss aversion coefficient at $\lambda_{LA} = 2.25$. In this context, we introduce the first term in the attractiveness level a^t as being proportional to the difference in accumulated wealth, $(w_t^f - w_t^c)$.

The second principle considered is based on one of the most empirically observed social factors in behavioural literature, known as herding (see Section 2.2). This is the investors' tendency to follow the strategies of other traders. It is one of the essential irreducible elements that an agent-based model of financial markets should posses (Alfi et al., 2009). The basic idea of herding is that joining a group becomes more attractive the more adherents it already has. In our model, this principle will be represented by a term proportional to the most recent difference between market fractions, $(n_t^f - n_t^c)$.

The last term of the attractiveness level is based on a predisposition towards one of the two strategies. This is simply expressed by a constant factor, α_0, which is positive (negative) if there is a priori preference toward fundamentalism (chartism). These three elements combined form the attractiveness level, formally defined as:

$$a_t = \alpha_0 + \alpha_n(n_t^f - n_t^c) + \alpha_w(w_t^f - w_t^c), \tag{4.5}$$

where α_0 is the predisposition parameter, $\alpha_n > 0$ captures the herding parameter and $\alpha_w > 0$ measures the influence of accumulated wealth of Equation 4.4.

In summary, the new behavioural model with loss aversion, denoted by WHP-LA, is governed by two central dynamic equations. Firstly, the price adjustments are made by a market maker according to Equations 3.6 and 3.7. Secondly, the changes of the market fractions n_t^f and n_t^c are described by Equations 3.8 and 4.5, based on the level of wealth (Equations 4.1-4.4) and encapsulating a herding mechanism. Note that the agents are assumed to be loss averse (Equation 4.4). In the next section we provide a summary of the parameter estimation method for our two behavioural models, HPM-DE and WHP-LA, similar to the one employed in Section 3.2 for the FW+ model. Next, we will focus on the changes brought by the two behavioural biases in Section 4.3.

Table 4.1: HPM-DE model parameters

ϕ	χ	σ_f	σ_c	α_0	α_n	α_p	μ_p	σ_p
0.18	2.30	0.790	1.900	-0.161	1.30	12.50	0.014	0.17

Table 4.2: WHP-LA model parameters

ϕ	χ	σ_f	σ_c	α_0	α_n	α_w	μ_p	σ_p	η
1.00	0.912	0.733	1.760	2.127	1.277	2679	0.014	0.17	0.987

4.2 Estimation of the model's parameters

The formal estimation of the behavioural models' parameters defined in the previous sections (HPM-DE and WHP-LA) is performed using the method of simulated moments (MSM) (see Section 3.2 for more details). Similarly to the estimation of the FW+ model parameters, the empirical data consists of $T = 6866$ daily observations of the $S\&P$ stock market from January 1980 to mid-march 2007. Moreover, we the models are estimated on the basis of the same nine moments as before.

In more detail, the estimation of the models' parameters is based on the minimisation of Equation 3.16, where the objective function J is defined by Equations 3.12, 3.13, 3.15 and the set of nine moments described earlier. Both behavioural models, HPM-DE and WHP-LA, were validated on the S&P 500 data (see Table 3.1 for the summary of empirical moments), leading to the set of parameters presented in Tables 4.1 and 4.2, respectively. These are the parameters values we use to generate all the results of the two models, discussed in the following sections. Note that the drift and volatility of fundamental price refer to a yearly time scale.

4.3 Behavioural implications

We now turn to explore the inner workings of our behavioural models and the impact of the disposition effect and loss aversion on the interactions between agents and price changes. First, we consider the effects of the disposition effect on the HPM-DE model in Section 4.3.1, followed by the changes brought by loss aversion on the WHP-LA model in Section 4.3.2.

4.3.1 Impact of the disposition effect

To assess the implications of investors' unwillingness to realise their losses in the HPM-DE model, we vary both the regret $\lambda_{DE} \in [0, 1]$ and the memory $\eta \in [0, 1)$ parameters, considering 100 equally spaced values for each parameter. We fix all the other parameters to their values reported in Table 4.1. Interestingly, Figures 4.1 (a) and (b) shows that

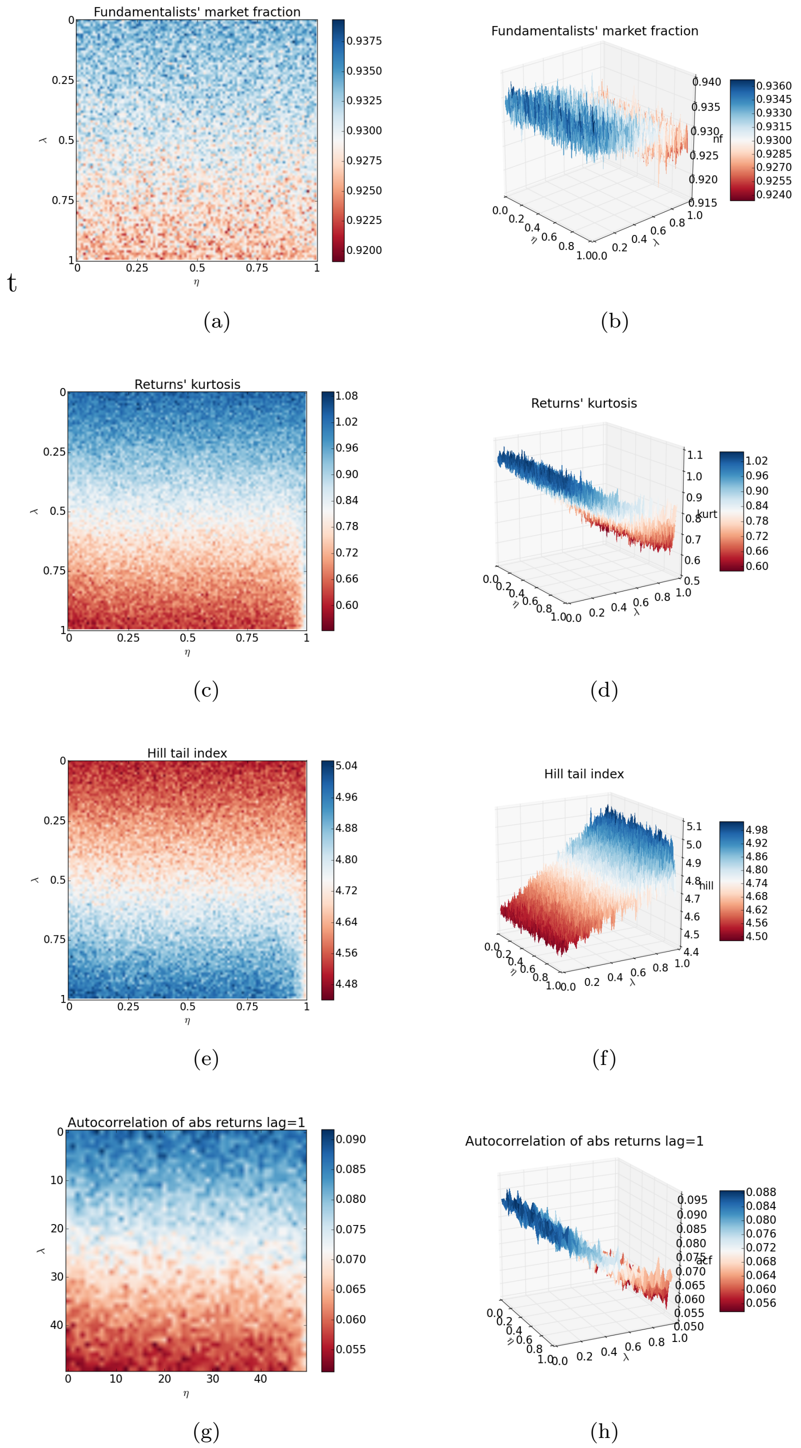

Figure 4.1: Impact of the disposition effect on HPM-DE agents' interactions and returns.

the agents' market fractions are not notably affected by either λ_{DE} or η, having a median value of $n_f \approx 0.93$. This is due to the stochastic behaviour of prices leading to the wealth changing from positive to negative values. Therefore, the changes in the transition probabilities balance out and market proportions don't change much, with both fundamentalist and chartists surviving the competition, as we often see in these settings.

We now turn our attention to the implications of memory and regret on the returns series produced by our behavioural model. The returns themselves are specified by percentage points, as in Equation 3.11. In Figures 4.1 (c), (d), (e), (f), (g) and (h) we plot the returns' excess kurtosis, its Hill tail index and the autocorrelation of absolute returns for different levels of λ_{DE} and η. We observe that the memory parameter doesn't affect the returns at all. This leads to the counter intuitive finding that the horizon over which wealth is measured doesn't change the agents' interactions or the model's dynamics (as discussed in Section 1.3). Therefore, the investors' unwillingness to cut their losses has the same effect on the overall dynamics irrespective of how often they evaluate their performance.

On the other hand, the regret parameter has a direct impact on the properties of returns. As λ_{DE} increases, its excess kurtosis drops 40% from 1.00 to 0.60, while the Hill tail index grows 11% from 4.50 to 5.00. This means the peakiness of returns and its tail heaviness decrease as the agents become more biased, leading to a flatter and more normal distribution of returns. In other words, we observe a decrease in the heavy tails of returns. Moreover, Figures 4.1 (g) and (h) show how the first autocorrelation of absolute returns decreases 42% from 0.93 to 0.54 as we increase λ_{DE}. This is a clear sign of decline in volatility clustering.

Therefore, an increase in the disposition effect leads to violations of some of the most discussed stylized facts of financial data[1]. Note that the outbursts in volatility appear as a consequence of chartism dominance periods. Accordingly, in a setting where p_t deviates from p_t^f, the fundamentalists lose money and their bias lowers the probability of shifting to chartism. Moreover, if the price follows the trend, the chartists make profit and their probability of switching to fundamentalism is not affected by the disposition effect. Accordingly, the fundamentalists' market fraction stays high, with fewer outbursts of chartism and volatility clustering. Overall, investors unwillingness to cut their losses has a significant impact on the time series produced by our model, altering the properties we usually observe in real-financial markets.

[1] For values of η close to 1 ($\eta > 0.995$), there are no disruptions in the returns' stylized facts. For such large values of η, the impact of the latest realized profit or loss will be very small on the total wealth. Since the price misalignment component ensures that the market price doesn't depart too much from the fundamental one, the agents' wealth would stay positive over time. This means that the behavioural bias won't be activated and the new model won't differ from the original one

4.3.2 Impact of loss aversion

Similar to the previous section, we now turn to explore the implications of loss aversion on the WHP-LA model. We start by providing the qualitative measures that demonstrate the existence of the most important stylized facts of financial markets. An in-depth analysis of the price, returns, volatility and volume series generated by the behavioural WHP-LA model will be performed in Section 4.5. The numerical parameters used for all the simulations are given in Table 4.1. Moreover, the loss aversion parameter is fixed at $\lambda_{LA} = 2.25$, as reported by Tversky and Kahneman (1992).

In Figure 4.2 we present a simple run of the model over 6867 days, covering the same time span as the empirical data. The top panel illustrates the (log) price series, p_t, generated by the model and the fundamental values, p_t^f. We can observe the irregular swings in the prices, with a considerable amplitude, similar to the behaviour of real financial series. The second panel shows the corresponding composition of the traders, or the fluctuations in the market fraction. The majority index, x_t, is moving continuously, changing from periods of fundamentalists domination ($x_t > 0.5$) to periods of chartists domination ($x_t < -0.5$). However, it shows that most of the time the market is dominated by fundamentalists, with sudden swings to chartism.

A quick visual inspection of the simulated price series and market fractions plotted in Figures 4.2 (a) and (b), respectively, reveals a high market fraction of chartism when the price deviates from its fundamental value. On the other hand, the fundamentalists increase in number when the price gets closer to the fundamental value. This behaviour is obvious from the model definition. The wealth element in the attractiveness level (Equation 4.5) makes fundamentalism more appealing when the price moves towards its fundamental value, making it a profitable strategy. An increasing number of fundamentalists begin to dominate the market and their actions start the mean reverting process of the price towards the fundamental value. As the current price gets close to the fundamental value, fundamentalism is no longer attractive or profitable, making the agents switch to chartism. The herding element takes over, leading to a majority fraction of chartism, whose actions will move the price away from the fundamental value.

In addition, Figure 4.2 (c) plots the accumulated wealth, w_t, of the two classes of traders. We can observe that when the actual price comes closer to the fundamental value the fundamentalists are more profitable, while the chartists make bigger profits when prices move away from the perceived fundamental value. Finally, the fourth panel in Figure 4.2 demonstrates the implications of these irregular switches in the agents' strategies on returns. Note that the volatility or swings in chartists wealth is expected to be higher than for fundamentalists, since they have a greater variability in demand. Similarly, by comparing Figures 4.2 (b) and (d) we observe that the level of returns during a chartism domination exceeds its level in a fundamentalism regime. Therefore, it appears that

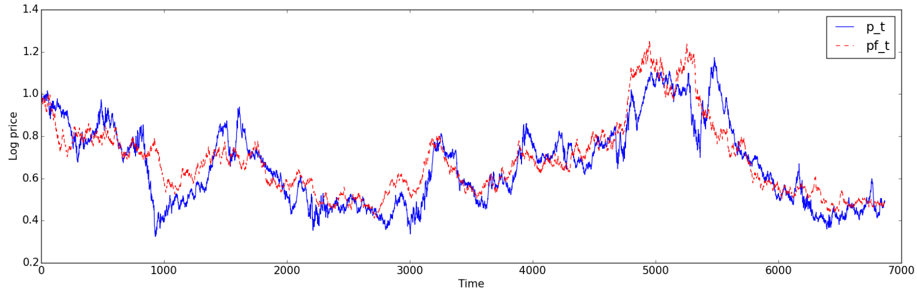

(a) Log prices

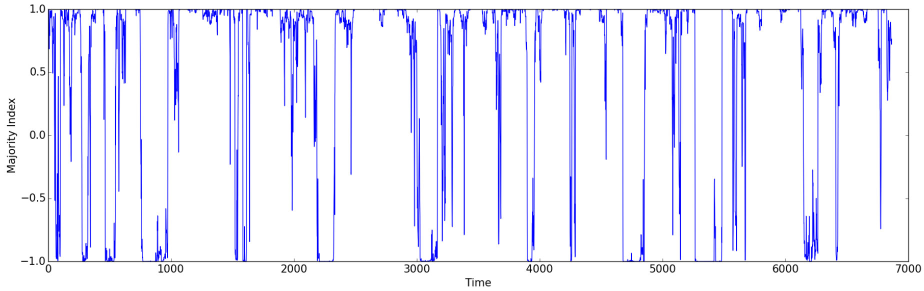

(b) Majority index

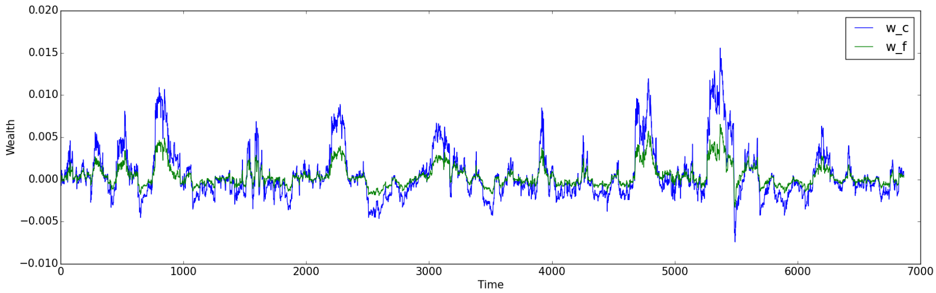

(c) Wealth

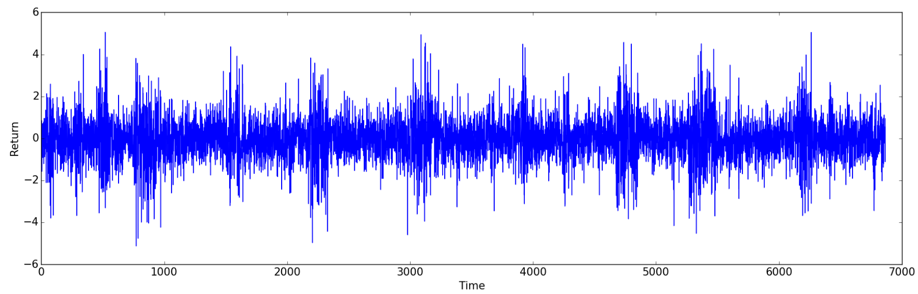

(d) Returns

Figure 4.2: Simulated WHP-LA price series, fundamental price series, majority index, wealth and returns over 6867 time periods.

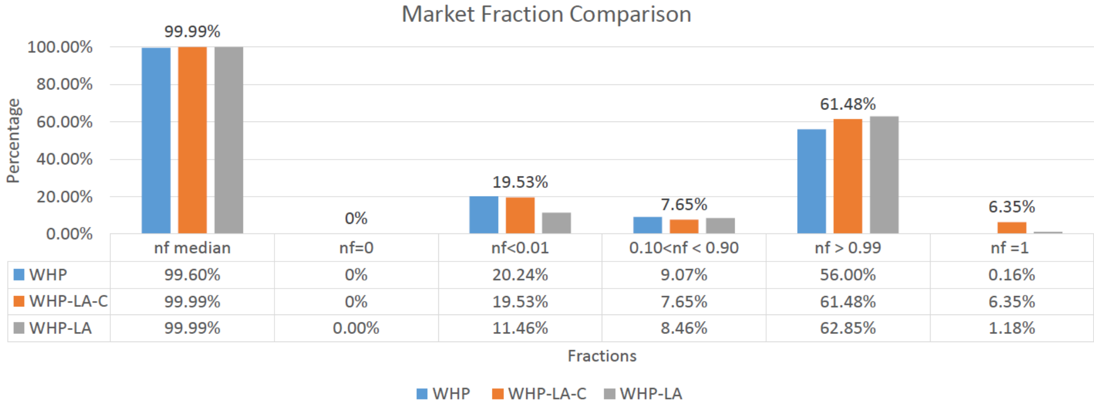

	nf median	nf=0	nf<0.01	0.10<nf < 0.90	nf > 0.99	nf =1
WHP	99.60%	0%	20.24%	9.07%	56.00%	0.16%
WHP-LA-C	99.99%	0%	19.53%	7.65%	61.48%	6.35%
WHP-LA	99.99%	0.00%	11.46%	8.46%	62.85%	1.18%

Figure 4.3: Market fractions comparison WHP vs. WHP-LA-C vs. WHP-LA

normal sequences of returns are interrupted by outbursts of increased volatility, when the majority of agents are chartists.

Clearly, our agent-based model is able to produce the real life dynamics of both price and return series. By inspecting Figure 4.2, we can observe the apparent randomness of the data, similar to what we witness in real financial markets. However, our aim here is not just to replicate the real life dynamics, but use such a model to test the implications of loss aversion (as discussed in Section 1.2). Next, we will show how this behavioural bias influences the interactions between agents and the results reproduced by our model.

In order to assess the impact of loss aversion to the interaction between agents, we perform a more in-depth analysis of the market fractions. For an accurate comparison, we consider 3 models that differ just in this respect. Namely, we compare our previously described behavioural model with loss aversion with the same model but without this bias in agents' behaviour. In particular, we look at the market fractions for three models, all having the switching mechanism dependent on wealth (W), herding (H) and a priori predisposition towards one of the two strategies (P). The first model to consider has the behavioural presence of loss aversion only for chartists (WHP-LA-C), the second model considers both fundamentalists and chartists loss aversion (WHP-LA) and the third one doesn't have the loss aversion bias at all (WHP).

For a more robust comparison, we perform 1000 simulation runs for WHP, WHP-LA and WHP-LA-C models and look at the median results, reported in Figure 4.3. As we have mentioned earlier, the market fractions stay at the ends of the interval, with the market being dominated by either chartists or fundamentalists. For the WHP model, we have median $0.1 < n_f < 0.9$ in 9% of the simulations, for WHP-LA-C 7.65%, while for WHP-LA it is in 8.46% of the simulations. Furthermore, for all three models, the market is heavily dominated by the fundamentalists (median $n_f = 0.99$) with sudden disruptions of chartism, $n_f < 0.01$ in more than 10% of the simulations. At first glance, the market fractions appear to be similar in our three models, WHP, WHP-LA and WHP-LA-C.

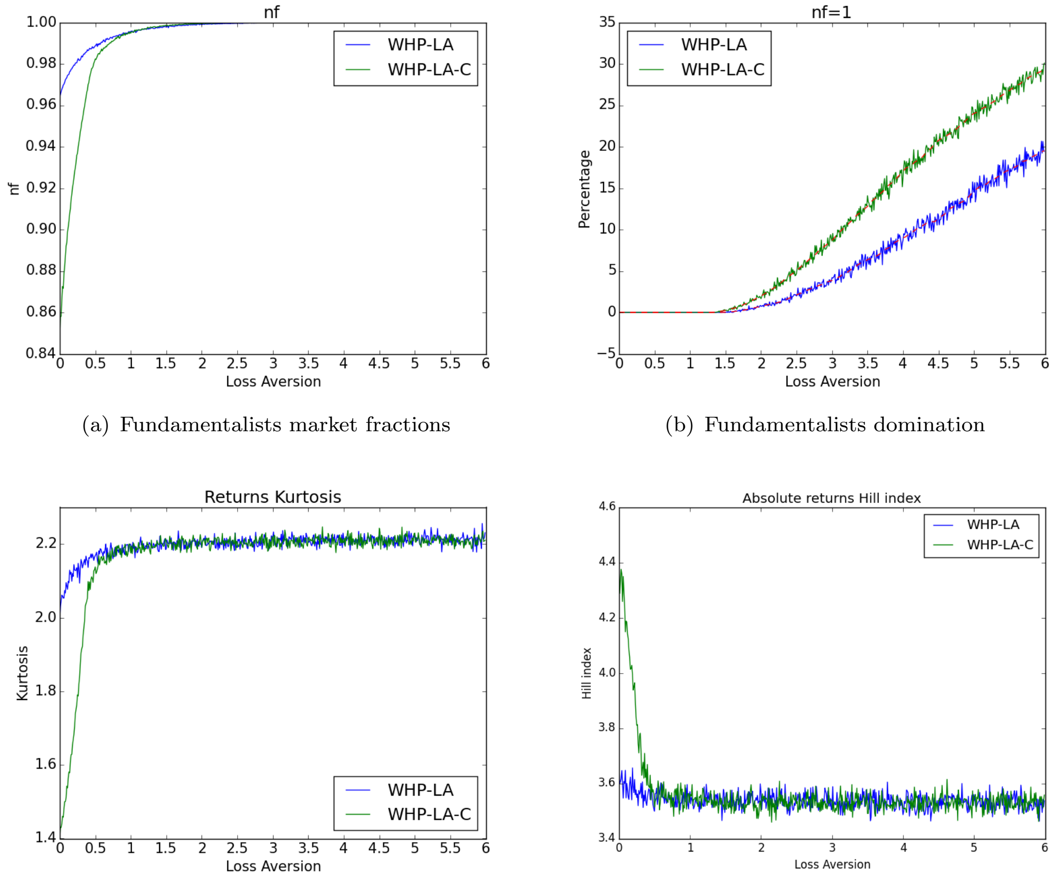

(a) Fundamentalists market fractions

(b) Fundamentalists domination

(c) Kurtosis

(d) Hill tail index

Figure 4.4: Impact of loss aversion parameter on WHP-LA agents' interactions and returns.

However, a closer look at the ends of the interval reveals major differences between the three settings. While fundamentalists never disappear from the market ($n_f > 0$), we notice that for the WHP-LA-C model, the chartists can be driven out of the market[2]. Indeed, the fundamentalists tend to occupy the entire market ($n_f \to 1$) in quite a few cases. In addition, over a median simulation of $T = 6867$ time periods (days), the market is occupied by fundamentalists 9.53% of the time. In contrast, for the WHP or WHP-LA models, the chartists almost never fully disappear from the market ($n_f \to 1$ in less than 1.2% of the simulations). Since the only difference between the three settings is the loss aversion for chartists, we conclude that it is this behavioural bias that pushes these agents entirely out of the market (as outlined in Section 1.3). This is a significant

[2]Note that for any positive finite level of the intensity of choice parameter β, a non-zero fraction of both classes of agents will always be in the market. However, in a setting where the market moves against the loss averse chartists for an extended period of time, they will completely overestimate their wealth and its effect on them. Therefore, their perceived level of wealth will be much lower than the actual one (see Equation 4.4). In turn, this leads to big differences in opinion that determine the attractiveness level in Equation 4.5 and the market fractions in Equation 3.8. The exponential function from Equation 3.8 will then tend to zero and the fundamentalists' fraction n_f will tend to 1. In such cases, we say that the chartists are driven out of the market.

result since even zero intelligence strategies, where agents buy or sell randomly, have been shown to survive the competition in agent-based settings (Ladley, 2012). It may be seen as a confirmation of behavioural finance literature and psychological experiments which argue that peoples' beliefs lead to instabilities in the market and systematic errors which can be exploited, leading to deviations from market efficiency (Shiller, 2015).

Generally speaking, loss aversion makes the chartist weight their losses more than the gains. Hence, when the price is going toward the fundamental value but fluctuates around, without following any trend, the chartists suffer heavy losses. In the presence of loss aversion, they overweight these losses, which makes chartism rapidly unattractive. Because of this significant difference in wealth, the switching mechanism then pushes more and more agents to fundamentalism. Moreover, if the prices stay around their fundamental values long enough, the herding mechanism accentuates the movement. This results in the complete disappearance of chartism. This is the first time this phenomenon has been observed in an agent-based model of financial markets. Usually, the agents interact with each other and their market fractions are continuously changing, but even the most simple strategies survive the competition. However, by introducing a behavioural heuristic to technical analysts, we observe how it leads to their complete extinction from the market. Therefore, in the complex setting of financial markets, where strategies compete with each other, a behavioural bias can be powerful enough to push the agents out of the market.

Note that so far we have kept the loss aversion parameter λ_{LA} equal to 2.25, as it is reported by prospect theory literature. With this level of aversion we have seen that if just one type of agents are biased (WHP-LA-C model), they tend to disappear from the market. However, when we make all agents loss averse (WHP-LA model), the effects induced by this heuristic are not strong enough to produce major disruptions in the model. This happens because of the stochastic behaviour of prices and the simplicity of agents' trading rules. Agents' trivial demands make their wealth change continuously from positive to negative levels. Moreover, with almost half of their time making or losing money, the effects of loss aversion are not prominent.

Given this, it is interesting to determine what levels of loss aversion are required to produce any effects on agents' interaction. To this end, we vary the magnitude of the loss aversion parameter in both the WHP-LA and WHP-LA-C models to test its impact not only on agents' market fractions but also on the system's macro dynamics. Accordingly, we vary the loss aversion parameter in Equation 4.4, considering 600 equally spaced values in the interval $\lambda_{LA} \in [0, 6]$. In Figure 4.4 (a) we plot the loss aversion parameters against the median value of the fundamentalists' market fraction for both WHP-LA and WHP-LA-C. Note that all the simulations have been repeated 1000 times, with the median value being considered. Observe how rapidly the fundamentalists' market fraction increases with the increase of the loss aversion parameter. For both WHP-LA-C and WHP-LA, the chartists have a very small median market fraction across the simulation time (T=6867 days). If $\lambda_{LA} > 2$ we see that chartists usually occupy less than 0.5% of the market.

To analyse the actual impact of the loss aversion parameter on the agents' interaction in greater detail, we look at the end intervals of the market fractions. To this end, Figure 4.4 (b) plots the median percentage of the simulations that see the chartists completely disappearing from the market. As before, each simulation has been carried out 1000 times, with the median value being reported. It is clear that more and more chartists are being thrown out of the market as their loss aversion parameter increases. In the case of the WHP-LA-C model, where just the chartists are biased, the 5% threshold is hit at $\lambda_{LA} = 2.43$. Hence, the interesting dynamics appear for a loss aversion parameter close to the value of 2.25 reported by Tversky and Kahneman (1992). At this level, the fundamentalist fully occupy the market 11.94% of the time, or approximately 3.25 years of the 27 years simulation time.

On the other hand, when we make all agents biased in the WHP-LA model, we observe how chartists disappear from the market at larger values of loss aversion. In this case, the 5% level is hit at $\lambda_{LA} = 3.10$. It is important to note that even when all agents are considered to be biased (not reported here in detail), the chartists don't survive the competition. Moreover, the chartists disappearance is not a rare event, the fundamentalist fully occupying the market 12% of the simulation time of 6966 periods. Interestingly, the fundamentalists always stay in the market, dominating it. Even when we consider a model with biased fundamentalists and unbiased chartists, the fundamentalists' market fractions still dominate the market and chartists survive, similar to what we see in the WHP model.

We now turn our attention to the implications of loss aversion on the simulated returns series produced by our behavioural models. In Figure 4.4 (c) we plot the kurtosis of returns for different levels of loss aversion. We observe how for our WHP-LA and WHP-LA-C models, the kurtosis rapidly increases as the loss aversion parameter becomes larger. Moreover, in both cases the kurtosis stabilises at a value around 2.2 for all $\lambda_{LA} > 1$. However, while for WHP-LA we see an increase of kurtosis of 10%, for the WHP-LA-C model the kurtosis has a 50% increase from 1.40 to 2.20.

Furthermore, in Figure 4.4 (d) we consider the Hill tail index of absolute returns as a function of the loss aversion parameter. This provides a quantitative measure of the heaviness of the tail for the distribution of returns. The Hill index stabilises around a value of 3.60 for a loss aversion parameter $\lambda_{LA} > 1$, experiencing a considerable 18% decrease for the WHP-LA-C model.

As we have demonstrated, loss aversion can have a considerable impact on our behavioural agent-based model. On one hand, if only the chartists are loss averse, even the standard value of $\lambda_{LA} = 2.25$ is powerful enough to drive them out of the market. On the other hand, if both fundamentalists and chartists are biased, the same effect

Table 4.3: WHP vs. WHP-LA model evaluation summary

	WHP model	WHP-LA model
p-value	**31.06**	**32.76**
joint MCR	27.40	28.68
% of joint MCR $\{r_t^{emp}\}$	**52.89**	**55.37**
v mean	83.10	82.64
r AC-1	70.34	70.34
v AC-1	99.30	99.36
v AC-5	66.76	68.74
v AC-10	85.78	86.12
v AC-25	81.76	82.72
v AC-50	79.3	80.78
v AC-100	78.42	78.94
Hill	95.8	95.48

is observed when we employ a sufficiently larger loss aversion parameter ($\lambda_{LA} > 3.10$). Overall, the number of chartists disappearing from the market grows as we increase the magnitude of loss aversion. These results are in line with behavioural finance literature arguing that investors' systematic biases can lead to unexpected market behaviour, with deviations from efficiency and rationality. Moreover, we have shown that loss aversion has an impact on the distribution of returns, increasing their peakiness and the heaviness of the tails. Therefore, while testing for the impact of loss aversion (one of our research aims discussed in Section 1.2), he have found that is has a direct impact on the interactions between agents and the models' price dynamics (as per Section 1.3).

4.4 Model comparison

In this section we are interested in assessing the quality of the statistics (moments) matching achieved by our model with loss aversion (WHP-LA) compared to the same model without the behavioural bias (WHP). After a straightforward estimation has given us the numerical values for the model's parameters (see Section 4.2), we investigate if it would be rejected by real data. This way we demonstrate one of the central objectives of an agent-based financial model, namely how good it is at matching and explaining the properties of real world data.

4.4.1 Moment-specific p-value

We start by exploring the bootstrapped moments of empirical returns, previously used to derive an estimated variance-covariance matrix of the moments (see Equation 3.13). The large sample of $B = 5000$ vectors obtained $(m^1, m^2, ..., m^B)$ can now be used to apply the distance function to them, thus obtaining an entire frequency distribution for J. From this, we can get a critical value of J and evaluate our model's numerical simulations against it. The idea is that such a bootstrapped distribution proxies the one of a hypothetical real-world data generation process. Therefore, if a model-generated return series and its moments have a value of J within the range of the bootstrapped J, we would not be able to differentiate it from a real-world time series. Accordingly, accounting for the rare events with a 5% significance level, a simulated return series will be considered inconsistent with the real-world data if its J-value exceeds the 95% quantile of the bootstrapped J function (denoted by $J_{0.95}$).

Moreover, we cannot consider the minimised J functions from Equation 3.18 since they were obtained over a simulation horizon S that is 10 times longer than the bootstrapped returns. For a more meaningful comparison to their values, the J-values should come from a model simulation of same length $T = 6866$. To this end, we perform a Monte Carlo experiment and repeat the simulations many times, giving rise to a distribution of model-generated J-values that can be compared with the bootstrapped ones.

In more detail, let $m^b = m^b(r_t^{emp})_{t=1}^T$ be the bootstrapped moments given by the empirical returns. Furthermore, let $\hat{\theta}$ be the estimated vector of a model's parameters, giving rise to Monte Carlo moment simulations $m = m^c(\hat{\theta}, T)$, c=1,...,5000. The two frequency distributions of J-values are given by

$$\begin{aligned} Bootstrap: \quad & J[m^b((r_t^{emp})_{t=1}^T)]_{b=1}^{5000} \\ MonteCarlo: \quad & J[m^c(\hat{\theta}, T)]_{c=1}^{5000}. \end{aligned} \tag{4.6}$$

We can now compare the performance of our model with loss aversion (WHP-LA) with the same model without loss aversion (WHP). After obtaining the critical value $J_{0.95}$ of the bootstrapped loss function, we calculate its analogous quantile in the two simulated distributions. While for the WHP model without loss aversion $J_{0.95}$ corresponds to its 31.06% quantile, introducing loss aversion with a coefficient $\lambda_{LA} = 2.25$ in agents' behaviour leads to a higher 32.76% quantile (Table 4.3). For a clear summary of the general moment matching we can say that the model has a p-value of 32.76%, with respect to parameters $\hat{\theta}$ from Table 4.2, the data of the S&P 500 index, and the specific nine moments that we have chosen. In other words, in almost one third of the simulations over $T = 6866$ it is not possible to reject our model. Moreover, the introduction of loss aversion improves the model, leading to a higher p-value.

4.4.2 Moment coverage ratio

In the previous section the model evaluation was based on the values of the objective function J. Now, we can use a more direct way to assess the degree of moment matching of a model, which also accounts for single moments, developed by Franke and Westerhoff (2012). In order to do this, we derive a confidence interval for the empirical moments

Table 4.4: Empirical moments and their 95% confidence intervals

	v mean	r AC-1	v AC-1	v AC-5	v AC-10	v AC-25	v AC-50	v AC-100	Hill
Measured	0.713	0.019	0.184	0.215	0.156	0.126	0.108	0.070	3.299
Lower bound	0.648	-0.003	0.093	0.150	0.103	0.072	0.030	0.004	3.054
Upper bound	0.788	0.044	0.241	0.239	0.181	0.148	0.107	0.068	3.976

from the bootstrapped ones $(m^1, ..., m^B)$, presented in Table 4.4. Intuitively, a simple qualitative assessment of a given return series is that it cannot be rejected as being incompatible with the data if all of its moments are contained in the confidence intervals.

However, since the sample variability of a single simulation is not sufficient to evaluate the model as a whole, we employ another Monte Carlo estimation. This leads to a coverage ratio defined as the percentage of Monte Carlo runs for which a single moment, or all moments jointly, are contained in their confidence intervals. As before, we are mainly interested in assessing the implications of loss aversion and we compute the moment coverage ratios for the model with loss aversion and the one without it. The central joint coverage ratio and individual ones are presented in Table 4.3.

In order to put this in a more quantitative perspective, consider the events sampling moments from a real world data generating process that fall into their confidence intervals. If we account for the fact that there exists certain dependence between them and only some are independent (say r AC-1, Hill, v Mean, v AC-1 and v AC-100) then a true data generating process would have a joint MCR of $0.95^5 = 77.4\%$. Using this as a conservative benchmark against our Monte Carlo ratio, the model without loss aversion corresponds to an effective MCR of $27.4/77.4 = 35.40\%$, and the model with loss aversion to an effective MCR of 37.05%.

We should note that all these statements are based on asymptotic theory (Franke and Westerhoff, 2012). It is unclear if the confidence intervals we derived are a good representation of the ones appropriate for small samples of a true data generating process. Under this constraint, we have used the bootstrapped empirical returns to get more information about the small sample properties. Once more, taking the 5000 bootstrap samples, we compute the joint MCR for them and use this as a benchmark against which we can measure our models' ratios (in percentage points). The result is a ratio of 51.80%, which is considerably lower than the hypothetical reference of 77.4%. On the other hand, the 51.80% is considerably higher than the MCRs of the models, as expected. Now, the most obvious way to relate our models' MCR to that of the bootstrap is to express them as a fraction of it. This is done in the third row of Table 4.4. Referring to these statistics, we see that the model with loss aversion reaches a level of 55.37%, higher than the 52.89% of the same model without loss aversion. Nonetheless, this measure is probably an overestimation of a more appropriate relative coverage statistic and should not be taken too literally. In any case, irrespective of the measure we refer to, the degree

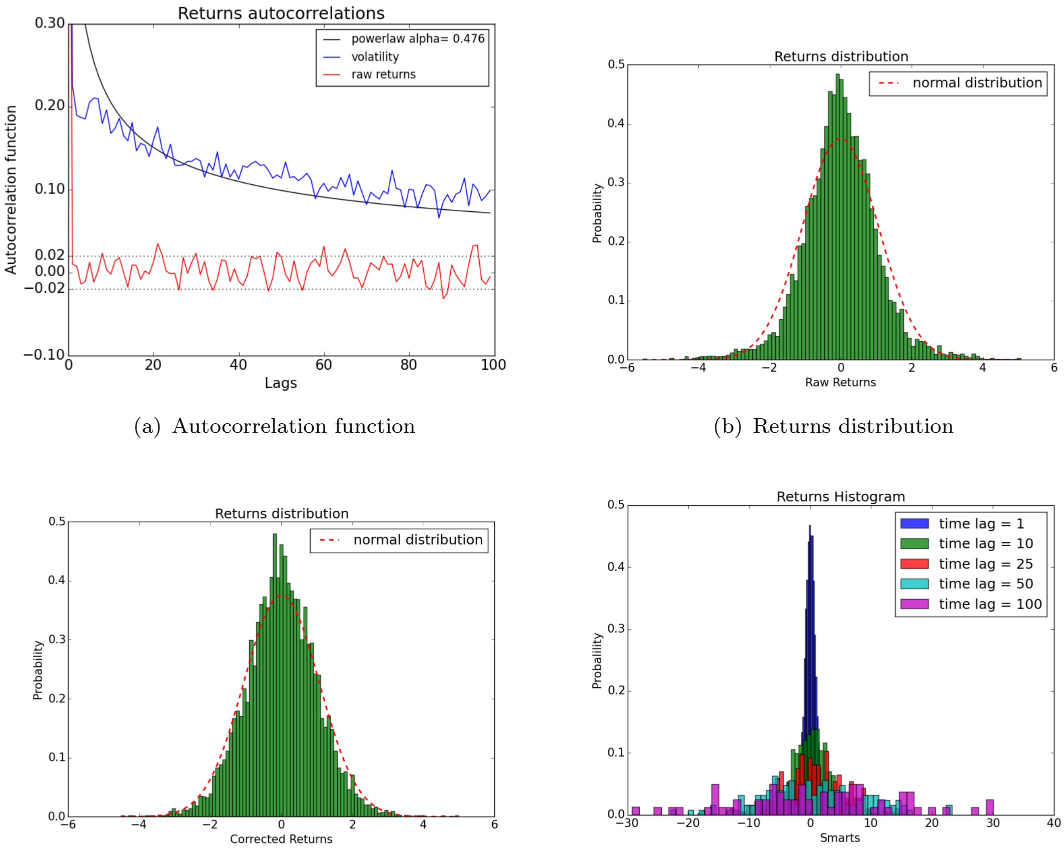

(a) Autocorrelation function

(b) Returns distribution

(c) Volatility corrected returns distribution

(d) Aggregate Gaussianity

Figure 4.5: WHP-LA returns' stylized facts

of the model's ability to match the statistics of real data is remarkable, with a plus for the one with loss aversion (as outlined in Section 1.3).

4.5 Stylized Facts

In this section we explore the statistical properties generated by our WHP-LA model, keeping the loss aversion parameter fixed at the usual value $\lambda = 2.25$. We show that our behavioural model matches a rich set of empirically observed stylized facts, demonstrating one of the main objectives of ABM (see Section 1.2). An in-depth analysis of the time series generated by the model will be performed. First, we demonstrate the existence of the most discussed properties, namely the absence of autocorrelations, heavy tails, volatility clustering and long memory. Moreover, our model is able to match some of the less discussed stylized facts such as conditional heavy tails, aggregate Gaussianity, gain-loss asymmetry, volume power-law, long memory and volume-volatility relations.

4.5.1 Absence of autocorrelations, volatility clustering and long memory

The lack of predictability is one of the most distinctive characteristics of financial time series, closely linked to its non-stationarity (see Section 2.5.1). Stock prices have probability distributions whose mean and variance change over time. This has motivated us to change the behaviour of the fundamental values to a GBM, one of our contributions (Section 3.1.3), such that the prices generated by our agent-based model are non-stationary.

In order to test the lack of predictability of the behavioural model price data, we apply the Augmented Dickey Fuller unit root test. Over 10000 different simulations, the test was rejected 1229 times with the p-value less than the critical value at 10% and 645 times with the critical value at 5%. Thus, we can say the price series generated by our behavioural model are non-stationary in more than 93% of the simulations (with 95% confidence interval).

A further, well known property of financial data states that prices are not autocorrelated (see Section 2.5.2). The prices autocorrelation function decays very sharply and is usually close to zero. This absence of autocorrelation can be easily observed by plotting the ACF of raw returns at lags from 1 to 100 in Figure 4.5 (a). We compare the ACF of simulated and empirical returns, showing that it becomes close to zero after the first lag and clearly demonstrating that price changes are not autocorrelated.

However, the absence of autocorrelations does not rule out the possibility of non-linear dependencies of returns, since the absence of serial correlation does not imply independence. A simple visual representation of return series in Figure 4.2 (b) shows that volatility measured as absolute returns is not constant in time. This property is known as heteroscedasticity. When analysing financial time series, volatility measured as non-linear representations of returns exhibit a much higher positive autocorrelation that persists over time. This is a quantitative signature of volatility clustering: large price variations are more likely to be followed by large price variations (see Section 2.5.4). Usually, the presence of volatility clustering is confirmed by considering the autocorrelation function of volatility, often expressed as the absolute returns (Chen et al., 2012). In the ACF of absolute returns plotted in Figure 4.5 (a), we observe a positive autocorrelation that persists over time, doubled by its slow decay. This clearly indicates the presence of volatility clustering.

Note that the spikes in volatility are a consequence of chartism dominance periods. Accordingly, if p_t deviates from p_t^f the fundamentalism becomes less appealing, leading to an increase in chartists' market fraction. Furthermore, since they have a greater variability in demand, the level of returns in a chartism dominated period exceeds the ones from a fundamentalism regime. Therefore, it appears that normal sequences of

Table 4.5: Excess kurtosis and skewness as we increase the time lags at which WHP-LA returns are computed.

time lag	1	10	25	50	100
kurtosis	2.32	1.94	1.42	0.95	0.56
skewness	-0.0046	0.0026	0.0349	0.038	0.0763

returns are interrupted by outbursts of increased volatility, when the majority of agents are chartists.

Finally, we can test the presence of volatility clustering by investigating how well a standard GARCH(1,1) model fits the returns. In Figure 4.6 (a) and (b) we plot the residuals time series of our simulated and empirical returns, their autocorrelations and partial autocorrelations and the QQ and Probability plots comparing the distribution of returns with a normal distribution. Both residuals are similar to white noise, having no significant autocorrelation. However, there is obvious conditional heteroskedasticity (conditional volatility) that the model has not captured. The shapes of the QQ and Probability plots indicate that the processes are close to normality but with heavy tails.

A further property closely related to volatility clustering is the actual decay of the autocorrelation function of absolute returns. The long memory effect specifically addresses this decay (see Section 2.5.5). We usually say that a process exhibits long memory if it has a slow decay, similar to a hyperbolic function. One way of observing the decay in the ACF of volatility is by fitting a power law of the form A/r^{α}. We notice that a close fit can be obtained, with an exponent $\alpha = 0.476$ (see Figure 4.4 (a)), in line with empirical studies reviewed in Section 2.5.5

Another widely used method of testing the long memory effect is by using the Hurst exponent, which falls in the range $1/2 < H < 1$ for a long memory process. Computing the Hurst exponent of the volatility, we obtain the values of 0.70. Similarly, the Hurst exponent in empirical absolute returns is equal to 0.69. Therefore, the presence of a long memory process is clearly demonstrated.

4.5.2 Heavy tails, conditional heavy tails and aggregate gaussianity

In this section, we consider the distribution of returns, another challenging topic in the econometrics literature. In Figure 4.4 (b) we plot the distribution of returns and clearly observe the deviations from the normal distribution. Our simulated returns, similarly to the ones of real financial assets, are bell shaped but contain more mass in the peak and the tail than normal. Specifically, an excess kurtosis of 2.35 implies a peakiness bigger than normal and a slow asymptotic decay of the probability density function. This non-normal decay is the so-called heavy tail (see Section 2.5.3). Moreover, we analyse the tail distribution of returns by computing the well-known Hill index of absolute returns.

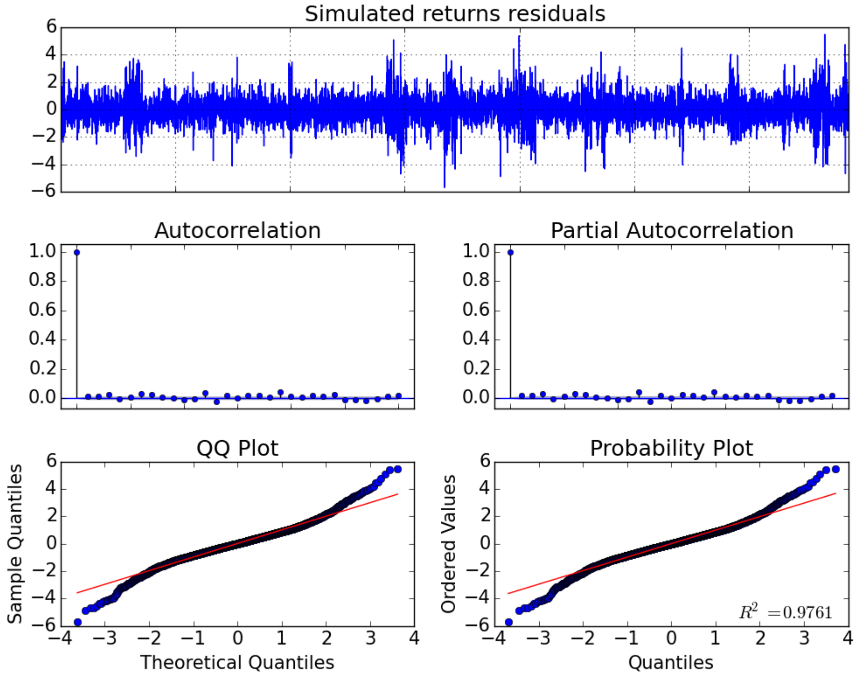

(a) Simulated WHP-LA returns

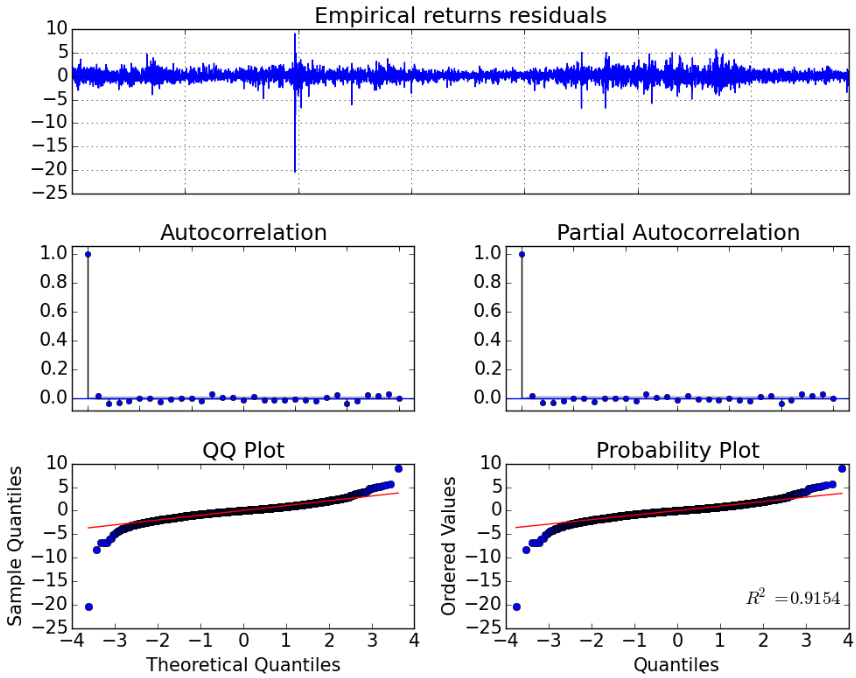

(b) Empirical S&P 500 returns

Figure 4.6: Residuals time series analysis of (a) simulated WHP-LA and (b) S&P 500 empirical data

We obtain a Hill tail index of 3.55, similar to empirical findings that usually report a tail index higher than two and less than five (see Section 2.5.3).

Next, we discuss a further property observed in financial markets that is rarely matched by agent-based models, the so-called conditional heavy tails of returns. It states that even after correcting returns for volatility clustering (e.g. via GARCH-type models), the residuals time series still has heavy tails (Cont, 2001). However, the tails are less heavy than the unconditional distribution. A visual representation of the property can be observed in Figure 4.4 (c), where we plot the distribution of the residual returns after correcting for volatility clustering via a GARCH(1,1) model. Compared with the normal distribution superimposed on it, we see that the residuals of returns have a higher peak and longer tails, confirmed by an excess kurtosis of 0.60 and a Hill tail index of 5.06. At the same time, the tail is less heavy than the one of raw returns.

We have clearly demonstrated the deviations from Gaussianity in returns' distribution. As we have seen, the distribution of returns has excessive kurtosis and heavy tails, with higher peak than the normal distribution and power law tails. However, this behaviour changes as we increase the time scale over which returns are calculated. Specifically, the distribution of returns changes its shape, becoming more and more like the normal distribution. This stylized fact, known as the aggregate Gaussianity (see Section 2.5.7), can be easily observed in different markets and time periods.

In more detail, the aggregate Gaussianity of the simulated returns generated by our model is clearly visible in Figure 4.4 (d). We plot the distribution of returns calculated for different time lags ($\tau = 1, 10, 25, 50, 100$). We can observe how the distribution of returns becomes more like the normal distribution as we increase the time lags. The peaks of the distributions decrease and the tails become wider. Specifically, when returns are calculated at small time scales, their distribution is extremely peaked in the middle with very heavy tails. However, as we increase the time scale, the return distributions become more flat, looking like the normal distribution. Finally, as can be seen from Table 2, the excess kurtosis clearly decreases.

4.5.3 Gain-loss asymmetry

In order to obtain a deeper understanding of the fluctuations of prices and returns, Simonsen et al. (2002); Jensen et al. (2003) proposed a different kind of approach involving inverse statistics. In the analysis of financial data, the inverse question can be formulated as follows: For a given return on an investment, what is the typical time span needed to obtain this return. Specifically, the gain/loss asymmetry of returns states that while the maximum of inverse statistics for a positive level of returns occurs at a specific time, the maximum for the same negative level of returns appears earlier. For

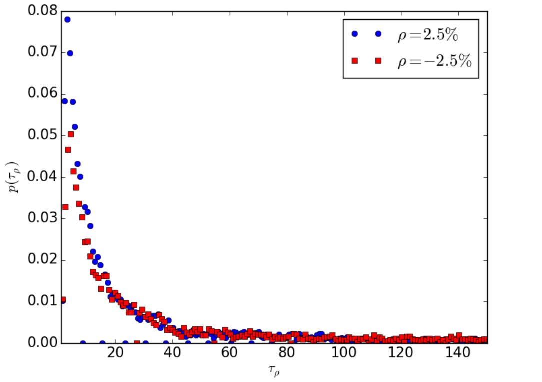

(a) Investment horizon distribtion for a return level $|\rho| =$ 2.5%. The insert is the optimal investment horizon τ_ρ^* for positive and negative levels of returns

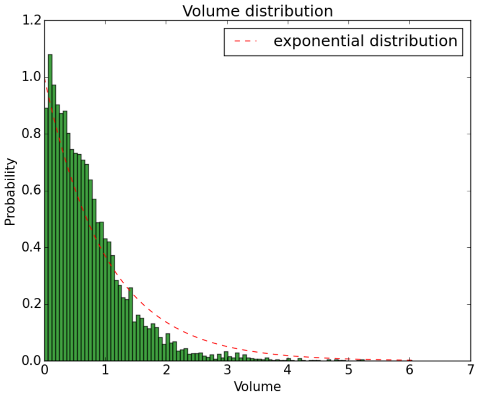

(b) Volume distribution

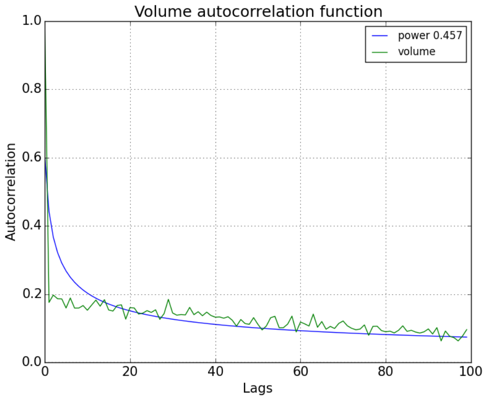

(c) Volume autocorrelation function

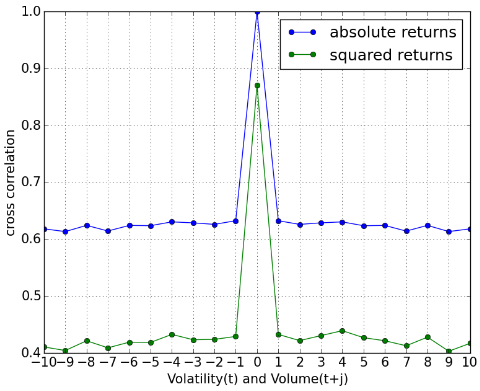

(d) Volume-volatility relations. Cross correlation between volatility measured as absolute retuns (blue) and squared returns (green) and volume distribution.

Figure 4.7: WHP-LA gain/loss asymmetry and volume related stylized facts

the first time in agent-based modelling, we will match this stylized fact with a 2-type design model.

In more detail, the level of return is $\pm\rho$, where the positive or negative sign corresponds to a gain or a loss, respectively. It is kept fixed and one looks for the shortest waiting time after t, $\tau_{\pm\rho}(t)$, for which the return is above or below the predefined $\pm\rho$. The distribution of waiting times for gains $\tau_{+\rho}$, and loses $\tau_{-\rho}$, is denoted $p(\tau_{\pm\rho})$ and corresponds to a fixed return level ρ. Therefore, the maximum of this distribution, τ_ρ^*, is the most probable time of production a return ρ, known as the optimal investment horizon.

In Figure 4.7 (a) we plot the probability distribution function, $p(\tau_{\pm\rho})$, of waiting times $\tau_{\pm\rho}$, with a return level $\rho = 2.5\%$. We can observe the asymmetry between simulated investment horizon distributions. In particular, for a negative level of return, there is a higher probability to short investment horizons, as compared to what is observed for the positive one. On average, while the optimal investment time for a negative return is 20 days, for the same level of positive returns we obtain an optimal investment time of 23 days.

4.5.4 Volume power-law, long memory and volume-volatility relations

We now turn our attention to volume, yet another highly discussed financial topic. Some of the empirical quantitative properties related to volume include a power-law behaviour, long memory and correlation to volatility. We will demonstrate all these stylized facts with our behavioural model. We define the volume at time t as the total absolute demand at t of both fundamentalists and chartists. In Figure 4.7 (b) we can see that the volume distribution has a bigger decay than the exponential distribution superimposed on it, indicating a power-law behaviour. Moreover, an excess kurtosis of 6.036 reinforces this behaviour.

In Figure 4.7 (c) we plot the autocorrelation function of volume. We observe a power-law decay with an exponent 0.457. The slow decay is a clear presence of long-memory of volume. Another widely used method for testing the long memory effect is by using the Hurst exponent. Computing the Hurst exponent of the simulated distribution of volume, we obtain the value of 0.68. Most empirical studies report a Hurst exponent in the interval $[0.5; 1]$ for long memory processes. Therefore, the presence of long memory in volume is clearly demonstrated. Finally, the dependence between volatility and volume traded has been noticed and documented across different financial instruments at different time scales. Specifically, the relationship between the two market descriptors can be observed via the cross correlation function. In Figure 4.7 (d) we observe a significantly positive cross correlation between volatility and volumes. This means that a small (large) trading volume is accompanied by a small (large) change in volatility. It is important to note that all measures of volatility are positively correlated with volume and that the dependence remains significant as we increase the time lag.

4.5.5 Price impact and extreme price events

The price impact addresses a different cross correlation, namely between volume and price changes. Specifically, we are interested in the impact of volume traded on the prices. Although there are different theories regarding this relation, the impact of a single order has been widely found to be a concave function of volume (see Section 2.5.8). In Figure 4.8 (a) we plot the simulated price changes as a function of volume, clearly obtaining a concave function of price impact. This behaviour is expected from the model definition, since the market maker computes the next price as a function of the agents' demand, multiplied by their market fraction. Clearly, if the absolute demands increase, the next price will also increase, but non-linearly.

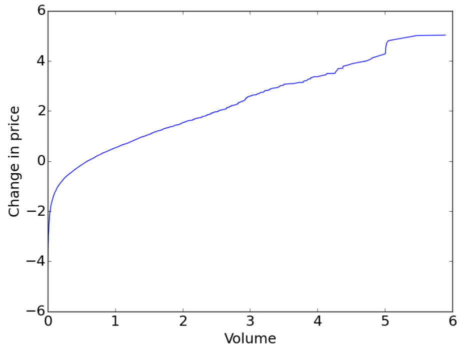

(a) Price impact

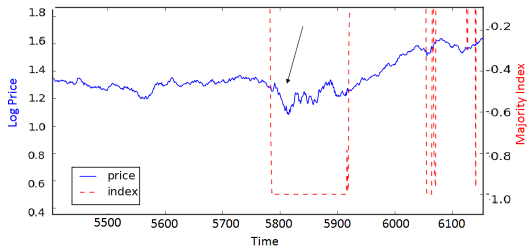

(b) Price series: crash from $t = 5799$ to $t = 5809$.

Figure 4.8: WHP-LA (a) Price impact function and (b) An extreme price event.

Finally, we tackle one of the recently observed anomalies in real financial markets, known as extreme price events. In Figure 4.8 (b) we present an example of crash in the market occurring between periods $t = 5799$ and 5809. We observe a drop of 10.52% from initial price at $t = 5799$. Unsurprisingly, the crash occurs in a period of switching from fundamentalists to chartists domination (majority index x_t close to -1).

4.5.6 Comparison with the model without loss aversion

In this section we perform a comparison of the statistics obtained by our behavioural agent-based model (WHP-LA) with the ones of the FW model (WHP). To this end, we have demonstrated that one of the main differences between the reference model and the new behavioural model with loss aversion relates to the movement of the majority index (see Section 4.3.2).

For a more robust comparison, we analyse some the most discussed stylized facts for each of the two models, which focus mainly on the returns' dynamics. Figure 4.9 plots the returns time series for WHP and WHP-LA, together with their distributions, autocorrelation functions, QQ plots and the complementary cumulative distribution functions.

First, we can easily observe the lack of autocorrelation of raw returns for both WHP and WHP-LA. Moreover, the volatility clustering can be seen in the returns time series, where low volatility is more likely to be followed by low volatility and high volatility is more likely to be followed by high volatility. A further quantitative measure of the well known volatility clustering is represented by the persistence of positive autocorrelations in absolute returns. However, fitting a power-law of the form A/r^{α}, we see that the WHP-LA model has an autocorrelation of absolute returns that persists longer. Moreover, the slow decay demonstrates the presence of long memory for both processes.

In addition, the two distributions of returns have higher peaks and heavier tails than the normal distribution. While the WHP model has an average excess kurtosis of 2.20

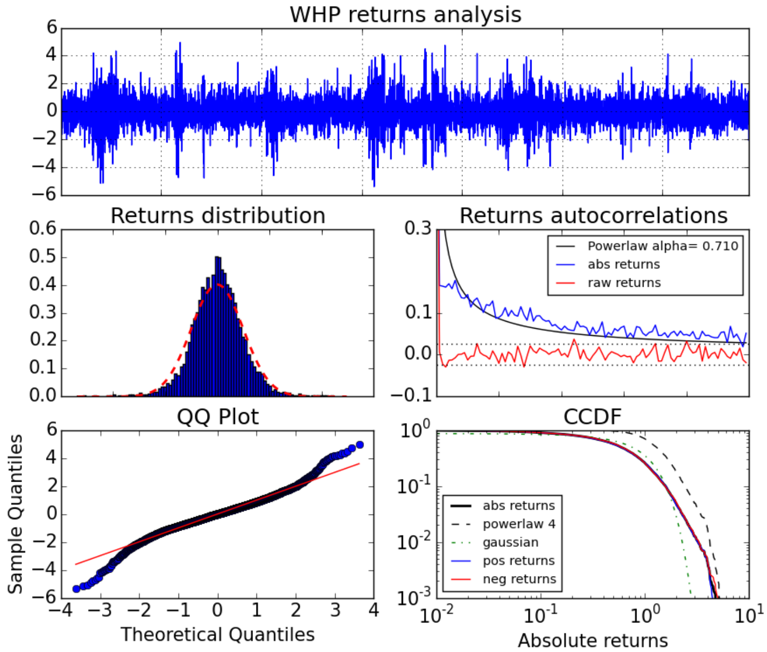

(a) WHP returns

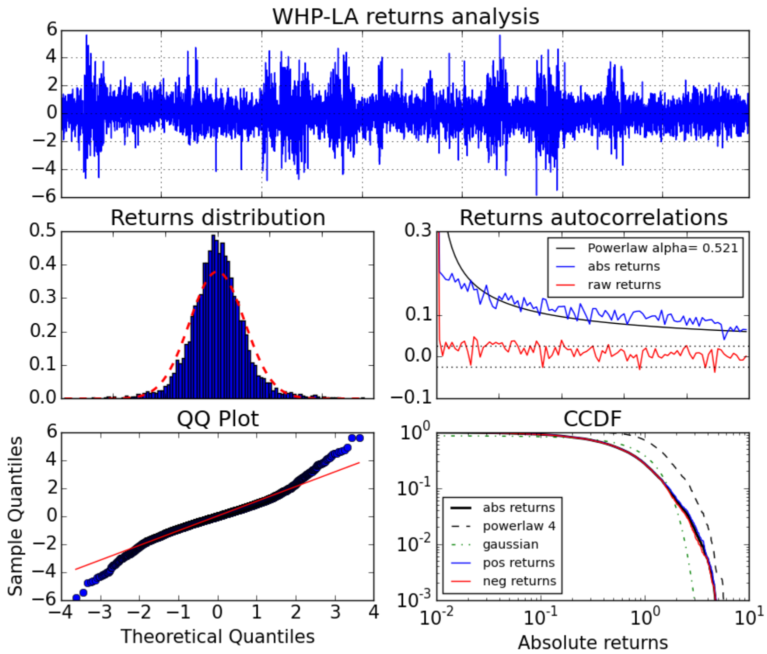

(b) WHP-LA returns

Figure 4.9: WHP vs. WHP-LA returns dynamics

and a Hill index of 3.60, the loss aversion in the WHP-LA model leads to an even more deformed distribution, with an excess kurtosis of 2.33 and a Hill index of 3.5. This behaviour is also confirmed by the QQ plots, where the elongated S shapes indicate excess kurtosis in the returns' distribution and heavy tails.

It is clear that the distributions of returns are not Gaussian. In reality, large fluctuations are more likely to happen in a stock market and dramatic crashes are observed more frequently. In order to characterise the probability of such events, we look at the complementary cumulative distribution function (CCDF), $F(x) = 1 - Prob(X < x)$, which describes the tails of the distribution $P(x)$ of returns. The CCDF of returns is usually found to be approximately a power-law with an exponent in the range 2-4 (Bouchaud and Potters, 2003). A slower decay than Gaussian reveals the presence of heavy tails. The distributions are also asymmetric, having different exponents characterizing their right tails (positive returns) and left tails (negative returns).

Comparing the most important stylized facts of the two models leads to no significant differences. While the FW model is able to replicate some of the properties of real financial data, the introduction of loss averse agents in the WHP-LA model mainly changes the interactions between market participants. In terms of the statistics, the loss aversion leads to a more pronounced volatility clustering and a distribution of returns with higher peaks and heavier tails. This comes as a consequence of the interaction between biased agents. As we have seen in Section 4.3.2, the biased agents tend to be driven out of the market. Therefore, the WHP-LA setting brings more dominant fundamentalists and more abrupt regime changes. These periods of alternating fundamentalism and chartism dominance lead to the disruptions we see in the returns series.

4.6 Summary

Heterogeneous agent-based models that rely on simple trading strategies have proven themselves very efficient in generating important dynamics of real financial markets. Indeed, the fundamentalists vs. chartists models have been shown to successfully capture empirically observed traders' behaviour. Moreover, a well-defined agent-based model able to match the financial markets' common properties is an important tool for testing various behavioural and economical theories and understanding how they influence the interactions between agents and prices (as per our research aims discussed in Section 1.2).

In this chapter we described, evaluated and extended one of the most successful approaches in capturing real-life market dynamics (see Chapter 3). We presented two new behavioural asset pricing model. First, we define a setting where agents suffer from the disposition effect and test the implications of this behavioural bias on the interactions between agents and the time series generated by the model. Our work leads to two important findings. On the one hand, we show that the horizon over which the agents consider alternative scenarios, an important concept of prospect theory, has no impact on the setting dynamics. This is surprising because in prospect theory it does have a decisive role in agents' interactions.

On the other hand, the level of the disposition effect directly impacts the returns series produced by the model, altering its stylized facts and leading to disruptive behaviour. Specifically, we demonstrate how these disruptions can appear and show that the disposition effect has a direct impact on the returns' distribution, its heavy tails and volatility clustering. This provides support for the behavioural finance literature arguing that investors systematic biases can lead to unexpected market behaviour, with deviations from efficiency and rationality (as outlined in Section 1.3).

Second, we presented a new behavioural model of asset pricing where the agents are loss averse. In the new setting, we observed major differences in the movements of the agents' market fractions compared to unbiased scenarios. Specifically, for certain levels of loss aversion, the biased chartists tend to be driven out of the market at different points in time (as per our research contributions discussed in Section 1.3). Additionally, by comparing the statistics of the base model without loss aversion and our new behavioural model, we have demonstrated how certain levels of this bias change the distribution of returns, leading to a more pronounced volatility clustering and a distribution of returns with higher peaks and heavier tails. Similar to the impact of the disposition effect, we can link our findings with the behavioural finance literature, which states that investors' systematic biases lead to surprising market behaviour, instabilities and errors.

Furthermore, one the central objectives of agent-based financial modelling is to propose an alternative to the apparent randomness of financial markets, trying to explain the most important properties of financial data (for a more detailed discussion on our research aims see Section 1.2). Specifically, we are interested in simple structures that can reproduce the empirical findings to a high degree and which are quantitatively close to the real ones. Trying to measure how close a model is to its empirical counterparts, we make use of two different statistics. First, the distance between their summary statistics (moments) is measured by a quadratic loss function. For the second, we used the concept of (joint) moment coverage ratio. Here, with the help of Monte Carlo runs, we count the number of times where all the simulated statistics are contained in the empirical confidence intervals. This can be seen as a quantitative measure of how may often the data from a model and the real market cannot be told apart. With a p-value of 31.06% and a coverage ratio of 55.37%, the introduction of loss aversion clearly improves the degree to which the model explains the real data, compared to the same model without loss aversion (see our research contributions discussed in Section 1.3).

We offer both a quantitative and a qualitative analysis of the simulated asset price series, its returns, volume traded and volatility. We provide a wide range of tests and arguments and demonstrate the presence of a rich set of stylized facts including absence of autocorrelations, heavy tails, conditional heavy tails, volatility clustering, long memory of returns, aggregate Gaussianity, gain-loss asymmetry, volume power law, long memory and volume-volatility relations. By doing so, our behavioural model is the first one to match such a rich set of the stylized facts.

Chapter 5

Regret in an Agent-Based Model of Asset Pricing

The main focus so far has been on two key objective of agent-based financial modelling (as per Section 1.2). On one hand, in Chapter 3 we were primarily interested in recreating the statistical properties observed in real financial markets, known as stylized facts. Moreover, a well defined agent-based asset pricing model able to match the widely observed properties of financial time series is valuable for testing the implications of various biases associated with investors' behaviour. By modelling different empirically observed behavioural biases in agents' decision rules, we have shown that the new setting can provide useful insight not only on the behaviour of individual agents, but also on the effects that emerge from their interactions (see Chapter 4).

However, a further major focus of heterogeneous agent models that are able to match the properties of real financial markets is on designing the model's mechanisms and offering an explanation for the observed market behaviour (as discussed in Section 1.2). The development of various models that are capable of recreating important stylized facts and the difficulty of comparing them, has resulted in a number of economical explanations for generating the interesting dynamics. In particular, Gaunersdorfer et al. (2008) consider the widely used adaptive beliefs system (see Section 2.3.2) and show that volatility clustering, one of the most common properties of financial returns can be characterised by a coexistence of stable attractors.

In this chapter, we depart from studying the stylized facts generated by an agent-based model (see Chapters 3 and 4) and rather focus on the deterministic mechanism responsible for their formation. Against this background, we define a new model that incorporates regret, a well-known psychological model (see Section 2.4.3 for more details), in agents behaviour and examine its implications on the overall market. It is important because we believe that for a better representation of how investors and traders behave, agent-based financial models should capture the behavioural factors observed in real markets (LeBaron, 2006). Specifically, as the agents try to model real life traders they should represent their behavioural biases to some extent.

In financial markets, regret is experienced when an investor's decision yields, ex post, a worse performance than an alternative one he could have chosen. In addition, the prospect of future regret and self-blame may influence investors into making suboptimal, non-rational decisions relative to their utility models.

In this vein, in this chapter we propose a simple heterogeneous agent-based model with modified expectations that account for regret. Specifically, an individual who experiences regret is expected to anticipate the feeling and take it into consideration when making a decision under uncertainty (Loomes and Sugden, 1982). In other words, rational investors make their optimal decision by maximising a modified expected utility, trying to account for past regret accordingly.

In order to understand the dynamics of our model we use stability and bifurcation theory to find the implications of this behavioural bias on the model's deterministic skeleton. In particular, we focus on the model's stability region and its coexistence of attractors, one of the most common mechanisms responsible for generating important properties in nonlinear agent-based financial models.

Focusing on this generic mechanism of volatility clustering (one of our research aims outlined in Section 1.2), we show that an increase in regret can dramatically affect the market dynamics. In detail, we demonstrate how this behavioural bias has the power to destabilise the market and transform a low volatility regime characterized by a stable steady state into a highly volatile one characterized by a stable limit circle. Since the interplay between these coexisting attractors leads to important properties observed in financial markets, we conclude that regret has the potential to explain some key aspects of financial markets. Consequently, a change in agents' psychology can destabilise the market and endogenously alter its volatility regime (for a discussion on our research contributions see Section 1.3).

The rest of the chapter is organised as follows. We formally define the agent-based model in Section 5.1, with a focus on incorporating regret into agents' expectation. Section 5.2 reduces the underlying deterministic model to a 6-dimensional nonlinear discrete-time dynamical system, whose unique steady state and local stability conditions are explained. Next, in Section 5.3 we perform bifurcation analysis with respect to the key parameters of regret through numerical experiments and explore the destabilising effect of this behavioural bias.

5.1 Model definition

The model described in this section is based on a standard discounted value asset pricing model with heterogeneous agents, closely related to the well known frameworks discussed in Section 2.1 and the adaptive beliefs system described in Section 2.3.2. More specifically, we focus on a setting similar to the one considered by Gaunersdorfer et al. (2008) and He et al. (2016), in relation to the mechanism responsible for generating important stylized facts. In particular, we follow the model of Dieci et al. (2006), which allows agents to either have fixed or adaptive strategies and was further calibrated on empirical data, showing that it explains an extensive set of financial properties (He and Li, 2017). In Section 5.1.1, we specify the model's price adjustments, as they were set by Dieci et al. (2006). Next, we extend the model to define the agents' expectations such that it incorporates regret in Section 5.1.2 and we present the complete model in 5.1.3.

5.1.1 Price dynamics

Consider the classical setting of an asset pricing model with one risky and one risk-free asset. It is assumed that the risk free asset is perfectly elastic and supplied at a gross return $R = 1 + r/K$, where r is the constant risk-free rate per annum and K is the trading frequency, measured in fractions of a year. Since our focus is on the statistics observed from daily price movements, we fix $K = 250$. Let p_t and y_t $(t = 0, 1, 2, ...)$ be the (ex dividend) price and stochastic dividend process of the risky asset, respectively. Given this, denote by $R_{t+1} = p_{t+1} + y_{t+1} - Rp_t$ the excess capital gain per share made on the risky asset at $t + 1$. It follows that the wealth of investor of type h at $t + 1$ is given by:

$$W_{h,t+1} = RW_{h,t} + (p_{t+1} + y_{t+1} - Rp_t)z_{h,t}, \tag{5.1}$$

where $W_{h,t}$ is the investor's wealth at time t and $z_{h,t}$ represents the number of shares of the risky asset (demand) held by investor type h, for one time period from t to $t + 1$. Note that by type h we refer to a subclass of investors all following the same particular strategy.

Let $E_{h,t}$ and $V_{h,t}$ be the beliefs of type h investors representing the conditional expectations and variance at $t + 1$ based of the available information they have at time t. As it is common in the literature, it is assumed that traders maximize their corresponding constant absolute risk aversion (CARA) utility wealth function $U_h(W) = -\exp(a_h W)$, where $a_h > 0$ is the risk aversion coefficient of type h traders. The optimal demand of a type h trader on the risky asset at time t is then given by:

$$z_{h,t} = \frac{E_{h,t}(R_{t+1})}{a_h V_{h,t}(R_{t+1})}. \tag{5.2}$$

As per most of the literature (see Section 1.1), it is assumed that the market is populated by two types of traders, fundamental traders (fundamentalists) and trend followers (chartists), denoted by type 1 and 2, respectively. Moreover, each class consists of adaptively rational agents that select different strategies over time according to a performance measure, and confident agents that do not switch and stay with their strategies over time. Specifically, let $q_{i,t} \in [0,1]$ $(i = 1,2)$ (replacing previously used $i = f, c$) be the agents' market fractions at time t, having both fixed and varying components (Dieci et al., 2006). Let $n_1 \in [0,1]$ and $n_2 \in [0,1]$ be their fixed proportions and among the $1-(n_1+n_2)$ switching agents, denote $n_{1,t} \in [0,1]$ and $n_{2,t} = 1 - n_{1,t} \in [0,1]$ as the varying proportions of fundamentalists and chartists at time t, respectively. Let $n_0 = n_1 + n_2$, $m_0 = (n_1 - n_2)/n_0$ and $m_t = n_{1,t} - n_{2,t}$. Combining the two components, it follows that the overall market fractions can be expressed as:

$$\begin{aligned} q_{1,t} &= \frac{1}{2}[n_0(1+m_0) + (1-n_0)(1+m_t)], \\ q_{2,t} &= \frac{1}{2}[n_0(1-m_0) + (1-n_0)(1-m_t)]. \end{aligned} \tag{5.3}$$

With no supply of outside shares, the total weighted average excess demand $z_{e,t}$ at time t is given by $z_{e,t} \equiv q_{1,t}z_{1,t} + q_{2,t}z_{2,t}$. Following some prominent examples in the literature, the market clearing price is determined by a market maker (see Section 2.1), that takes a long position in the presence of excess demand $(z_{e,t} < 0)$ and a short position in the presence of excess supply $(z_{e,t} > 0)$. The market price is adjusted such that $p_{t+1} = p_t + \mu z_{e,t} + \epsilon_t$, where $\mu > 0$ denotes the speed of adjustment of the market maker and the term $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ is an i.i.d. random disturbance accounting for market noise, independent of p_t. The price equation becomes:

$$p_{t+1} = p_t + \frac{\mu}{2}\left\{ q_{1,t}\frac{E_{1,t}[R_{t+1}]}{a_1 V_{1,t}[R_{t+1}]} + q_{2,t}\frac{E_{2,t}[R_{t+1}]}{a_2 V_{2,t}[R_{t+1}]} \right\} + \epsilon_t, \tag{5.4}$$

where $q_{1,t}$ and $q_{2,t}$ are given by Equation 5.3.

Next, we define the agents' expectations about future price movements. The first class of investors, fundamentalists, act as a stabilizing force in the market. They believe in the existence of a fundamental price and invest relatively to its value. The fundamental value of the risky asset, p_t^f, is assumed to follow a random walk:

$$p_{t+1}^f = p_t^f + \eta_t, \quad \eta_t \sim N(0, \sigma_1^2), \tag{5.5}$$

where η_t is independent of p_t^f and p_t. Fundamentalist believe the price series will eventually move or converge towards this long term equilibrium. Specifically, the conditional mean and variance of the fundamental traders are given by:

$$\begin{aligned} E_{1,t}[p_{t+1}] &= p_t + (1-\alpha)(p^f_{t+1} - p_t), \\ V_{1,t}[p_{t+1}] &= \sigma_1^2, \end{aligned} \tag{5.6}$$

where $(1-\alpha)$ denotes the speed of the fundamentalists' price adjustments, with $\alpha \in (0,1)$, and σ_1^2 is a constant variance. Note that a high (low) value of α represents a slow (quick) adjustment of the expected price towards its perceived fundamental value.

The second class of agents, technical analysts or chartists, forecast the future prices entirely by modelling historical data. They do not take into consideration the market fundamentals and base their decisions entirely on observed historical patterns in past prices. In more detail, trend followers extrapolate recent price changes and adjust their beliefs accordingly. Specifically, their conditional mean and variance are given by:

$$\begin{aligned} E_{2,t}[p_{t+1}] &= p_t + \gamma(p_t - u_t) \\ V_{2,t}[p_{t+1}] &= \sigma_1^2 + b_2 v_t, \end{aligned} \tag{5.7}$$

where $\gamma \geq 0$ quantifies their extrapolation rate, u_t and v_t are the sample mean and variance of some learning process, with $b_2 \geq 0$. Here, high (low) values of γ correspond to strong (weak) extrapolations from trend followers. Different learning schemes can be used to calculate the sample mean and variance (see e.g. Chiarella and He (2002, 003a)). We follow He and Li (2017) and assume that:

$$\begin{aligned} u_t &= \delta u_{t-1} + (1-\delta)P_t, \\ v_t &= \delta v_{t-1} + \delta(1-\delta)(P_t - u_{t-1})^2, \end{aligned} \tag{5.8}$$

where $\delta \in (0,1)$. This is equivalent to a limiting geometric decay process when the memory lag tends to infinity (Chiarella et al., 2006). In practice, traders tend to put more weight on the most recent prices when they estimate the sample mean and variance. Moreover, this choice of learning process makes it mathematically tractable.

5.1.2 Expectations with regret

We now turn our attention towards modelling regret (see Section 2.4.3) and incorporating it in agents' heterogeneous beliefs. Here, we depart from the classical setting of fundamentalists vs. chartists models (defined in the previous section) and modify the agents' expectations such that they include the disutility from having chosen an ex-post suboptimal alternative. As we will demonstrate, this psychological factor has major implications on the agents' interactions and can change the model's locally stable attractors.

In financial markets, the idea of regret has a natural application since investors usually get their performance compared against exogenous benchmarks or other market participants (Bleichrodt and Wakker, 2015). Therefore, it is natural for an investor to feel disappointed if their own portfolio underperformed compared to other alternatives and that can affect future performance. The main assumption of regret theory is that people may feel regret after making their decisions under uncertainty if those decisions prove to be wrong or suboptimal, even if they appeared to be correct with the available information ex ante. In this case, an investor's future expectations should also depend on the realisation of not having chosen better alternatives.

To this end, we model regret such that investors' decisions are influenced by the comparison with other market participants, as this is more realistic portrait of the real world. Therefore, agents feel regret towards a level of wealth they could have obtained from a foregone alternative and modify their future expectations accordingly. Note that in our agent based model investors can only be one of two types, fundamentalists or chartists, so they can only feel regret towards the other type's profits.

Let $\pi_{h,t} = R_t z_{h,t-1} = (p_t + y_t - Rp_{t-1})z_{h,t-1}$ be the excess realised return (realised profit or gain) of agent type $h = 1, 2$ at time t based on the demand submitted at time $t-1$, $t = 1, 2....$ Intuitively, we differentiate between two different cases of experiencing regret at time t, based on the latest movements of demands and excess returns (from $t-1$ to t):

1. If the risky asset had a positive excess capital gain per share ($R_t > 0$), then the agents with the largest previous demand realised the biggest profit ($\pi_{h,t} = R_t z_{h,t-1}$). Conversely, the agents with a smaller demand regret not having bought more shares at time $t-1$. Assuming $z_{1,t-1} \geq z_{2,t-1}$, it follows that agent type 1 made a bigger profit than agent type 2 in the latest trading period, $\pi_{1,t} \geq \pi_{2,t}$. Accordingly, agent type 2 regrets not having submitted a bigger demand at time $t-1$. Agent type 1 has no regret.

2. If the risky asset had a negative excess capital gain per share ($R_t < 0$), then the agents with the smallest previous demand realised the biggest profit ($\pi_{h,t} = R_t z_{h,t-1}$). Conversely, the agents with a larger demand regret not having bought less shares at time $t-1$. Assuming $z_{1,t-1} \geq z_{2,t-1}$, it follows that agent type 2 made a bigger profit (or smaller loss) than agent type 1 in the latest trading period, $\pi_{1,t} \leq \pi_{2,t}$. Accordingly, agent type 1 regrets not having submitted a smaller demand at time $t-1$. Agent type 2 has no regret.

Therefore, by comparing their performance with the other market participants ex-post, some agents feel the regret of not having had different demands ex-ante. Moreover, regret has such a strong impact on investors' psychology that they will try to avoid it in the future by deviating from their usual strategies. In other words, after feeling the regret of realising less profits than others, an agent will modify its expectations about future price movements by altering them in the direction set by regret.

Formally, we define the regret of trader type 1, compared to type 2, at time t as:

$$\mathrm{Rgrt}_{1,2;t} = \mathrm{sgn}(R_t)(\pi_{2,t} - \pi_{1,t})^+, \tag{5.9}$$

where $(\pi_{2,t} - \pi_{1,t})^+ = \max(0, \pi_{2,t} - \pi_{1,t})$ and $\mathrm{sgn}(R_t) = -1$ if $R_t < 0$, 0 if $R_t = 0$ and 1 if $R_t > 0$. $\mathrm{Rgrt}_{2,1;t}$ can be defined in a similar way. Note that we define regret using the non-differentiable functions sgn and max. However, in order to study the impact of regret on the overall dynamics of the model using stability and bifurcation analysis, we are required to compute the system's derivatives at the fundamental equilibrium. In this context, we can approximate the non-differentiable regret function defined in Equation 8 as:

$$\mathrm{Rgrt}^k_{1,2;t} = \frac{(\pi_{2,t} - \pi_{1,t})}{1 + \exp[-k(\pi_{2,t} - \pi_{1,t})]} \tanh(kR_t). \tag{5.10}$$

Note that $\mathrm{Rgrt}^k_{1,2;t} \to \mathrm{Rgrt}_{1,2;t}$ as $k \to \infty$ since $\tanh(kR_t) \to \mathrm{sgn}(R_t)$ and $\frac{(\pi_{2,t}-\pi_{1,t})}{1+\exp[-k(\pi_{2,t}-\pi_{1,t})]}$ $\to \max(0, \pi_{2,t} - \pi_{1,t})$ as $k \to \infty$. In addition, we smooth the impact of the wild changes in $\max(0, \pi_{2,t} - \pi_{1,t})$ on regret by considering $\log[\max(0, \pi_{2,t} - \pi_{1,t}) + 1]$. Therefore, our differentiable regret function is given by:

$$\mathrm{Rgrt}^k_{1,2;t} = \log\left[\frac{(\pi_{2,t} - \pi_{1,t})}{1 + \exp[-k(\pi_{2,t} - \pi_{1,t})]} + 1\right] \tanh(kR_t) \tag{5.11}$$

In this context, we incorporate the disutility from having chosen an ex-post suboptimal alternative in agents' future price expectations as:

$$\begin{aligned} \hat{E}_{1,t}[p_{t+1}] &= p_t + (1-\alpha)(p^f_{t+1} - p_t) + \lambda_{\mathrm{Rgrt}}\,\mathrm{Rgrt}^k_{1,2;t}, \\ \hat{E}_{2,t}[p_{t+1}] &= p_t + \gamma(p_t - u_t) + \lambda_{\mathrm{Rgrt}}\,\mathrm{Rgrt}^k_{2,1;t}, \end{aligned} \tag{5.12}$$

where $\lambda_{\mathrm{Rgrt}} \geq 0$ represents the impact of regret at time t and $\hat{E}$ is the expectation with regret. Note that regret makes the biased agents' depart from their rational behaviour, which in this setting is represented by their predefined expectations about future price changes.

5.1.3 The complete model

Following Dieci et al. (2006) we assume that the dividend process follows a normal distribution $y_t \sim N(\bar{y}, \sigma_y^2)$, where the expected long-term fundamental price $\bar{p} = \bar{y}/(R-1)$ and the variances of price and dividends are related by $\sigma_y^2 = r^2\sigma_1^2$. Now, since p_t and y_t are independent, $R_{t+1} = (p_{t+1} + y_{t-1} - Rp_t)$ and using Equations 10 and 11, it follows that the traders' conditional expected excess returns (incorporating regret) and conditional variances can be computed as:

$$\begin{aligned}
E_{1,t}[R_{t+1}] &= (\alpha-1)(p_t - p^f_{t+1}) - (R-1)(p_t - \bar{p}) + \\
&\quad + \lambda_{\mathrm{Rgrt}} \log\left[\frac{(\pi_{2,t} - \pi_{1,t})}{1+\exp[-k(\pi_{2,t} - \pi_{1,t})]} + 1\right] \tanh(kR_t), \\
V_{1,t}[R_{t+1}] &= (1+r^2)\sigma_1^2, \\
E_{2,t}[R_{t+1}] &= \gamma(p_t - u_t) - (R-1)(p_t - \bar{p}) + \\
&\quad + \lambda_{\mathrm{Rgrt}} \log\left[\frac{(\pi_{1,t} - \pi_{2,t})}{1+\exp[-k(\pi_{1,t} - \pi_{2,t})]} + 1\right] \tanh(kR_t), \\
V_{2,t}[R_{t+1}] &= \sigma_1^2(1 + r^2 + bv_t),
\end{aligned} \tag{5.13}$$

where $b = b_2/\sigma_1^2$.

Therefore, the agents' demands are given by:

$$\begin{aligned}
z_{1,t} &= \frac{(\alpha-1)(p_t - p^f_{t+1}) - (R-1)(p_t - \bar{p}) + \lambda_{\mathrm{Rgrt}} \log\left[\frac{(\pi_{2,t}-\pi_{1,t})}{1+\exp[-k(\pi_{2,t}-\pi_{1,t})]} + 1\right]\tanh(kR_t)}{a_1(1+r^2)\sigma_1^2}, \\
z_{2,t} &= \frac{\gamma(p_t - u_t) - (R-1)(p_t - \bar{p}) + \lambda_{\mathrm{Rgrt}} \log\left[\frac{(\pi_{1,t}-\pi_{2,t})}{1+\exp[-k(\pi_{1,t}-\pi_{2,t})]} + 1\right]\tanh(kR_t)}{a_2(1+r^2+bv_t)\sigma_1^2}.
\end{aligned} \tag{5.14}$$

In order to complete the model we need to specify the movement of the adaptive market fractions for the two types of agents. Following the adaptive belief system (see Section 2.3.2), the proportions of switching or adaptive agents of type h at time $t+1$ is determined by a multinomial logit model:

$$n_{h,t+1} = \frac{\exp[\beta(\pi_{h,t+1} - C_h)]}{\sum_i \exp[\beta(\pi_{i,t+1} - C_i)]}, \quad h = 1, 2, \tag{5.15}$$

where $C_h \geq 0$ is a fixed cost associated with using strategy of type h, and β is the intensity of choice measuring the switching sensitivity of the population of adaptive traders to the more profitable strategy. Note that since $m_{t+1} = n_{1,t+1} - n_{2,t+1}$, we have:

$$m_{t+1} = \tanh\left\{\frac{\beta}{2}\left[(\pi_{1,t+1} - \pi_{2,t+1}) - (C_1 - C_2)\right]\right\}. \tag{5.16}$$

To sum up, the model can be expressed as the following random discrete-time dynamic system ($t = 0, 1, ...$):

$$
\begin{aligned}
p_{t+1} &= p_t + \frac{\mu}{2}\{[(n_0(1+m_0) + (1-n_0)(1+m_t)]z_{1,t} + \\
&\quad + [n_0(1-m_0) + (1-n_0)(1-m_t)]z_{2,t}\} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), \\
u_t &= \delta u_{t-1} + (1-\delta)P_t, \\
v_t &= \delta v_{t-1} + \delta(1-\delta)(p_t - u_{t-1})^2, \\
m_t &= \tanh\left\{\frac{\beta}{2}\left[(z_{1,t-1} - z_{2,t-1})(p_t + y_t - Rp_t) - (C_1 - C2)\right]\right\}, \\
y_t &= \bar{y} + \sigma_y \upsilon_t, \quad \upsilon_t \sim N(0,1), \\
z_{1,t} &= \frac{(\alpha-1)(p_t - p^f_{t+1}) - (R-1)(p_t - \bar{p}) + \lambda_{\text{Rgrt}} \log\left[\frac{(\pi_{2,t} - \pi_{1,t})}{1+\exp[-k(\pi_{2,t} - \pi_{1,t})]} + 1\right] \tanh(kR_t)}{a_1(1+r^2)\sigma_1^2}, \\
z_{2,t} &= \frac{\gamma(p_t - u_t) - (R-1)(p_t - \bar{p}) + \lambda_{\text{Rgrt}} \log\left[\frac{(\pi_{1,t} - \pi_{2,t})}{1+\exp[-k(\pi_{1,t} - \pi_{2,t})]} + 1\right] \tanh(kR_t)}{a_2(1+r^2+bv_t)\sigma_1^2}.
\end{aligned}
\tag{5.17}
$$

5.2 Dynamics of the deterministic system

In order to better understand the interactions and inner workings of the model, we study the dynamics of the corresponding deterministic model. Indeed, since we are interested in the mechanism responsible for generating important properties at the macro level, the focus here is on the deterministic skeleton and its particular micro dynamics. Note that the deterministic skeleton of the stochastic model is obtained by setting the noise terms equal to zero, by assuming a constant dividend $\bar{y}$ per time period and a constant fundamental p_t^* that is equal to the long-run fundamental price $p_t^* = \bar{p} = \bar{y}/(R-1)$. Accordingly, we obtain a 6D deterministic dynamical system driven by the following 6D map $T^k : (p, u, v, m, z_1, z_2) \to (p', u', v', m', z_1', z_2')$

$$
T^k := \begin{cases}
P' &= p + \frac{\mu}{2}\left\{[(n_0(1+m_0) + (1-n_0)(1+m)]z_1 + [n_0(1-m_0) + (1-n_0)(1-m)]z_2\right\}, \\
u' &= \delta u + (1-\delta)p', \\
v' &= \delta v + \delta(1-\delta)(p'-u)^2, \\
m' &= \tanh\left\{\frac{\beta}{2}\left[(z_1 - z_2)(p' + \bar{y} - Rp) - C_1 + C_2\right]\right\}, \\
z_1' &= \frac{(\alpha-R)(p'-\bar{p}) + \phi \log\left[\frac{(p'+\bar{y}-Rp)(z_2-z_1)}{1+\exp[-k(p'+\bar{y}-Rp)(z_2-z_1)]} + 1\right] \tanh[k(p'+\bar{y}-Rp)]}{a_1(1+r^2)\sigma_1^2}, \\
z_2' &= \frac{\gamma(p'-u') - (R-1)(p'-\bar{p}) + \phi \log\left[\frac{(p'+\bar{y}-Rp)(z_1-z_2)}{1+\exp[-k(p'+\bar{y}-Rp)(z_1-z_2)]} + 1\right] \tanh[k(p'+\bar{y}-Rp)]}{a_2(1+r^2+bv')\sigma_1^2}.
\end{cases}
\tag{5.18}
$$

where $z_{1,t}$ and $z_{2,t}$ are given by Equation 5.14, for a fixed value of $t = 1, 2, ...$

We now turn to the existence and uniqueness of the system's steady-state and its local stability analysis, summarised in the following Proposition.

Proposition 5.1. *(i) The 6D map given by Equation 5.18 has a unique steady state (also called fundamental steady state or fundamental equilibrium)* $(p, u, v, m, z_1, z_2) = (\bar{p}, \bar{p}, 0, \bar{m}, 0, 0)$*, where* $\bar{p} = \bar{y}/(R-1)$ *and* $\bar{m} = \tanh(\frac{\beta}{2}(C_2 - C_1))$*. Moreover, at the steady state, the equilibrium fractions are given by* $(\bar{q}_1, \bar{q}_2)$*, where*

$$\begin{cases} \bar{q}_1 = [n_0(1+m_0) + (1-n_0)(1+\bar{m})]/2, \\ \bar{q}_2 = [n_0(1-m_0) + (1+n_0)(1-\bar{m})]/2. \end{cases}$$

(ii) Let

$$\begin{aligned} M &= \frac{\mu(a_1\bar{q}_2(R-1) + a_2\bar{q}_1(R-\alpha))}{a_1 a_2 \sigma^2 (r^2+1)}, \\ \gamma_0 &= (R-1) + \frac{a_2\bar{q}_1}{a_1\bar{q}_2}(R-\alpha) + \frac{a_2\sigma^2(r^2+1)(1-\delta)}{\bar{q}_2\mu\delta}, \end{aligned} \tag{5.19}$$

and assume that $0 < M < 2$*. Then the fundamental steady state* $(\bar{p}, \bar{p}, 0, \bar{m}, 0, 0)$ *is locally asymptotically stable for* $\gamma \in (0, \gamma_0)$*, and it undergoes a Neimark-Sacker bifurcation at* $\gamma = \gamma_0$*. In particular, there is an invariant curve near the fundamental steady state.*

Proof. See Appendix A. □

Note that the stability conditions are independent of k since the realised excess return is zero at the fundamental equilibrium, meaning that regret is also zero at the fundamental steady state. Intuitively, since there is no excess return for agents to make, there is no incentive to invest in the risky asset, demands are zero and hence there is no ex-post disappointment for not having chosen a better strategy. For this reason, when evaluated at the equilibrium, the 6D map from Equation 15 reduces to the 4D map of Dieci et al. (2006) and has the same stability region around the fundamental equilibrium. Mathematically, the Jacobian matrices of the 6D and 4D maps, evaluated at steady state, have the same two non-zero eigenvalues. Similarly, when the impact of regret is zero, $\lambda_{\text{Rgrt}} = 0$, our system reduces to the one of Dieci et al. (2006).

Proposition 5.1 implies that the fundamental steady state is locally stable when the impact of trend followers, measured by γ, is not strong enough ($\gamma < \gamma_0$). Note that γ_0 increases as $\bar{q}_1$ or a_2 increase ($R = 1 + r/K > 1 > \alpha$), or as $\bar{q}_2$, δ, α or μ decrease. In other words, when the fundamental traders dominate the market, either via a high market fraction $\bar{q}_1$, high price adjustment $(1-\alpha)$ or low risk aversion a_1, the fundamental price is stabilised. On the other hand, when the trend followers dominate the market,

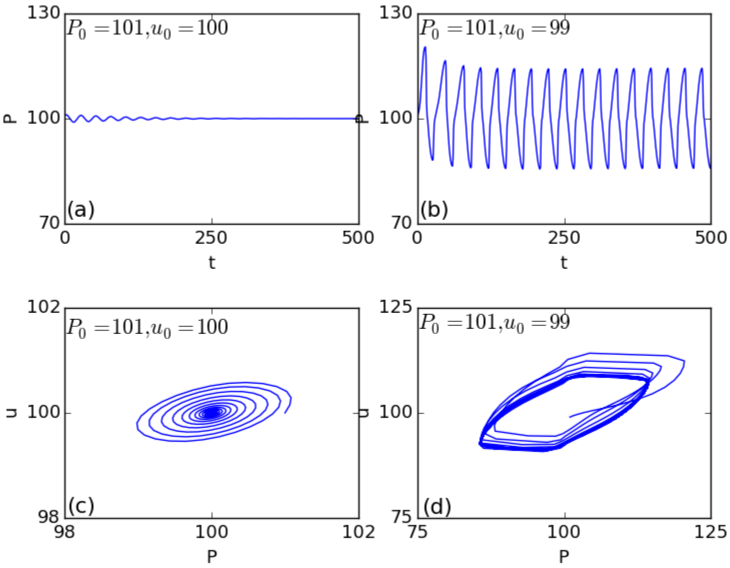

Figure 5.1: Dynamics generated from parameters chosen close to the boundary, in the case of coexisting stable steady state and stable limit cycle. Parameters are: $a_1 = a_2 = a = 0.5, \alpha = 0.3, \gamma = 0.8, \delta = 0.85, \mu = 1, \beta = 0.5, b = 0.03125, n_0 = 0.25, m_0 = 0$. (a,b) Deterministic trajectory of price versus time under slightly different initial conditions($p_0 = \bar{p} + 1 = 101, v_0 = 0, m_0 = \bar{m} \equiv \tanh \frac{\beta}{2}(C_2 - C_1), u_0 = 100$ in (a) and $u_0 = 99$ in (b); (c,d) Phase-plots in the plane of state variables p, u for the two settings of (a) and (b), respectively.

characterised by a high market fraction $\bar{q}_2$, high risk aversion a_2, high weight on the latest price trend δ, strong extrapolation γ or high price adjustment μ, the steady state price is destabilised. Moreover, if the cost of the fundamental strategy C_1 is higher than the cost of trend following C_2, then an increase in the switching intensity β decreases γ_0. That is, the fundamental price becomes less stable when traders switch between strategies more often, which is the well known rational route to randomness of Brock and Hommes (1997, 1998).

Similar results have also been explored in Dieci et al. (2006) and He and Li (2007). Moreover, near the bifurcation point $\gamma = \gamma_0$, the invariant circle around the fundamental steady state was shown to be either stable or unstable. In particular, there exists a critical value $\hat{\gamma}$, such that the stable steady state coexists with the stable forward circle for $\hat{\gamma} < \gamma < \gamma_0$ (see He et al. (2016) for more details). Inside this interval, even when the fundamental steady state is locally stable, prices don't necessarily converge to the fundamental value and can settle down to a stable limit circle. Furthermore, when noise terms are added, the interaction between coexisting stable attractions in the deterministic system and noise processes can endogenously generate important properties observed in financial markets such as volatility clustering and long range dependence in volatility.

Specifically, agent-based modelling argues that financial stylized facts emerge endogenously from the interaction between market participants (Dieci and He, 2018). It has been shown that when γ is close to its bifurcation value, inside a well-defined interval, a coexistence of attractors of the deterministic model leads to important nonlinear macro dynamics (He et al., 2016). In this context, a low volatility regime characterized by a stable steady state coexists with a high volatility one characterized by a stable limit circle. In the next section, we address one of the key aims of our research (see Section 1.2) and focus on the implications of regret inside this particularly interesting volatility clustering interval, where irregular switches between two volatility regimes lead to important dynamics. In particular, we show that regret has the power to destabilise the market, by transforming one volatility regime into the other. That is, a change in market psychology can explain important emerging properties such as volatility clustering.

5.3 The destabilising impact of regret

In this section we provide numerical simulations to examine the model's global dynamics and analyse the effects of regret on the market price behaviour. In all of the numerical experiments, we model a daily trading frequency and therefore we choose $r = 0.05, K = 250$, $R = 1 + r/k = 1.0002, \bar{y} = 0.02, \bar{p} = \bar{y}/(R-1) = 100, \sigma = 0.02, \sigma_1^2 = (\bar{p}\sigma)^2/K = 1.6$ and $C \equiv C_1 - C_2 = 0.5$.

In our first example we verify the coexistence of locally stable attractors. We choose $a_1 = a_2 = a = 0.5, \alpha = 0.3, \gamma = 0.8, \delta = 0.85, \mu = 1, \beta = 0.5, b_2 = 0.05, b = b_2/\sigma_1^2 = 0.03125, n_0 = 0.5, m_0 = 0$. For this choice, there are 50% non-switching agents in the market, where 50% are fundamentalists and 50% are chartists. The remaining 50% of agents change between the two strategies according to the adaptive switching mechanism. In order to demonstrate the coexistence of attractors we consider the deterministic trajectory of price with different initial conditions. Specifically, Figure 5.1 (a) considers the system with initial values $(p_0, u_0, v_0, m_0, z_{1,0}, z_{2,0}) = (\bar{p}+1, \bar{p}, 0, \bar{m}, 0, 0)$, showing that the steady state is locally asymptotically stable. However, if we decrease the initial point of u_0 to $\bar{p} - 1$ and consider the initial conditions $(P_0, u_0, v_0, m_0, z_{1,0}, z_{2,0}) = (\bar{p} + 1, \bar{p} - 1, 0, \bar{m}, 0, 0)$, the price trajectory converges to a closed invariant circle, plotted in Figure 5.1 (b). This phenomenon has also been observed in Dieci et al. (2006).

The coexistence of the locally stable steady state and invariant circle shows that the price dynamics depend on the initial values. Moreover, as with most fundamentalists vs. chartists models, the two types of agents have different impacts on the dynamics. In particular, when the trading activities of either the fundamental investors or the trend followers dominate the market, the price fluctuates around either the fundamental value with low volatility or a cyclical price movement with high volatility (Dieci and He, 2018).

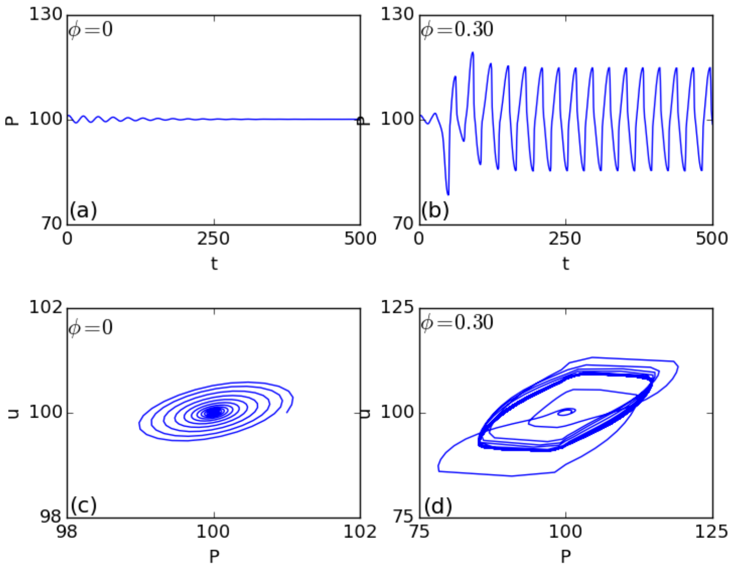

Figure 5.2: Dynamics generated from parameters $a_1 = a_2 = a = 0.5, \alpha = 0.3, \gamma = 0.8, \delta = 0.85, \mu = 1, \beta = 0.5, b = 0.03125, n_0 = 0.25, m_0 = 0$. (a,b) Deterministic trajectory of price versus time in the cases $\lambda_{\text{Rgrt}} = 0$ and $\lambda_{\text{Rgrt}} = 0.3$, respectively, starting with the same initial point close to the steady state ($p_0 = \bar{p} + 1 = 101, u_0 = \bar{p} = 100, v_0 = 0, m_0 = \bar{m} \equiv \tanh \frac{\beta}{2}(C_2 - C_1) \approx -0.12435$); (c,d) Phase-plots in the plane of state variables p, u for $\lambda_{\text{Rgrt}} = 0$ and $\lambda_{\text{Rgrt}} = 0.3$, respectively

.

Next, we focus on the effects of regret on the system's dynamics. In the upcoming example consider a setting with $a_1 = a_2 = a = 0.5, \alpha = 0.3, \gamma = 0.8, \delta = 0.85, \mu = 1, \beta = 0.5, b_2 = 0.05, b = b_2/\sigma_1^2 = 0.03125, n_0 = 0.25, m_0 = 0$. For this setting, there are 25% non-switching agents in the market, with 50% of them fundamentalists. Starting with the same initial values, $(p_0, u_0, v_0, m_0, z_{1,0}, z_{2,0}) = (\bar{p}+1, \bar{p}, 0, \bar{m}, 0, 0)$, the impact of regret parameter λ_{Rgrt} has major implications on the dynamics. On one hand, Figures 5.2 (a) and (c) show that the price converges to its fundamental steady state when agents have no regret, $\lambda_{\text{Rgrt}} = 0$. However, the system becomes unstable and the price converges to a limit cycle as agents increase their impact of regret. Notice the abrupt change in the dynamics following a small change in the impact of regret $\lambda_{\text{Rgrt}} = 0.3$ (Figures 5.2 (b) and (d)). Numerical simulations show that the bifurcation value of the impact of regret, assuming the rest is fixed, is increasing λ_{Rgrt} with $0.10 < \lambda^*_{\text{Rgrt}} < 0.12$ (the bifurcation value is denoted by λ^*_{Rgrt}).

Similarly, if we consider the parameter setting from Figure 5.1 (a) and start increasing the impact of regret, we can demonstrate that for a certain value of λ_{Rgrt} the model changes its attractors. In general, when the impact of regret is beyond a threshold value, the system changes its dynamics and instead of a trajectory which converges to the steady state, it settles to a wide attracting closed curve. However, the reverse is not

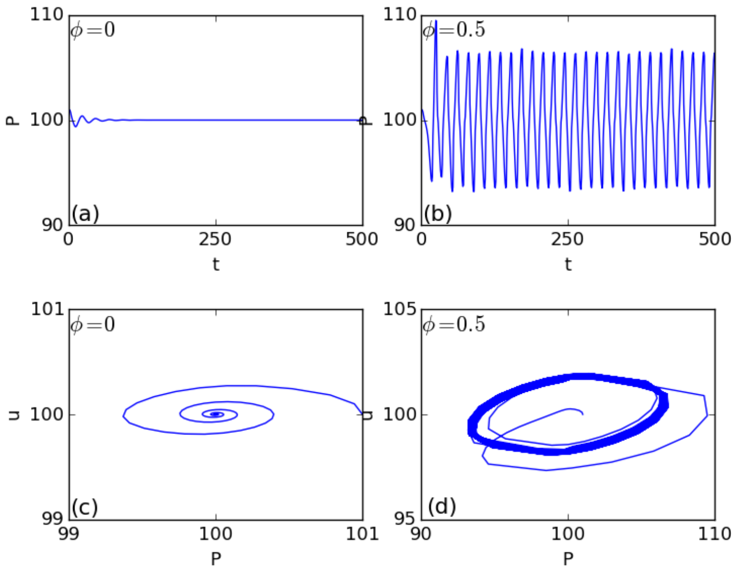

Figure 5.3: Dynamics generated from parameters $a_1 = a_2 = a = 0.5, \alpha = 0.25, \gamma = 0.75, \delta = 0.9, \mu = 1, \beta = 1.5, b = 0.03125, n_0 = 0.8, m_0 = 0.75$. (a,b) Deterministic trajectory of price versus time in the cases $\lambda_{\text{Rgrt}} = 0$ and $\lambda_{\text{Rgrt}} = 0.5$, respectively, starting with the same initial point close to the steady state ($p_0 = \bar{p} + 1 = 101, u_0 = \bar{p} = 100, v_0 = 0, m_0 = \bar{m} \equiv \tanh \frac{\beta}{2}(C_2 - C_1) \approx -0.12435$); (c,d) Phase-plots in the plane of state variables p, u for $\lambda_{\text{Rgrt}} = 0$ and $\lambda_{\text{Rgrt}} = 0.5$, respectively.

true. If we start with a system that is already unstable, the impact of regret has no power in stabilising the price.

However, in the previous examples we have used a set of parameters describing a market with a high proportion of switching agents (75%) and equal proportions of fixed-strategy agents ($m_0 = 0$). Under similar conditions, (Dieci et al., 2006) show that small changes in the parameters n_0 or m_0 can change the state from stable to unstable. In other words, both market mood and proportion of switching agents play an important role in the model. We are now interested to check if a similar phenomenon can be caused by the impact of regret under less intuitive and extreme situations.

The question here is whether a small change in the impact of regret can destabilise an otherwise stable market. That is, even when we are dealing with a relatively quiet setting when the fixed fraction is high (thus the proportion of switching agents is small) and the market mood is mostly characterised by fundamentalist beliefs. To model such a setting, the parameters used in this example are the following: $a_1 = a_2 = a = 0.5, \alpha = 0.25, \gamma = 2.5, \delta = 0.9, \mu = 1, \beta = 1.5, b = 0.03125$; and the fraction of confident agents is high, $n_0 = 0.80$, mostly occupied by fundamentalists, $m_0 = 0.75$. This corresponds to a setting where the fraction n_1 of confident fundamentalists is 70% of the whole population, while the fraction n_2 of confident trend followers is only 10%. We present the price trajectories for $\lambda_{\mathrm{Rgrt}} = 0$ and $\lambda_{\mathrm{Rgrt}} = 0.5$ in Figures 5.3 (a) and (b), respectively, starting with the same initial conditions. While for $\lambda_{\mathrm{Rgrt}} = 0$ the system is locally asymptotically stable and the price converges to its fundamental value (Figures 5.3 (a) and (c)), an increase in the impact of regret to $\lambda_{\mathrm{Rgrt}} = 0.5$ is sufficient to destabilise the market. Figures 5.3 (b) and (d) plot the price's wild fluctuations around the fundamental value, and its projection in the (p,u) plane, respectively. Numerical simulations show that the sudden change in regime occurs when λ_{Rgrt} is increased above a particular threshold value $\bar{\lambda}_{\mathrm{Rgrt}}$ with $0.18 < \bar{\lambda}_{\mathrm{Rgrt}} < 0.2$.

However, this interaction between the coexistence of the deterministic dynamics and noise processes can endogenously generate volatility clustering and long range dependence (Dieci and He, 2018). Economically, market prices fluctuate around either the fundamental value with low volatility or a cyclical price movement with high volatility. We show that a change in the impact of regret, or a change in agents' psychology, can dramatically affect the market dynamics. In particular, the behavioural bias has an indirect effect on the market efficiency and can destabilise it by transforming a low volatility regime into a high volatility one. Finally, since regret alters the mechanism responsible for generating key stylized facts, it has the potential to explain important properties of financial markets.

5.4 Summary

In this chapter we consider a nonlinear discrete-time asset price model with heterogeneous agents that interact in a market with one risky and one risk free asset, under a market maker scenario. The fundamentalists vs. chartists model allows agents to adapt in the market and change their type based on realised profits, but also assumes that a given proportion of agents have fixed strategies over time. The resulting nonlinear dynamical system has been proven capable of generation a complete range of complex behaviour (He and Li, 2017).

Motivated by the desire to have a better representation of investors' behaviour (see Section 1.2), we incorporate regret in agents' behaviour. The main assumption of regret theory is that people may feel disappointed after making a decision under uncertainty, if that decision proves to be wrong or suboptimal ex-post. It is a psychological model in which agents' concerns are not limited to their future expected payoffs and should also depend on the realisation of not having chosen better alternatives. In most cases an investor can observe the performance of others' stocks or portfolios in which he or she could have invested but decided not to. Moreover, an individual who experiences regret is expected to try to anticipate the feeling and take it into consideration when making a decision under uncertainty. For this reason, we modify the agents' expectations of future returns such that they account for regret accordingly. In particular, we model regret as proportional to the level of wealth an agent could have obtained from one time period to another. In the two type fundamentalists vs. chartist setting, agents can only feel disappointed if the other type made a better profit than them in the last trading period.

We focus on the deterministic skeleton of the model and by applying stability and bifurcation analysis, we obtain the conditions of local asymptotic stability of the steady state. Moreover, we verify the endogenous mechanism of volatility clustering by demonstrating the coexistence of a locally stable steady state and a stable closed invariance circle. When doubled with noise, irregular switches between the two attractors, corresponding to two different volatility regimes, can endogenously generate volatility clustering and long range dependence in volatility observed in real financial markets. In this context, as agents become more and more regret averse, we show that their interactions change the local attractors of the deterministic model from a stable steady state to a closed invariant circle.

Specifically, we show that regret has the potential to destabilise the market and transform a relative quiet regime into a highly volatility one (as outlined in our research contributions Section 1.3). Therefore, we demonstrate that the level of regret felt by agents has an indirect impact on the mechanism responsible for generating key properties of financial markets. This helps to provide a deeper understanding of the nonlinear properties of the financial and economic systems and confirms that market psychology plays a crucial role in the overall dynamics.

Chapter 6

Conclusions and Future Work

Heterogeneous agent-based models that rely on simple automated trading agents have proven themselves very efficient in generating important dynamics of real financial markets. Indeed, the fundamentalists vs. chartists models have been shown to successfully capture empirically observed traders' behaviour. In this thesis, we described, evaluated and extended some of the most recent and successful models in capturing the real-life market dynamics, with a focus on key mechanisms and their particular roles in the models' structure.

One of our main objectives in the field of agent-based financial modelling is to propose an alternative to the apparent randomness of financial markets, trying to explain the most important properties of financial data. Specifically, we are interested in simple structures that can reproduce the empirical findings to a high degree and which are quantitatively close to the real ones. To this end, we offer both a quantitative and qualitative analysis of the simulated asset price series, returns, volume traded and volatility. Using econometric tests, we demonstrate the presence of a rich set of stylized facts including absence of autocorrelations, heavy tails, volatility clustering, long memory, volume volatility relations, aggregate Gaussianity, price impact and extreme price events. By doing so, we show that a rich set of the stylized facts can arise endogenously, from the interaction of market participants.

However, a well-defined agent-based model able to match the financial markets' properties is ideal for testing various behavioural and economical theories. We believe that the fields of agent-based modelling and behavioural finance complement each other and in this thesis we demonstrated how they can be used together to provide a deeper understanding on the highly nonlinear economic systems. Because of the complexities of real financial markets, it is difficult to assess the implications of investors' behavioural heuristics on the overall market dynamics. Nevertheless, an agent based model that can recreate the stylized facts of real financial markets and which is simple enough to ensure tractability can be used as a tool to test the effects of micro level interactions at the macro level.

Thus, a further objective of our work is to test the influence of various heuristics in agents' behaviour on their interactions and price movements. In this respect, we present two new behavioural models of asset pricing, one incorporating loss aversion and another the disposition effect. In the new settings, we observed major differences in the movements of the agents' market fractions. Specifically, as investors become more and more loss averse, we show that this behavioural bias can lead to their complete disappearance. Moreover, measuring how close a model is to its empirical counterparts, we demonstrate that the introduction of loss aversion clearly improves the degree to which the model explains the real data, compared to the same model without loss aversion. Similarly, focusing on the implications of the disposition effect, we demonstrate that the introduction of this behavioural bias leads to violations of some of the most discussed stylized facts of financial data.

An additional objective of this thesis is to offer a better understanding of the mechanisms responsible for generating important stylized facts and their relation to investors' behaviour. In this context, we focus on the effects of regret on the well-known adaptive beliefs system and its coexistence of attractors, which was shown to trigger volatility clustering of returns. We demonstrate that regret has a direct impact on the mechanism generating important properties observed in real financial markets. In other words, we show that investors' psychology plays an important role in the overall market dynamics. Finally, the properties observed in real financial markets can be generated endogenously by agents' interactions, with investors' behaviour and deviations from rationality playing an important role in the key mechanisms responsible for their existence.

6.1 Future Work

At this stage, we have discussed models where the market participants, i.e. fundamentalists and chartists, are modelled in the form of a representative agent on top of which heterogeneity is added through noise variables. For a better representation and understanding of agents' interactions we can take a step further and model individual agents (see e.g. Schmitt (2018)). This way, we can test the implications of various behavioural biases or economic policies through a higher level of heterogeneity, with their heuristics being modelled individually. Taking a step further, we can test the implications of investors' behaviour in a more granular way and assess the consequences of individual choices on the other market participants and its internal dynamics.

Moreover, in this thesis we model agents' interactions and switching between strategies according to a series of factors including wealth, herding, price misalignment or predisposition towards a specific one of them. However, we are interested in exploring other factors that force the traders to change their beliefs regarding future price movements and test their consequences. This is one of the main objectives of well-defined ABMs. As we have seen in Chapter 4, introducing a loss averse preference had major implications on the interaction between agents. Therefore, more behavioural components can be incorporated in the agents' interactions. We want to fully exploit our ABM and investigate how different behavioural factors and biases such as overconfidence, underconfidence, representativeness or anchoring, to name a few, can be integrated into the models' mechanism and assess their implications. Because of the tractability offered by agent-based modelling, this can provide us with consequences that cannot be observed in the complex world of real financial markets, where is difficult to measure the implications of individual behaviour.

One of the major challenges in the field of agent-based financial modelling is the way the model parameters are evaluated (Lussange et al., 2018). The method of simulated moments and the maximum likelihood dominate the literature and appear to be the best available option. They require the minimisation or maximisation of an objective function capturing the difference between a set of empirical and simulated properties or moments (see Section 3.2). However, these methods often fail because the objective function does not always behave well enough to guarantee a global optimum solution (Fabretti, 2013). In particular, because of the random variables in the model definition, the distance function is highly stochastic and nonlinear. Hence, it is not smooth and has a multitude of local minima that are hard to overcome. Moreover, the algorithm usually used for the minimisation problem, i.e. Nelder-Mead algorithm, is not always capable of solving this issues. Since the algorithms cannot get out of the multiple valleys, their solution is almost always close to the initial guess. Efficiently calibrating agent-based models to empirical data is an open challenge, with various methods being proposed over the years (see e.g. Fagiolo et al. (2017) and Lamperti et al. (2018)).

The calibration procedure, in the sense of parameter estimation, represents a crucial step in order to fit a model to specific empirical data. But given the uncertainty about the most adequate model, choosing the exact values is not sufficient. One must take into account a further step, namely a model selection procedure. To this end, some agent-based models have been rejected by later work (Winker and Jeleskovic, 2007), authors concluding that the replication of few stylized facts is not completely satisfactory and is attributed to the design of the model itself. Although model design is still quite an open problem in ABM, a commonly agreed upon set of core concepts is still missing. Moreover, the issue of discriminating between various models and their comparison has not been completely solved and requires further research (Barde, 2015).

In the previous chapters, we have seen that our models are able to generate the stylized facts of financial markets and it is also possible to offer explanations of their origins. However, in our models, as in most of the models in the literature, these properties occur only in a very specific and limited region of the model parameters. This is rarely discussed in the literature (Alfi et al., 2009) but the fundamental question is why the market dynamics evolves spontaneously, or self-organises, in the specific region of parameters which corresponds to the stylized facts? Usually, self-organisation is observed in physics, where nonlinear dynamics spontaneously drive the system towards the critical state. This can occur from a variety of initial configurations and parameters and define a state that is always the same, the critical one. For financial models, an analogy would be very tempting and we should investigate its possibility and implications. Moreover, in this thesis we focused on one particular mechanism responsible for generating important stylized facts, namely the coexistence of attractors. For future work, we can investigate how various behavioural biases affect other mechanisms such as the herding mechanism of Lux or the structural stochastic volatility of Franke and Westerhoff.

Finally, the agents in our model are presented with the option of trading a single asset at each time period. A further modification can be done by increasing the number of assets available for trading (see Dieci et al. (2018)). Obviously, in real life the investors can change from trading one asset to another depending on their current prices or other characteristics in the market. One way of preserving heterogeneity would be by allowing the agents to invest in various assets on different markets. Furthermore, instead of modelling the participants trading at each time period, we could allow for different actions, depending on the observed state of the market and agents' character. In this way, our model will capture the empirically observed traders' behaviour to a higher level and better represent the real world interactions. A similar direction of future work is to model the relationship between different financial markets. We can further expand our model so that the agents could trade in different stock markets. Therefore, we will focus not only on the stylized facts of individual markets, but also on the relationship between them. For example, agents' behaviour may cause correlated movements of asset prices, returns and volumes traded on different markets as we observe in real life. However, increasing heterogeneity can lead to a decrease in the models' tractability, therefore losing the primary advantage of agent-based modelling compared to the real financial markets.

Appendix A

Appendix A

In this section we provide the proof of Proposition 5.1 regarding the existence and stability of the fundamental steady state of the transformation map given by Equation 5.18.

Proof. (i) One can easily check that $(\bar{p}, \bar{p}, 0, \bar{m}, 0, 0)$ is a fixed point of the map (5.18) and hence a steady state point. We show that there is only one such steady state. Suppose that there is another steady point that is the fixed point of (5.18), and we use the symbol “ $\hat{}$ ” to indicate the equilibrium quantities that define this point. This immediately implies that $\hat{p} = \hat{u}$, $\hat{v} = 0$, the total excess demand is zero, i.e. $\hat{q}_1\hat{z}_1 + \hat{q}_2\hat{z}_2 = 0$, and the equilibrium demands are given by

$$\hat{z}_1 = \frac{(\alpha - R)(\hat{p} - \bar{p})}{a_1(1 + r^2)\sigma_1^2},$$
$$\hat{z}_2 = \frac{(1 - R)(\hat{p} - \bar{p})}{a_2(1 + r^2)\sigma_1^2}.$$

Since $R = 1 + r/K > 1 > \alpha$, both $(\alpha - R)$ and $(1 - R)$ are negative, and so $\hat{z}_1, \hat{z}_2$ must have the same sign. Given that $\hat{q}_1, \hat{q}_2 \geq 0$, both $\hat{q}_1\hat{z}_1$ and $\hat{q}_2\hat{z}_2$ admit the same sign as well, and since $\hat{q}_1\hat{z}_1 + \hat{q}_2\hat{z}_2 = 0$, their addition is zero if and only if both of them are zero, i.e. $\hat{q}_1\hat{z}_1 = \hat{q}_2\hat{z}_2 = 0$. Since $\hat{q}_1, \hat{q}_2 \geq 0$ and $\hat{q}_1 + \hat{q}_2 = 1$, it is implied that $\hat{z}_1 = 0$ or $\hat{z}_2 = 0$ which in either case leads us to $\hat{p} = \bar{p}$. This proves the uniqueness.

(ii) In order to find the stability conditions we split the proof into 4 steps.

Step 1. The conditions of local asymptotic stability of the steady state are determined by the eigenvalues of the Jacobian matrix of the map T^k in (5.18), evaluated at the steady state itself. Note that at the fundamental steady state, $p' = p = u = \bar{p}$, $v' = v = 0$, $z_1 = z_2 = 0$ and $p' + \bar{D} - Rp = 0$. Therefore, the transition map T^k is independent of k at the fundamental steady state, where regret does not exist. Next, we evaluate the Jacobian matrix.

- First, we have,

$$
\begin{aligned}
&\frac{\partial z_1}{\partial p} = \frac{\alpha - R}{a_1\sigma^2(r^2+1)}, \quad \frac{\partial z_2}{\partial p} = \frac{\gamma - R + 1}{a_2\sigma^2(r^2+1)},\\
&\frac{\partial z_2}{\partial u} = \frac{-\gamma}{a_2\sigma^2(r^2+1)}, \quad \frac{\partial z_2}{\partial v} = 0.
\end{aligned}
$$

- Consider the partial derivatives of the agents' demands with respect to the state variables, evaluated at the fundamental steady state,

$$
\begin{aligned}
&\frac{\partial z_1'}{\partial p} = \frac{\alpha - R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial p}, \quad \frac{\partial z_1'}{\partial u} = \frac{\alpha - R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial u},\\
&\frac{\partial z_1'}{\partial v} = 0, \quad \frac{\partial z_1'}{\partial m} = 0,\\
&\frac{\partial z_1'}{\partial z_1} = \frac{(\alpha - R)}{a_1(r^2+1)\sigma^2}\frac{\partial p'}{\partial z_1}, \quad \frac{\partial z_1'}{\partial z_2} = \frac{(\alpha - R)}{a_2(r^2+1)\sigma^2}\frac{\partial p'}{\partial z_2},
\end{aligned}
$$

$$
\begin{aligned}
&\frac{\partial z_2'}{\partial p} = \frac{\gamma\delta - R + 1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial p},\\
&\frac{\partial z_2'}{\partial u} = \frac{-\gamma\delta + (\gamma\delta + R - 1)\frac{\partial p'}{\partial u}}{a_2(r^2+1)\sigma^2},\\
&\frac{\partial z_2'}{\partial v} = 0, \quad \frac{\partial z_2'}{\partial m} = 0,\\
&\frac{\partial z_2'}{\partial z_1} = \frac{\gamma\delta - R + 1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_1},\\
&\frac{\partial z_2'}{\partial z_2} = \frac{\gamma\delta - R + 1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_2}.
\end{aligned}
$$

- Consider the partial derivatives of p' with respect to the state variables, evaluated at the fundamental steady state,

$$
\begin{aligned}
&\frac{\partial p'}{\partial p} = 1 + \mu(\bar{q_1}\frac{\partial z_1}{\partial p} + \bar{q_2}\frac{\partial z_2}{\partial p}),\\
&\frac{\partial p'}{\partial u} = \mu\bar{q_1}\frac{\partial z_1}{\partial p},\\
&\frac{\partial p'}{\partial v} = 0, \quad \frac{\partial p'}{\partial m} = 0,\\
&\frac{\partial p'}{\partial z_1} = \mu\bar{q_1}, \quad \frac{\partial p'}{\partial z_2} = \mu\bar{q_2}.
\end{aligned}
$$

- Consider the partial derivatives of u' with respect to the state variables, evaluated at the fundamental steady state,

$$\frac{\partial u'}{\partial p} = (1-\delta)\frac{\partial p'}{\partial p},$$
$$\frac{\partial u'}{\partial u} = \delta + (1-\delta)\frac{\partial p'}{\partial u},$$
$$\frac{\partial u'}{\partial v} = 0, \quad \frac{\partial u'}{\partial m} = 0,$$
$$\frac{\partial u'}{\partial z_1} = (1-\delta)\frac{\partial p'}{\partial z_1}, \quad \frac{\partial u'}{\partial z_2} = (1-\delta)\frac{\partial p'}{\partial z_2}.$$

- Consider the partial derivatives of v' with respect to the state variables, evaluated at the fundamental steady state,

$$\frac{\partial v'}{\partial p} = 2\delta(1-\delta)(p'-u)\frac{\partial p'}{\partial p} = 0,$$
$$\frac{\partial v'}{\partial u} = 2\delta(1-\delta)(p'-u)(\frac{\partial p'}{\partial u} - 1) = 0,$$
$$\frac{\partial v'}{\partial v} = \delta + 2\delta(1-\delta)(p'-u)\frac{\partial p'}{\partial v} = \delta,$$
$$\frac{\partial v'}{\partial m} = 2\delta(1-\delta)(p'-u)\frac{\partial p'}{\partial m} = 0,$$
$$\frac{\partial v'}{\partial z_1} = 2\delta(1-\delta)(p'-u)\frac{\partial p'}{\partial z_1} = 0,$$
$$\frac{\partial v'}{\partial z_2} = 2\delta(1-\delta)(p'-u)\frac{\partial p'}{\partial z_2} = 0.$$

- Consider the partial derivatives of m' with respect to the state variable, evaluated at the fundamental steady state,

$$\frac{\partial m'}{\partial p} = 0, \quad \frac{\partial m'}{\partial u} = 0, \quad \frac{\partial m'}{\partial v} = 0,$$
$$\frac{\partial m'}{\partial m} = 0, \quad \frac{\partial m'}{\partial z_1} = 0, \quad \frac{\partial m'}{\partial z_2} = 0.$$

Therefore, the Jacobian matrix of the map T^k is given by,

$$J = \begin{bmatrix} \frac{\partial p'}{\partial p} & \frac{\partial p'}{\partial u} & 0 & 0 & \frac{\partial p'}{\partial z_1} & \frac{\partial p'}{\partial z_2} \\ (1-\delta)\frac{\partial p'}{\partial p} & \delta + (1-\delta)\frac{\partial p'}{\partial u} & 0 & 0 & (1-\delta)\frac{\partial p'}{\partial z_1} & (1-\delta)\frac{\partial p'}{\partial z_2} \\ 0 & 0 & \delta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{\alpha-R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial p} & \frac{\alpha-R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial u} & 0 & 0 & \frac{\alpha-R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_1} & \frac{\alpha-R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_2} \\ \frac{\gamma\delta-R+1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial p} & \frac{-\gamma\delta+(\gamma\delta-R+1)\frac{\partial p}{\partial u}}{a_2(1+r^2)\sigma^2} & 0 & 0 & \frac{\gamma\delta-R+1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_1} & \frac{\gamma\delta-R+1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_2} \end{bmatrix},$$

and the structure of the Jacobian implies that two eigenvalues are $\lambda_1 = 0$ and $\lambda_2 = \delta \in (0,1)$. The remaining eigenvalues are given by the matrix,

$$J^{'} = \begin{bmatrix} \frac{\partial p'}{\partial p} & \frac{\partial p'}{\partial u} & \frac{\partial p'}{\partial z_1} & \frac{\partial p'}{\partial z_2} \\ (1-\delta)\frac{\partial p'}{\partial p} & \delta + (1-\delta)\frac{\partial p'}{\partial u} & (1-\delta)\frac{\partial p'}{\partial z_1} & (1-\delta)\frac{\partial p'}{\partial z_2} \\ \frac{\alpha - R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial p} & \frac{\alpha - R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial u} & \frac{\alpha - R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_1} & \frac{\alpha - R}{a_1\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_2} \\ \frac{\gamma\delta - R+1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial p} & \frac{-\gamma\delta + (\gamma\delta - R+1)\frac{\partial p'}{\partial u}}{a_2(r^2+1)\sigma^2} & \frac{\gamma\delta - R+1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_1} & \frac{\gamma\delta - R+1}{a_2\sigma^2(r^2+1)}\frac{\partial p'}{\partial z_2} \end{bmatrix}.$$

The characteristic polynomial of $J^{'}$ gives two more zero eigenvalues $\lambda_3 = \lambda_4 = 0$, with the last two eigenvalues satisfying the equation

$$\begin{aligned} F(\lambda) &= \lambda^2 + A(\gamma)\lambda + B(\gamma), \\ A(\gamma) &= -\left\{\delta + 1 + \frac{\mu}{a_1 a_2\sigma^2(r^2+1)}\left[a_1\bar{q}_2(\gamma\delta - R + 1) + a_2\bar{q}_1(\alpha - R)\right]\right\}, \\ B(\gamma) &= \delta\left\{1 + \frac{\mu}{a_1 a_2\sigma^2(r^2+1)}\left[a_1\bar{q}_2(\gamma - R + 1) + a_2\bar{q}_1(\alpha - R)\right]\right\}. \end{aligned} \tag{A.1}$$

Step 2. Let

$$\begin{aligned} M &= \frac{\mu(a_1\bar{q}_2(R-1) + a_2\bar{q}_1(R-\alpha))}{a_1 a_2\sigma^2(r^2+1)}, \\ \gamma_0 &= (R-1) + \frac{a_2\bar{q}_1}{a_1\bar{q}_2}(R-\alpha) + \frac{a_2\sigma^2(r^2+1)(1-\delta)}{\bar{q}_2\mu\delta}, \end{aligned}$$

and suppose that $0 < M < 2$. In this step, we show that all the roots of the $F(x)$ lie inside the unit circle for $\gamma \in [0, \gamma_0)$, and if $\gamma = \gamma_0$, there is a pair of conjugate roots, $\cos\theta_0 \pm i\sin\theta_0$ for some $\theta_0 \in (0, \pi/2)$, of $F(x)$.

Suppose that $0 \le \gamma < \gamma_0$, the roots of $F(x)$ lies on the unit circle if and only if $|B(\gamma)| \le 1$ and $|A(\gamma)| \le B(\gamma) + 1$[1], and using (A.1), both inequalities could be easily checked. This proves that all the roots lie on the unit circle.

Next, suppose that $x = \cos\theta_0 \pm i\sin\theta_0$, $\theta_0 \in (0, \pi)$ are roots of the characteristic polynomial. Then we have

$$\cos^2\theta_0 - \sin^2\theta_0 + A(\gamma)\cos\theta_0 + B(\gamma) = 0,$$

$$2\sin\theta_0\cos\theta_0 + A(\gamma)\sin\theta_0 = 0,$$

[1] It can be shown that the roots of any quadratic equation such as $z^2 + az + b = 0$ (where a and b are real coefficients) lie on the unit circle if and only if $|b| \le 1$ and $|a| \le b + 1$.

which are equivalent to $2\cos^2\theta_0 + A(\gamma)\cos\theta_0 + B(\gamma) - 1 = 0$ and $\sin\theta_0(2\cos\theta_0 + A(\gamma)) = 0$. Moreover, $2\cos\theta_0 + A(\gamma) = 0$ since $\sin\theta_0 \neq 0$. Therefore, we have that $B(\gamma) = 1$, or equivalently, $\gamma = \gamma_0$.

Substituting γ_0 back into the second equality, we have $\cos\theta_0 = -\frac{A(\gamma_0)}{2}$. However, $A(\gamma_0) = -(2 + M(\delta - 1))$, which implies that $\cos\theta_0 \in (\delta, 1)$ and therefore $\theta_0 \in (0, \pi/2)$.

Step 3. Suppose that $\rho(\gamma)(\cos\theta(\gamma) \pm i\sin\theta(\gamma))$ are the roots of $F(x)$ satisfying $\rho(\gamma_0) = 1$ and $\theta(\gamma_0) = \theta_0$. In this step, we show that $\rho'(\gamma_0) > 0$. By differentiating

$$\lambda(\gamma) = \rho(\gamma)(\cos\theta(\gamma) + i\sin\theta(\gamma)),$$

with respect to γ, we obtain

$$\begin{aligned}\lambda'(\gamma) &= \rho'(\gamma)\cos\theta(\gamma) - \rho(\gamma)\theta'(\gamma)\sin\theta(\gamma) + i\rho'(\gamma)\sin\theta(\gamma) + i\rho(\gamma)\theta'(\gamma)\cos\theta(\gamma) \\ &= U_1(\gamma) + iV_1(\gamma),\end{aligned} \tag{A.2}$$

where

$$U_1(\gamma) = \rho'(\gamma)\cos\theta(\gamma) - \rho(\gamma)\theta'(\gamma)\sin\theta(\gamma),$$

$$V_1(\gamma) = \rho'(\gamma)\sin\theta(\gamma) + \rho(\gamma)\theta'(\gamma)\cos\theta(\gamma),$$

which implies that

$$\rho'(\gamma) = U_1(\gamma)\cos\theta(\gamma) + V_1(\gamma)\sin\theta(\gamma). \tag{A.3}$$

Similarly, by differentiating

$$0 = \lambda^2(\gamma) + A(\gamma)\lambda(\gamma) + B(\gamma),$$

with respect to γ, we have

$$\begin{aligned}\lambda'(\gamma) &= -\frac{B'(\gamma) + A'(\gamma)\lambda(\gamma)}{2\lambda(\gamma) + A(\gamma)} \\ &= -\frac{B'(\gamma) + A'(\gamma)\rho(\gamma)(\cos\theta(\gamma) + i\sin\theta(\gamma))}{A(\gamma) + 2\rho(\gamma)(\cos\theta(\gamma) + i\sin\theta(\gamma))} \\ &= U_2(\gamma) + iV_2(\gamma),\end{aligned} \tag{A.4}$$

where

$$U_2(\gamma) = -\frac{(B'(\gamma) + A'(\gamma)\rho(\gamma)\cos\theta(\gamma))(A(\gamma) + 2\rho(\gamma)\cos\theta(\gamma)) + 2A'(\gamma)\rho^2(\gamma)\sin^2\theta(\gamma)}{(A(\gamma) + 2\rho(\gamma)\cos\theta(\gamma))^2 + (2\rho(\gamma)\sin\theta(\gamma))^2},$$

$$V_2(\gamma) = \frac{2\rho(\gamma)\sin\theta(\gamma)(B'(\gamma) + A'(\gamma)\rho(\gamma)\cos\theta(\gamma)) - A'(\gamma)\rho(\gamma)\sin\theta(\gamma)(A(\gamma) + 2\rho(\gamma)\cos\theta(\gamma))}{(A(\gamma) + 2\rho(\gamma)\cos\theta(\gamma))^2 + (2\rho(\gamma)\sin\theta(\gamma))^2}.$$

From (A.2) and (A.4), we have that $U_1(\gamma) = U_2(\gamma)$, $V_1(\gamma) = V_2(\gamma)$. Therefore, by (A.3), we obtain $\rho'(\gamma) = U_2(\gamma)\cos\theta(\gamma) + V_2(\gamma)\sin\theta(\gamma)$, and using the fact that $\rho(\gamma_0) = 1$ and $A(\gamma_0) + 2\rho(\gamma_0)\cos\theta(\gamma_0) = 0$, we have

$$\begin{aligned}
\operatorname{sgn}\{\rho'(\gamma_0)\} &= \operatorname{sgn}\{\frac{-2A'(\gamma_0)\rho^2(\gamma_0)\sin^2\theta(\gamma_0)\cos\theta(\gamma_0)}{(2\rho(\gamma)\sin\theta(\gamma))^2} \\
&+ \frac{2\rho(\gamma_0)\sin^2\theta(\gamma_0)(B'(\gamma_0) + A'(\gamma_0)\rho(\gamma_0)\cos\theta(\gamma_0))}{(2\rho(\gamma)\sin\theta(\gamma))^2}\} \\
&= \operatorname{sgn}\{B'(\gamma_0)\} \\
&= \operatorname{sgn}\{\frac{\delta\mu\bar{q}_2}{a_2\sigma^2(1+r^2)}\} > 0.
\end{aligned}$$

Step 4. Finally, it follows from (Kuznetsov, 2004) that the steady state $(\bar{p}, \bar{p}, 0, \bar{m}, 0, 0)$ of the map T^k is asymptotically stable if $\gamma \in (0, \gamma_0)$, and the map T^k undergoes a Neimark-Sacker bifurcation at $\gamma = \gamma_0$ which implies that there is an isolated closed invariant curve near the origin.

$\square$

References

Alfarano, S., Lux, T., and Wagner, F. (2005). Estimation of agent-based models: the case of an asymmetric herding model. *Computational Economics*, 26(1):19–49.

Alfarano, S., Lux, T., and Wagner, F. (2006). Estimation of a simple agent-based model of financial markets: An application to australian stock and foreign exchange data. *Physica A: Statistical Mechanics and its Applications*, 370(1):38–42.

Alfarano, S., Lux, T., and Wagner, F. (2007). Empirical validation of stochastic models of interacting agents. *The European Physical Journal B*, 55(2):183–187.

Alfi, V., Cristelli, M., Pietronero, L., and Zaccaria, A. (2009). Minimal agent based model for financial markets i. *The European Physical Journal B*, 67(3):385–397.

Allen, H. and Taylor, M. P. (1990). Charts, noise and fundamentals in the london foreign exchange market. *The Economic Journal*, pages 49–59.

Amilon, H. (2008). Estimation of an adaptive stock market model with heterogeneous agents. *Journal of Empirical Finance*, 15(2):342–362.

Andrews, D. W. (2004). The block–block bootstrap: improved asymptotic refinements. *Econometrica*, 72(3):673–700.

Anufriev, M. and Hommes, C. (2012). Evolution of market heuristics. *The Knowledge Engineering Review*, 27(2):255–271.

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American economic review*, pages 406–411.

Barber, B. M., Lee, Y.-T., Liu, Y.-J., and Odean, T. (2007). Is the aggregate investor reluctant to realise losses? evidence from taiwan. *European Financial Management*, 13(3):423–447.

Barberis, N., Huang, M., and Santos, T. (1999). Prospect theory and asset prices. Technical report, National bureau of economic research.

Barberis, N., Huang, M., and Thaler, R. H. (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American economic review*, 96(4):1069–1090.

Barberis, N. and Thaler, R. (2003). A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128.

Barberis, N. and Xiong, W. (2009). What drives the disposition effect? an analysis of a long-standing preference-based explanation. *the Journal of Finance*, 64(2):751–784.

Barde, S. (2015). Direct calibration and comparison of agent-based herding models of financial markets. Technical report, School of Economics Discussion Papers.

Beja, A. and Goldman, M. B. (1980). On the dynamic behavior of prices in disequilibrium. *The Journal of Finance*, 35(2):235–248.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations research*, 30(5):961–981.

Benartzi, S. and Thaler, R. H. (1993). Myopic loss aversion and the equity premium puzzle. Technical report, National Bureau of Economic Research.

Bikhchandani, S. and Sharma, S. (2000). Herd behavior in financial markets. *IMF Staff papers*, 47(3):279–310.

Bleichrodt, H. and Wakker, P. P. (2015). Regret theory: A bold alternative to the alternatives. *The Economic Journal*, 125(583):493–532.

Bondarenko, O. (2003). Statistical arbitrage and securities prices. *Review of Financial Studies*, 16(3):875–919.

Boswijk, H. P., Hommes, C. H., and Manzan, S. (2007). Behavioral heterogeneity in stock prices. *Journal of Economic dynamics and control*, 31(6):1938–1970.

Bouchaud, J.-P., Mézard, M., and Potters, M. (2008). Statistical properties of stock order books: empirical results and models. *arXiv preprint cond-mat/0203511*.

Bouchaud, J.-P. and Potters, M. (2003). *Theory of financial risk and derivative pricing: from statistical physics to risk management*. Cambridge university press.

Bracke, P. and Tenreyro, S. (2016). History dependence in the housing market.

Braun, M. and Muermann, A. (2004). The impact of regret on the demand for insurance. *Journal of Risk and Insurance*, 71(4):737–767.

Brennan, T. J. and Lo, A. W. (2011). The origin of behavior. *The Quarterly Journal of Finance*, 1(01):55–108.

Brock, W., Lakonishok, J., and LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of finance*, pages 1731–1764.

Brock, W. A. and Hommes, C. H. (1997). A rational route to randomness. *Econometrica: Journal of the Econometric Society*, pages 1059–1095.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control*, 22(8):1235–1274.

Brock, W. A., Hommes, C. H., and Wagener, F. O. (2005). Evolutionary dynamics in markets with many trader types. *Journal of Mathematical Economics*, 41(1):7–42.

Brooks, R. A. (1991). Intelligence without representation. *Artificial intelligence*, 47(1):139–159.

Chakraborti, A., Toke, I. M., Patriarca, M., and Abergel, F. (2011). Econophysics review: I. empirical facts. *Quantitative Finance*, 11(7):991–1012.

Chandra, A. and Thenmozhi, M. (2017). Behavioural asset pricing: Review and synthesis. *Journal of Interdisciplinary Economics*, 29(1):1–31.

Chen, G., Kim, K. A., Nofsinger, J. R., and Rui, O. M. (2007). Trading performance, disposition effect, overconfidence, representativeness bias, and experience of emerging market investors. *Journal of Behavioral Decision Making*, 20(4):425–451.

Chen, S.-H., Chang, C.-L., and Du, Y.-R. (2012). Agent-based economic models and econometrics. *The Knowledge Engineering Review*, 27(02):187–219.

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

Chiarella, C., Dieci, R., and He, X. (2008). Heterogeneity, market mechanisms, and asset price dynamics. *Quantitative Finance Research Centre Research Paper*, (231).

Chiarella, C. and He, X.-Z. (2002). Heterogeneous beliefs, risk and learning in a simple asset pricing model. *Computational Economics*, 19(1):95–132.

Chiarella, C. and He, X.-Z. (2003a). Dynamics of beliefs and learning under al-processesthe heterogeneous case. *Journal of Economic Dynamics and Control*, 27(3):503–531.

Chiarella, C., He, X.-Z., Hung, H., and Zhu, P. (2006). An analysis of the cobweb model with boundedly rational heterogeneous producers. *Journal of Economic Behavior & Organization*, 61(4):750–768.

Chuang, C.-C., Kuan, C.-M., and Lin, H.-y. (2009). Causality in quantiles and dynamic stock return–volume relations. *Journal of Banking & Finance*, 33(7):1351–1360.

Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica: journal of the Econometric Society*, pages 135–155.

Connolly, T. and Zeelenberg, M. (2002). Regret in decision making. *Current directions in psychological science*, 11(6):212–216.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues.

Cont, R. (2005). Long range dependence in financial markets. In *Fractals in Engineering*, pages 159–179. Springer.

Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88.

Cont, R., Potters, M., and Bouchaud, J.-P. (1997). Scaling in stock market data: stable laws and beyond. In *Scale invariance and beyond*, pages 75–85. Springer.

Copeland, T. (1976). A model of asset trading under the assumption of sequential information arrival*. *The Journal of Finance*, 31(4):1149–1168.

Day, R. H. and Huang, W. (1990). Bulls, bears and market sheep. *Journal of Economic Behavior & Organization*, 14(3):299–329.

De Bondt, W. (2005). *The psychology of world equity markets*. Edward Elgar.

De Bondt, W., Mayoral, R. M., and Vallelado, E. (2013). Behavioral decision-making in finance: An overview and assessment of selected research. *Spanish Journal of Finance and Accounting/Revista Española de Financiación y Contabilidad*, 42(157):99–118.

De Bondt, W. F., Muradoglu, Y. G., Shefrin, H., and Staikouras, S. K. (2008). Behavioral finance: Quo vadis? *Journal of Applied Finance (Formerly Financial Practice and Education)*, 18(2).

De Bondt, W. F. and Thaler, R. H. (1994). Financial decision-making in markets and firms: A behavioral perspective. Technical report, National Bureau of Economic Research.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of political Economy*, pages 703–738.

De Vries, C. G. and Leuven, K. (1994). Stylized facts of nominal exchange rate returns.

Dekkers, A. L. and De Haan, L. (1989). On the estimation of the extreme-value index and large quantile estimation. *The Annals of Statistics*, pages 1795–1832.

Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association*, 74(366a):427–431.

Dieci, R., Foroni, I., Gardini, L., and He, X.-Z. (2006). Market mood, adaptive beliefs and asset price dynamics. *Chaos, Solitons & Fractals*, 29(3):520–534.

Dieci, R. and He, X.-Z. (2018). Heterogeneous models in finance agent. *Heterogeneous Agent Modeling*, 4:257.

Dieci, R., Schmitt, N., and Westerhoff, F. (2018). Steady states, stability and bifurcations in multi-asset market models. *Decisions in Economics and Finance*, pages 1–22.

Diks, C. and Van Der Weide, R. (2005). Herding, a-synchronous updating and heterogeneity in memory in a cbs. *Journal of Economic dynamics and control*, 29(4):741–763.

Ding, Z., Granger, C. W., and Engle, R. F. (1993). A long memory property of stock market returns and a new model. *Journal of empirical finance*, 1(1):83–106.

Dodonova, A. and Khoroshilov, Y. (2005). Applications of regret theory to asset pricing.

Engle, R. (2001). Garch 101: The use of arch/garch models in applied econometrics. *Journal of economic perspectives*, 15(4):157–168.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrics*, 50:987–1007.

Evstigneev, I. V., Hens, T., and Schenk-Hoppé, K. R. (2008). Evolutionary finance. *Swiss Finance Institute Research Paper*, (08-14).

Fabretti, A. (2013). On the problem of calibrating an agent based model for financial markets. *Journal of Economic Interaction and Coordination*, 8(2):277–293.

Fagiolo, G., Guerini, M., Lamperti, F., Moneta, A., and Roventini, A. (2017). Validation of agent-based models in economics and finance. Technical report, LEM Working Paper Series.

Fama, E. F. (1965). The behavior of stock-market prices. *Journal of business*, pages 34–105.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work*. *The journal of Finance*, 25(2):383–417.

Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256):685.

Farmer, J. D. and Joshi, S. (2002). The price dynamics of common trading strategies. *Journal of Economic Behavior & Organization*, 49(2):149–171.

Fielitz, B. D. (1971). Stationarity of random data: some implications for the distribution of stock price changes. *Journal of Financial and Quantitative Analysis*, 6(03):1025–1034.

Fisher, R. A. and Tippett, L. H. C. (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 24, pages 180–190. Cambridge Univ Press.

Fogel, S. O. and Berry, T. (2006). The disposition effect and individual investor decisions: the roles of regret and counterfactual alternatives. *The journal of behavioral finance*, 7(2):107–116.

Franke, R. (2009). Applying the method of simulated moments to estimate a small agent-based asset pricing model. *Journal of Empirical Finance*, 16(5):804–815.

Franke, R. and Westerhoff, F. (2011). Estimation of a structural stochastic volatility model of asset pricing. *Computational Economics*, 38(1):53–83.

Franke, R. and Westerhoff, F. (2012). Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 36(8):1193–1211.

Franke, R. and Westerhoff, F. (2016). Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation. *Journal of Economic Interaction and Coordination*, 11(1):1–34.

Frankel, J. A. and Froot, K. (1990a). Exchange rate forecasting techniques, survey data, and implications for the foreign exchange market. Technical report, National Bureau of Economic Research.

Frankel, J. A. and Froot, K. A. (1987). Short-term and long-term expectations of the yen/dollar exchange rate: evidence from survey data. *Journal of the Japanese and International Economies*, 1(3):249–274.

Frankel, J. A. and Froot, K. A. (1990b). Chartists, fundamentalists, and trading in the foreign exchange market. *The American Economic Review*, pages 181–185.

Frazzini, A. (2006). The disposition effect and underreaction to news. *The Journal of Finance*, 61(4):2017–2046.

Gabaix, X., Gopikrishnan, P., Plerou, V., and Stanley, H. E. (2003). A theory of power-law distributions in financial market fluctuations. *Nature*, 423(6937):267–270.

Gaunersdorfer, A., Hommes, C. H., and Wagener, F. O. (2008). Bifurcation routes to volatility clustering under evolutionary learning. *Journal of Economic Behavior & Organization*, 67(1):27–47.

Genesove, D. and Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. *The Quarterly Journal of Economics*, 116(4):1233–1260.

Geweke, J. and Porter-Hudak, S. (1983). The estimation and application of long-memory times series models.

Gilli, M. and Winker, P. (2003). A global optimization heuristic for estimating agent based models. *Computational Statistics & Data Analysis*, 42(3):299–312.

Gopikrishnan, P., Meyer, M., Amaral, L. N., and Stanley, H. E. (1998). Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B-Condensed Matter and Complex Systems*, 3(2):139–140.

Gopikrishnan, P., Plerou, V., Gabaix, X., and Stanley, H. E. (2000). Statistical properties of share volume traded in financial markets. *Physical Review E*, 62(4):R4493.

Granger, C. W. and Ding, Z. (1995). Some properties of absolute return: An alternative measure of risk. *Annales d'Economie et de Statistique*, pages 67–91.

Grinblatt, M. and Keloharju, M. (2001). How distance, language, and culture influence stockholdings and trades. *The Journal of Finance*, 56(3):1053–1073.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, pages 393–408.

Gu, G.-F., Chen, W., and Zhou, W.-X. (2008). Empirical distributions of chinese stock returns at different microscopic timescales. *Physica A: Statistical Mechanics and its Applications*, 387(2):495–502.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, pages 667–686.

Han, B. and Hsu, J. (2004). Prospect theory and its applications in finance. *Imagine*.

Hansen, L. P. and Singleton, K. J. (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *The Journal of Political Economy*, pages 249–265.

He, X.-Z., Li, K., and Wang, C. (2016). Volatility clustering: A nonlinear theoretical approach. *Journal of Economic Behavior & Organization*, 130:274–297.

He, X.-Z. and Li, Y. (2007). Power-law behaviour, heterogeneity, and trend chasing. *Journal of Economic Dynamics and Control*, 31(10):3396–3426.

He, X.-Z. and Li, Y. (2017). The adaptiveness in stock markets: testing the stylized facts in the dax 30. *Journal of Evolutionary Economics*, 27(5):1071–1094.

Heath, C., Huddart, S., and Lang, M. (1999). Psychological factors and stock option exercise. *The Quarterly Journal of Economics*, 114(2):601–627.

Hill, B. M. et al. (1975). A simple general approach to inference about the tail of a distribution. *The annals of statistics*, 3(5):1163–1174.

Hirshleifer, D. (2001). Investor psychology and asset pricing. *The Journal of Finance*, 56(4):1533–1597.

Hirshleifer, D. and Shumway, T. (2003). Good day sunshine: Stock returns and the weather. *The Journal of Finance*, 58(3):1009–1032.

Holden, K. (2010). The emotions and cognitions behind financial decisions: The implications of theory for practice. *Center for Financial Security WP*, 10:4.

Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic dynamics and control*, 35(1):1–24.

Hommes, C. et al. (2017). Booms, busts and behavioural heterogeneity in stock prices. *Journal of Economic Dynamics and Control*, 80:101–124.

Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of computational economics*, 2:1109–1186.

Hsu, D.-A. (1977). Tests for variance shift at an unknown time point. *Applied Statistics*, pages 279–284.

Hull, J. C. (2006). *Options, futures, and other derivatives.* Pearson Education India.

Hurst, H. E. (1951). {Long-term storage capacity of reservoirs}. *Trans. Amer. Soc. Civil Eng.*, 116:770–808.

Jansen, D. W. and De Vries, C. G. (1991). On the frequency of large stock returns: Putting booms and busts into perspective. *The review of economics and statistics*, pages 18–24.

Jennings, R. H., Starks, L. T., and Fellingham, J. C. (1981). An equilibrium model of asset trading with sequential information arrival. *The Journal of Finance*, 36(1):143–161.

Jensen, M. H., Johansen, A., and Simonsen, I. (2003). Inverse statistics in economics: the gain–loss asymmetry. *Physica A: Statistical Mechanics and its Applications*, 324(1):338–343.

Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J., and Tivnan, B. (2013). Abrupt rise of new machine ecology beyond human response time. *Scientific reports*, 3.

Jondeau, E. and Rockinger, M. (2003). Testing for differences in the tails of stock-market returns. *Journal of Empirical Finance*, 10(5):559–581.

Jordan, D. and Diltz, J. D. (2004). Day traders and the disposition effect. *The Journal of Behavioral Finance*, 5(4):192–200.

Kahneman, D. (2011). *Thinking, fast and slow.* Macmillan.

Kahneman, D. and Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological review*, 93(2):136.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.

Kahneman, D. and Tversky, A. (1982). The psychology of preferences. *Scientific American*, 246(1):160–173.

Kahneman, D. and Tversky, A. (1984). Choices, values, and frames. *American psychologist*, 39(4):341.

Karpoff, J. M. (1987). The relation between price changes and trading volume: A survey. *Journal of Financial and quantitative Analysis*, 22(01):109–126.

Kendall, M. G. and Hill, A. B. (1953). The analysis of economic time-series-part i: Prices. *Journal of the Royal Statistical Society. Series A (General)*, 116(1):11–34.

Keynes, J. M. (1937). The general theory of employment. *The quarterly journal of economics*, 51(2):209–223.

King, W. I. et al. (1930). Index numbers elucidated.

Kirman, A. (1991). Epidemics of opinion and speculative bubbles in financial markets. *Money and financial markets*, pages 354–368.

Kirman, A. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, pages 137–156.

Kirman, A. and Teyssière, G. (2002). *Bubbles and long range dependence in asset prices volatilities*. Springer.

Kuznetsov, Y. A. (2004). Elements of applied bifurcation theory.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335.

Ladley, D. (2012). Zero intelligence in economics and finance. *The Knowledge Engineering Review*, 27(2):273–286.

Lamperti, F., Roventini, A., and Sani, A. (2018). Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90:366–389.

LeBaron, B. (2002). Building the santa fe artificial stock market. *Physica A*.

LeBaron, B. (2006). Agent-based computational finance. *Handbook of computational economics*, 2:1187–1233.

LeBaron, B., Arthur, W. B., and Palmer, R. (1999). Time series properties of an artificial stock market. *Journal of Economic Dynamics and control*, 23(9):1487–1516.

LeBaron, B. and Samanta, R. (2005). Extreme value theory and fat tails in equity markets. *Available at SSRN 873656*.

Li, B., Shang, W., Li, H., Huo, L., and Xu, S. (2014). Disposition effect in an agent-based financial market model. *Procedia Computer Science*, 31:680–690.

Lillo, F. and Farmer, J. D. (2004). The long memory of the efficient market. *Studies in nonlinear dynamics & econometrics*, 8(3).

Liu, Y., Cizeau, P., Meyer, M., Peng, C.-K., and Stanley, H. E. (1997). Correlations in economic time series. *Physica A: Statistical Mechanics and its Applications*, 245(3):437–440.

Liu, Y., Gopikrishnan, P., Stanley, H. E., et al. (1999). Statistical properties of the volatility of price fluctuations. *Physical Review E*, 60(2):1390.

Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management, Forthcoming*.

Lo, A. W. (2007). Efficient markets hypothesis.

Locke, P. and Mann, S. (2000). Do professional traders exhibit loss realization aversion.

Longin, F. M. (1996). The asymptotic distribution of extreme stock market returns. *Journal of business*, pages 383–408.

Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal*, 92(368):805–824.

Loretan, M. and Phillips, P. C. (1994). Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets. *Journal of empirical finance*, 1(2):211–248.

Lucas Jr, R. E. (1978). Asset prices in an exchange economy. *Econometrica: Journal of the Econometric Society*, pages 1429–1445.

Lussange, J., Belianin, A., Bourgeois-Gironde, S., and Gutkin, B. (2018). A bright future for financial agent-based models. *arXiv preprint arXiv:1801.08222*.

Lux, T. (1995). Herd behaviour, bubbles and crashes. *The economic journal*, pages 881–896.

Lux, T. (1996). The stable paretian hypothesis and the frequency of large returns: an examination of major german stocks. *Applied financial economics*, 6(6):463–475.

Lux, T. (1998). The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. *Journal of Economic Behavior & Organization*, 33(2):143–165.

Lux, T. (2008). Stochastic behavioral asset pricing models and the stylized facts. Technical report, Economics working paper/Christian-Albrechts-Universität Kiel, Department of Economics.

Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719):498–500.

Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: a microsimulation of interacting agents. *International journal of theoretical and applied finance*, 3(04):675–702.

MacKinnon, J. G. (2010). Critical values for cointegration tests. Technical report, Queen's Economics Department Working Paper.

Mandelbrot, B. et al. (1963). The variation of certain speculative prices. *The Journal of Business*, 36.

Mandelbrot, B. B. (1971). When can price be arbitraged efficiently? a limit to the validity of the random walk and martingale models. *The Review of Economics and Statistics*, pages 225–236.

Mandelbrot, B. B. and Taqqu, M. S. (1979). *Robust R/S analysis of long run serial correlation*. IBM Thomas J. Watson Research Division.

Marsh, T. A. and Wagner, N. (2004). Return-volume dependence and extremes in international equity markets. In *EFA 2003 Annual Conference Paper*, number 284.

Mason, D. M. (1982). Laws of large numbers for sums of extreme values. *The Annals of Probability*, pages 754–764.

Massaro, D. (2013). Heterogeneous expectations in monetary dsge models. *Journal of Economic Dynamics and Control*, 37(3):680–692.

Mastromatteo, I., Toth, B., and Bouchaud, J.-P. (2014). Agent-based models for latent liquidity and concave price impact. *Physical Review E*, 89(4):042805.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, pages 995–1026.

Mehra, R. and Prescott, E. C. (1988). The equity risk premium: A solution? *Journal of Monetary Economics*, 22(1):133–136.

Menkhoff, L., Rebitzky, R. R., and Schröder, M. (2009). Heterogeneity in exchange rate expectations: Evidence on the chartist–fundamentalist approach. *Journal of Economic Behavior & Organization*, 70(1):241–252.

Menkhoff, L. and Taylor, M. P. (2007). The obstinate passion of foreign exchange professionals: technical analysis. *Journal of Economic Literature*, pages 936–972.

Michenaud, S. and Solnik, B. (2008). Applying regret theory to investment choices: Currency hedging decisions. *Journal of International Money and Finance*, 27(5):677–694.

Moro, E., Vicente, J., Moyano, L. G., Gerig, A., Farmer, J. D., Vaglica, G., Lillo, F., and Mantegna, R. N. (2009). Market impact and trading profile of hidden orders in stock markets. *Physical Review E*, 80(6):066102.

Muermann, A., Mitchell, O. S., and Volkman, J. M. (2006). Regret, portfolio choice, and guarantees in defined contribution schemes. *Insurance: Mathematics and Economics*, 39(2):219–229.

Muzy, J.-F., Delour, J., and Bacry, E. (2000). Modelling fluctuations of financial time series: from cascade process to stochastic volatility model. *The European Physical Journal B-Condensed Matter and Complex Systems*, 17(3):537–548.

Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The computer journal*, 7(4):308–313.

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of finance*, 53(5):1775–1798.

Osborne, M. M. (1959). Brownian motion in the stock market. *Operations research*, 7(2):145–173.

Pagan, A. (1996). The econometrics of financial markets. *Journal of empirical finance*, 3(1):15–102.

Pape, B. (2007). *Asset allocation, multivariate position based trading, and the stylized facts*. Universitas Wasaensis.

Plerou, V. and Stanley, H. E. (2008). Stock return distributions: Tests of scaling and universality from three distinct stock markets. *Physical Review E*, 77(3):037101.

Polach, J. and Kukacka, J. (2017). Prospect theory in the heterogeneous agent model. *Journal of Economic Interaction and Coordination*, pages 1–28.

Rekik, Y. M., Hachicha, W., and Boujelbene, Y. (2014). Agent-based modeling and investors behavior explanation of asset price dynamics on artificial financial markets. *Procedia Economics and Finance*, 13:30–46.

Ritov, I. and Baron, J. (1995). Outcome knowledge, regret, and omission bias. *Organizational Behavior and human decision processes*, 64(2):119–127.

Samuelson, P. A. (1963). Risk and uncertainty: A fallacy of large numbers.

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *IMR; Industrial Management Review (pre-1986)*, 6(2):41.

Sbordone, A. M., Tambalotti, A., Rao, K., and Walsh, K. J. (2010). Policy analysis using dsge models: an introduction.

Schmitt, N. (2018). *Heterogeneous expectations and asset price dynamics*. Number 134. BERG Working Paper Series.

Selim, K. S., Okasha, A., and Ezzat, H. M. (2015). Loss aversion, adaptive beliefs, and asset pricing dynamics. *Advances in Decision Sciences*, 2015.

Shapira, Z. and Venezia, I. (2001). Patterns of behavior of professionally managed and independent investors. *Journal of Banking & Finance*, 25(8):1573–1587.

Shefrin, H. (2002). Behavioral decision making, forecasting, game theory, and role-play. *International journal of forecasting*, 18(3):375–382.

Shefrin, H. (2008). *A behavioral approach to asset pricing*. Elsevier.

Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of finance*, 40(3):777–790.

Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of economic perspectives*, 17(1):83–104.

Shiller, R. J. (2015). *Irrational exuberance*. Princeton university press.

Shimokawa, T., Suzuki, K., and Misawa, T. (2007). An agent-based approach to financial stylized facts. *Physica A: Statistical Mechanics and its Applications*, 379(1):207–225.

Simões, G., McDonald, M., Williams, S., Fenn, D., and Hauser, R. (2018). Relative robust portfolio optimization with benchmark regret. *Quantitative Finance*, pages 1–13.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118.

Simonsen, I., Jensen, M. H., and Johansen, A. (2002). Optimal investment horizons. *The European Physical Journal B-Condensed Matter and Complex Systems*, 27(4):583–586.

Stanley, M. H., Amaral, L. A., Buldyrev, S. V., Havlin, S., Leschhorn, H., Maass, P., SALINGER, M. A., and EUGENE STANLEY, H. (1996). Can statistical physics contribute to the science of economics? *Fractals*, 4(03):415–425.

Stephens, M. A. (1974). Edf statistics for goodness of fit and some comparisons. *Journal of the American statistical Association*, 69(347):730–737.

Tauchen, G. E. and Pitts, M. (1983). The price variability-volume relationship on speculative markets. *Econometrica: Journal of the Econometric Society*, pages 485–505.

Taylor, M. P. and Allen, H. (1992). The use of technical analysis in the foreign exchange market. *Journal of international Money and Finance*, 11(3):304–314.

Taylor, S. J. (2011). *Asset price dynamics, volatility, and prediction.* Princeton university press.

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial life*, 8(1):55–82.

Thaler, R. H. (2000). From homo economicus to homo sapiens. *Journal of economic perspectives*, 14(1):133–141.

Thaler, R. H. (2005). *Advances in behavioral finance*, volume 2. Princeton University Press.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131.

Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, pages 1039–1061.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.

Westerhoff, F. (2010). *A simple agent-based financial market model: direct interactions and comparisons of trading profits.* Springer.

Westerhoff, F. and Franke, R. (2013). Agent-based models for economic policy design.

Westerhoff, F. H. (2008). The use of agent-based financial market models to test the effectiveness of regulatory policies. *Jahrbücher für Nationalökonomie und Statistik*, 228(2-3):195–227.

Westerhoff, F. H. and Dieci, R. (2006). 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.

Wichern, D. W., Miller, R. B., and Hsu, D.-A. (1976). Changes of variance in first-order autoregressive time series models-with an application. *Applied Statistics*, pages 248–256.

Winker, P., Gilli, M., et al. (2001). *Indirect estimation of the parameters of agent based models of financial markets.* FAME International center for financial asset management and engineering.

Winker, P., Gilli, M., and Jeleskovic, V. (2007). An objective function for simulation based inference on exchange rate data. *Journal of Economic Interaction and Coordination*, 2(2):125–145.

Winker, P. and Jeleskovic, V. (2007). Dependence ofand long memory inexchange rate returns: statistics, robustness. Technical report, time aggregation. Technical report WP011-07, University of Essex, Colchester.

Wong, K. P. (2015). A regret theory of capital structure. *Finance Research Letters*, 12:48–57.

Wooldridge, M. and Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(02):115–152.

Working, H. (1934). A random-difference series for use in the analysis of time series. *Journal of the American Statistical Association*, 29(185):11–24.

Ying, C. C. (1966). Stock market prices and volumes of sales. *Econometrica: Journal of the Econometric Society*, pages 676–685.

Zeeman, E. C. (1974). On the unstable behaviour of stock exchanges. *Journal of mathematical economics*, 1(1):39–49.