DISS. ETH NO.:

AGENT-BASED MODELS TO UNDERSTAND, EXPLOIT AND PREVENT FINANCIAL BUBBLES

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH
(Dr. sc. ETH Zurich)

presented by

REBECCA WESTPHAL

MSc UZH ETH in Quantitative Finance

born on 19.08.1993

citizen of Germany

accepted on the recommendation of

Prof. Dr. Didier Sornette

Prof. Dr. Thorsten Hens

2021

Abstract

Financial crashes that follow the formation of bubbles are widely recognized as disruptive events with severe negative consequences for the economy and society. Therefore, this cumulative dissertation studies financial markets comprising bubbles and crashes. Specifically, the aim of this thesis is to study the diagnostic of financial bubbles and the predictability of crashes as well as the market impact of strategies that utilize this information to either exploit bubbles or prevent them. Realistic agent-based models provide the framework to simulate endogenous bubbles and to quantify the market impact of various agents. In addition, agent-based models of multi-asset markets are developed to understand how bubbles emerge across multiple assets.

The first part of the thesis endows an investor with the ability to diagnose bubbles and predict the time of the crash better than chance. The trader uses this information to arbitrage bubbles and maximize his wealth. Using an agent-based model with fundamentalists and noise traders, we find that the arbitrageur can outperform the other traders in the market. In small numbers, the arbitrageurs reduce market inefficiencies and stabilise the market. However, a large number of arbitrageurs in the market tends to destabilise the asset price, because their diagnostics of bubbles become increasingly self-referencing, leading to volatile fluctuations that are further amplified by the noise traders.

The second part of the thesis investigates the intended and unintended consequences of an agent who uses a bubble diagnostic with the goal to prevent bubbles and crashes. This agent represents a policy maker, such as a central bank, that uses open market intervention to fight market exuberance. Using the agent-based model with fundamentalists and noise traders, we find that the policy maker succeeds at preventing bubbles and crashes using countercyclical market interventions. In the agent-based model, the policy maker improves all analysed market return metrics, volatility, skewness, kurtosis and Value-at-Risk (VaR), without affecting the long-term growth of the asset price. Even for substantially miscalibrated long-term expected returns, the results are robust.

The third part of the thesis presents two generalizations of the previous agent-based model both containing multiple assets. A Fixed Income Market model is introduced that includes a stock and multiple bonds with different maturities. The prices are defined by supply and demand and the model can reproduce several empirical characteristics observed in Fixed Income Markets. The second market model consists of several stocks and one risk-free asset that are traded by fundamentalists and noise traders. The models can reproduce several stylized facts of financial markets, such as volatility clustering and fat-tails of the return distribution. The noise traders' opinion dynamics are implemented as an $O(n)$ model allowing for an accurate mathematical description of the emerging phenomena. The extended market model can explain the emergence of synchronization among financial bubbles.

Kurzfassung

Finanzcrashs, die auf die Bildung von Spekulationsblasen folgen, werden weithin als unerwünschte Ereignisse mit schwerwiegenden negativen Folgen für Wirtschaft und Gesellschaft gesehen. Daher untersucht diese kumulative Dissertation Finanzmärkte, die Finanzblasen und Crashs aufweisen. Konkret geht es um die Diagnostik von Finanzblasen und die Vorhersehbarkeit von Crashs, sowie die Marktauswirkungen von Strategien, die diese Informationen nutzen, um Blasen entweder finanziell auszunutzen oder sie zu verhindern. Realistische agentenbasierte Modelle bieten den Rahmen, um endogene Blasen zu simulieren und die Marktauswirkungen verschiedener Marktakteure zu quantifizieren. Um zu verstehen, wie Finanzblasen über mehrere Vermögenswerte hinweg entstehen, werden agentenbasierte Modelle von Märkten mit mehreren Assets entwickelt.

Der erste Teil der Arbeit stattet einen Investor mit der Fähigkeit aus, Finanzblasen zu diagnostizieren und den Zeitpunkt des Crashs besser vorherzusagen als durch Zufall. Der Anleger nutzt diese Informationen, um Blasen auszunutzen und sein Vermögen zu maximieren. Unter Verwendung eines agentenbasierten Modells mit Fundamental- und Noise-Tradern stellen wir fest, dass der Arbitrageur die anderen Investoren auf dem Markt übertreffen kann. In geringer Anzahl reduzieren die Arbitrageure Marktineffizienzen und stabilisieren den Markt. Eine große Anzahl von Arbitrageuren auf dem Markt tendiert jedoch dazu, den Preis zu destabilisieren, da ihre Diagnostik der Blasen zunehmend selbstbezogen wird. Dies führt zu volatilen Schwankungen, die durch die Noise-Trader weiter verstärkt werden.

Der zweite Teil der Arbeit untersucht die beabsichtigten und unbeabsichtigten Folgen eines Agenten, der die Diagnostik von Finanzblasen mit dem Ziel einsetzt, Blasen und Crashs zu verhindern. Dieser Agent repräsentiert einen politischen Entscheidungsträger, z.B. eine Zentralbank, der Marktinterventionen zur Bekämpfung von Marktüberhitzung einsetzt. Mit Hilfe des agentenbasierten Modells mit Fundamental- und Noise-Tradern stellen wir fest, dass es dem politischen Entscheidungsträger gelingt, durch antizyklische Marktinterventionen Finanzblasen und Crashs zu verhindern. Im agentenbasierten Modell verbessert der politische Entscheidungsträger alle analysierten Renditeeigenschaften, Volatilität, Schiefe, Kurtosis und Value-at-Risk, ohne das langfristige Wachstum des Preises zu beeinflussen. Selbst bei erheblich fehlkalibrierten langfristigen Renditeerwartungen sind die Ergebnisse robust.

Im dritten Teil der Arbeit werden zwei Verallgemeinerungen des bisherigen agentenbasierten Modells vorgestellt, die beide mehrere Vermögenswerte enthalten. Es wird ein Marktmodell vorgestellt, das eine Aktie und mehrere Anleihen mit unterschiedlichen Laufzeiten enthält. Die Preise werden durch Angebot und Nachfrage definiert und das Modell kann mehrere empirische Merkmale reproduzieren, die auf Anleihemärkten beobachtet wurden. Das zweite Marktmodell besteht aus mehreren Aktien und einem risikofreien Vermögenswert, die von Fundamental- und Noise-Tradern gehandelt werden. Die Modelle können mehrere stilisierte Fakten der Finanzmärkte reproduzieren, wie z.B. Volatilitätsclusterbildung und Fattails der Renditeverteilung. Die Meinungsbildung der Noise-Trader wird als ein $O(n)$-Modell implementiert, das eine genaue mathematische Beschreibung der aufkommenden Phänomene ermöglicht. Das erweiterte Marktmodell kann das Entstehen einer Synchronisation zwischen Finanzblasen beschreiben.

Acknowledgements

This dissertation has benefited from the input and support of countless people. First and foremost, I am extremely grateful to my supervisor Professor Didier Sornette for his valuable guidance and advice. Both as a person and as a scientist, I benefited greatly from his tireless support, inspiring ideas and extensive knowledge. A special thanks goes to Professor Thorsten Hens from University of Zurich for being the co-referee of my PhD.

I very much appreciate the fruitful collaborations with Antoine Kopp and Davide Cividino as well as Basile Despond and Jian-Hong Lin. I enjoyed the constructive and interesting discussions on the design and implementation of our agent-based models.

It was a pleasure to works with the present and former members of the Chair of Entrepreneurial Risk, Ali Ayoub, Chahat Bhatia, Ming Chen, Katharina Fellnhofer, Giuseppe Ferro, Jan-Christian Gerlach, Sandro Lera, Sumit Kumar Ram, Michael Schatz, Ran Wei, Spencer Wheatley, Ke Wu, and Dongshuai Zhao as well as the former master students Jean-Remy Conti, Emily Damiani and Sayuli Drouard. I am thankful for all of the excellent discussions and exchanges of ideas. Furthermore, I would like to acknowledge the help of Isabella Bieri, Adriana Schellenbaum-Lenner and Judith Holzheimer who have ensured an enjoyable and smooth administrative support.

I am especially grateful for the support from family and friends throughout the last years and would like to express my deepest appreciation to Nico Zimmermann, my sister Theresa and my parents Sabine and Hans-Jürgen Westphal for their continuous help and support over the years.

Contents

Abstract iii

Kurzfassung v

Acknowledgements vii

1 Introduction 1
- 1.1 Motivation . . . 1
- 1.2 Financial bubbles . . . 2
- 1.3 Agent-based models in finance . . . 5
- 1.4 Detailed overview of the thesis . . . 8
  - 1.4.1 Market impact and performance of arbitrageurs of financial bubbles in an agent-based model . . . 8
  - 1.4.2 How market intervention can prevent bubbles and crashes . . . 8
  - 1.4.3 Agent-based models of multi-asset markets . . . 8

2 Market impact and performance of arbitrageurs of financial bubbles in an agent-based model 11
- 2.1 Introduction . . . 11
- 2.2 The market model . . . 13
  - 2.2.1 Set-up of the investment universe (assets) . . . 14
  - 2.2.2 Fundamentalists . . . 14
  - 2.2.3 Noise traders . . . 16
  - 2.2.4 Dragon riders (DR) . . . 17
  - 2.2.5 Price setting via equilibrium of supply and demand . . . 19
  - 2.2.6 Parameters and initial values . . . 19
- 2.3 Properties of the market and performance of fundamentalists, noise traders and dragon riders . . . 21
  - 2.3.1 Properties of the price and wealth dynamics . . . 21
  - 2.3.2 Performance of fundamentalists, noise traders and dragon riders . . . 23
- 2.4 Impact of Dragon Rider strategy on the market price properties . . . 29
  - 2.4.1 Impact on frequency and size of bubbles . . . 29
  - 2.4.2 Impact of DR on the stability of the market . . . 32
- 2.5 Conclusion . . . 35
  - 2.5.1 LPPLS fits . . . 36
- 2.6 The Equilibrium Market Price . . . 37
- 2.7 Table with parameter values . . . 40

3 How market intervention can prevent bubbles and crashes 41
- 3.1 Introduction . . . 41
- 3.2 The market model . . . 44

3.2.1 Market set-up . . . 45
3.2.2 Fundamentalist strategy . . . 45
3.2.3 Noise trader strategy . . . 46
3.2.4 Dragon-slayer strategy . . . 48
3.2.5 Market clearing and price equation . . . 50
3.3 Impact of the dragon-slayer on the price dynamics . . . 50
3.3.1 General conditions of the simulations . . . 50
3.3.2 Dragon-slayers with negligible market impact . . . 51
3.3.3 Dragon-slayers with significant market impact . . . 52
3.4 Impact of the dragon-slayer on the traders' performance . . . 61
3.5 Conclusion . . . 65
3.6 Appendix . . . 66
3.6.1 The Equilibrium Market Price . . . 66
3.6.2 Parameter values . . . 69
3.6.3 Additional figures . . . 70

4 Agent-based models of multi-asset markets 75
4.1 Agent-based model generating stylized facts of Fixed Income Markets . . . 75
4.1.1 Introduction . . . 75
4.1.2 Model setup . . . 76
4.1.3 Model dynamics . . . 88
4.1.4 Analyses . . . 96
4.1.5 Conclusion . . . 101
4.1.6 Appendix . . . 101
4.2 Multi-asset financial bubbles in an agent-based model with noise traders' herding described by an n-vector Ising model . . . 109
4.2.1 Introduction . . . 109
4.2.2 The agent-based market model . . . 111
4.2.3 Time series analysis and stylized facts of financial markets . . . 122
4.2.4 Emergence and time synchronization of bubbles among the risky assets . . . 126
4.2.5 Conclusion . . . 134

Bibliography 136

References 136

Chapter 1

Introduction

1.1 Motivation

Financial markets such as stock, bond, commodity or foreign exchange markets, which provide the opportunity to trade financial assets, play a key role in our economy and society. Well functioning financial markets are essential to facilitate the flow of money by determining the value of investments via supply and demand. This ensures that money is directed towards where it creates most economic value and leads to an efficient allocation of resources. Companies can raise capital for example by publicly issuing shares at a stock market to position themselves for future economical and technological developments. The access to capital fosters innovations and helps companies to develop products, hire employees, grow their business and invest into plants and equipment. Financial markets also allow to share the risk a company faces between multiple investors. On the other hand, individuals who invest in the stock market can accumulate wealth, which may result in an increase of their consumption and a larger economic activity. The wealth effect describes the increase in an investor's consumption if the perceived wealth in terms of the market price of their asset grows. The additional spending attributed to capital gains in corporate equities is estimated to be up to 19 cents per dollar (Juster et al., 2006). This is advantageous for society by fueling economic growth.

While well-functioning financial markets are beneficial for the economy, malfunctioning financial markets can create serious harm. For example, financial bubbles and especially the consecutive crashes can severely threaten the economy. Historical examples of stock market crashes such as the great depression following the Wall Street Crash in October 1929 show that societies potentially suffer over years from the negative consequences of a stock market crash. In the four years after the crash, the stock market value decreased by 90% and one fifth of the U.S. banks had failed due to banking panics when many depositors simultaneously demanded their money back after loosing confidence in the solvency of the bank (Romer, 2003). The negative consequences provoked by the crash propagated far beyond the banking sector. The industrial production in the U.S. declined by 46.8% and GDP declined by 30% (Romer, 2003). Consequently, the unemployment rate increased from 3.2% of the civilian labor force in 1929 to 24.9% in 1933 (Lebergott, 1957). It took more than a decade until the economy fully recovered to its pre-crisis level (Bernstein, 1987). This example reveals the potentially devastating impact of financial bubbles and crashes and how important it is to mitigate the consequences and possibly prevent them. While the economic consequences of the more recent stock market crashes have been less severe than the above-mentioned example, there has been an increasing number of market turbulences in the last decades. This shows that the origin of bubbles is still not fully understood.

The main goal of this thesis is to contribute to a better understanding of financial bubbles, bubble diagnostic and the possibility to exploit or prevent them. More specifically, the following questions will be addressed: What are the mechanisms behind financial bubbles and which properties of financial markets or investors facilitate the emergence of bubbles? How do bubbles spread

across different assets? Can we diagnose bubbles and predict crashes? Assuming that investors can diagnose bubbles: What is the market impact of traders exploiting this information? Can we design policy measures that actively intervene in the market to prevent bubbles and crashes? These questions are approached using agent-based models (ABMs) of different financial markets such as a market consisting of a bond and a stock, a market with multiple stocks and one bond, and a Fixed Income market. Heterogeneous traders invest into these assets and define the assets' prices endogenously via demand and supply. The herding and momentum following of some traders creates bubbles in the asset prices. In general, ABMs describe complex systems by specifying the interactions and characteristics of heterogeneous agents at the micro level. Thus, ABMs provide the means to approach a realistic description of real economic and financial systems by allowing the introduction of arbitrary levels of complexity and heterogeneity. They are not constrained to stationary conditions and can therefore describe transient dynamics such as bubbles and crashes (Sornette, 2014). These properties distinguish ABMs as suitable testing environment to quantify the dynamics of an inherently nonlinear impact of trading strategies or a policy in the presence of complex feedback loops. In practice, monetary policies cannot be tested on a small sample of the population before they are introduced. Thus, despite their huge impact on the economy, it is not clear how individuals and institutions react to it. Realistic ABMs provide the opportunity to test and compare policies in an experimental set-up.

The following two subsections provide an overview of financial bubbles and ABMs, the two essential components of this thesis.

1.2 Financial bubbles

The diagnostic of financial bubbles is a controversial topic and it is frequently deemed impossible. Alan Greenspan famously said at a symposium sponsored by the Federal Reserve Bank of Kansas City: "As events evolved, we recognized that, despite our suspicions, it was very difficult to definitively identify a bubble until after the fact–that is, when its bursting confirmed its existence" (Greenspan, 2004). Similarly, Kindleberger and Aliber (2011) define a financial bubbles as "an upward price movement over an extended range that then implodes". Thus, they determine the bubble by the collapse following a spectacular growth. However, in order to provide the opportunity to react timely and appropriately to a bubble, it is necessary to identify a bubble before it bursts.

Bubbles are also frequently defined as a deviation of the asset price from its fundamental value, which is the intrinsic or "true" value of the investment. In this sense, bubbles are for example defined as a "price not equal to its market fundamentals for some period of time for reasons other than random shocks" (Rosser, 2000) or as "the part of asset price movement that is unexplainable based on fundamentals" (Garber, 1990). However, these definitions only shift the problem from defining a bubble to defining the fundamental value or defining a threshold that qualifies for example as "large deviation". A large variety of bubble definitions is available, but the definitions are often inconsistent. Gürkaynak (2008) concludes a review of different econometric tests for a bubble with the observation that these definitions are in fact often contradicting. The problem with most of these definitions is that bubbles are defined as an abnormal behavior compared to a normal one. However, neither the "abnormal" nor the "normal" is clearly defined and often set arbitrarily. Since financial bubbles occur in very different assets and markets, it is difficult to distinguish bubbles from time-varying fundamentals (Gürkaynak, 2008). In order to outline typical characteristics that are common among financial bubbles, the following section provides an overview of some of the largest and most famous bubbles in history.

One of the first well documented bubbles is the Dutch Tulip Bubble also known as the Dutch tulipmania. Originally imported to the Netherlands as a gift from the Sultan of the Ottoman Empire, a few particular types of tulips became very popular in the Netherlands and its neighboring countries. Tulip bulbs became the fourth most important good for export for the Netherlands (Ayres, 2020) and over decades people made money trading tulip bulbs. However, the increasing value

of the currency together with the increasingly prosperous Dutch economy made investors overly confident that the prices of tulip bulbs would continue to increase (Sornette, 2003). The hight of the tulipmania was reached from mid-1626 to February 1637. The trading of tulip bulbs had shifted from professionals to non-professionals and forward contracts were traded instead of the actual bulbs which had to stay in the ground during winter. At the peak of the bubble rare bulbs such as the Semper Augustus were sold at a price of 5500 guilders corresponding to a price of approximately $66000 today and even the price of common bulbs increased 20 times before returning to the original price after the crash (Garber, 1989).

The South Sea Bubble emerged around an international trading company in the 18th century. The South Sea Company was equipped with special trading rights in the South Sea by the British government. After the executives spread rumors about excessive prospective returns of the company due to the trade monopoly in combination with monetary expansion, the share price increased far above the economic value of the company in the years between 1716 and 1720 (Kindleberger & Aliber, 2011). The rise in the share price was accompanied by rising share prices of similar companies and a speculative mania in the whole country. Simultaneous to the build-up of the South Sea Bubble in Britain, the Mississippi bubble around the Mississippi Company evolved in France. Both bubbles crashed resulting in an economic crisis (Kindleberger & Aliber, 2011). South Sea Company itself was saved by the bank of England and continued to be profitable for the next hundred years (Ayres, 2020).

Two bubbles that occurred fifty years apart, but were similar in many aspects, are the canal bubble (1793-1795) and the railway bubble (1835-1850) in the U.K.. Both of these bubbles accompanied advancements in infrastructure development and were fueled by hopes that the revolutionary way of transport would open new markets and spectacular profits. The first canals were build by the merchants or mine owners who actually wanted to ship goods. However, the exaggerated expectations on profits lured speculators in as well. For example the share price of Grand Junction Canal increased from £100 to £472 within one month (Ayres, 2020). Most of the invested money was lost and some canals were never completed, but in the long term the country benefited from those canals that were finished. During the railway mania, a large number of new railway companies formed and the expected profits made many investors from the middle class invest for the first time which further pumped up the share prices. The crash was triggered by a poor harvest creating fear of a famine. The downward tilt of the prices in combination with concerns about the extend of new railway promotions by the government created a fear among investors and a downward spiral of the prices (G. Campbell & Turner, 2010).

A type of bubbles that existed at different times and at different places all around the world are real estate bubbles. An early example is the Florida housing bubble between 1900 and 1926. The construction of the Florida East Coast Railway in combination with the creation of new land by "draining the everglades" and newly built luxury hotels convinces many people from the north to travel or move to Florida (Ayres, 2020). The hype around Florida eventually faded and real estate prices began to fall after a combination of negative events in 1926 such as a railway embargo, a ship sinking in Miami harbor and a hurricane (Ayres, 2020).

In the "roaring twenties" following the first world war, the economy boomed due to the development of mass production and the spread of new technologies such as the car and the radio. However, the stock market exaggerated the economic growth because of the assumption of many that it would continue to grow forever (Gordon, 2005). The Wallstreet crash in 1929 triggering the great depression was one of the sharpest and most abrupt collapses of the stock market as described in section 1.1. Besides the initial stock market crash, the tight monetary policy of the Federal Reserve Board pushed the economy further into a recession (F. S. Mishkin & White, 2002).

The largest stock market drop within a single day occurred on October 19, 1987, the so-called Black Monday. Within a single day, the Dow Jones dropped by 22.6%. The crash occurred after an overall acceleration of the price return between 1985 and 1987. In the nine months before the crash stock prices in the U.S. increased by 31.4% (Sornette, 2003). Many possible reasons for

the stock market crash, such as computer trading, portfolio insurances, derivatives, illiquidity or overvaluation have been broad up, but Barro and Kamphuis (1988) concluded in their systematic analysis that there is no clear cause of the crash. This suggests that the crash had an endogenous origin rather than an external cause. The acceleration of stock prices that emerge from collective "crowd" behavior of many interacting agents was accompanied by systematic oscillatory-like corrections implying discrete scale invariance in the trading structure (Sornette & Johansen, 1997). The acceleration eventually resulted in a critical point at which the crash occurred.

The dot-com bubble growing in the late 1990s on the enthusiasm about the development of the world wide web shows many parallels to the bubble in the 1920s. The invention of new technology created expectations of future profits and in combination with a low inflation stimulated investments that became excessive and unsustainable over the years (Gordon, 2005). The share prices of companies related to the Internet, software, telecommunication or computer hardware diverged from the share prices of companies in the so-called "Old Economy". While the share price of stocks in the S&P's technology sector quadrupled over the years 1998 and 1999, the S&P 500 index only increased by 50% (Sornette, 2003). The NASDAQ composite index reached a high of 5133 on March 10, 2000, and crashed to 3300 a month later (Sornette, 2003). Both, the fast acceleration of price growth in the winter of 2000 as well as the crash in the consecutive year occurred without any substantial positive or negative fundamental news (De Long & Magin, 2006). Thus, the burst of the speculative bubble is rather the natural end of the boom (Sornette, 2003).

In an attempt to reduce the impact of the crash of the dot-com bubble, the Fed rate was reduced from 6.5% in 2000 to 1% in 2003 (Sornette & Woodard, 2010). The easy access to money together with the belief that home prices could never fall and the invention of new financial products, so called Collateralized Debt Obligations (CDOs) fueled a real estate bubble in the U.S. (R. J. Shiller, 2012). So called Mortgage-backed securities (MBS) allowed banks to sell tranches of real estate loans to investors similar to a bond passing over the risk of default. However, Sornette and Woodard (2010) and Sornette and Cauwels (2014) point out that the root of the speculation is much deeper than the real estate bubble in the U.S.. The accumulation and mutual reinforcement of multiple bubbles in the preceding decades created an illusion of a "perpetual money machine" resulting in an unsustainable and artificial growth of the stock market detached from economic growth that is based on productivity and innovation. The default of some house owners to repay their debt created a chain reaction of falling real estate prices and default on mortgages. With the bankruptcy of Lehman Brothers in September 2008, at that time the fourth-largest investment bank by asset size, a wave of asset sales was triggered that resulted in a global financial crisis (F. S. Mishkin, 2011). The fact that the relatively mild financial disruption from the housing market could transform into a global financial crisis exposed the large level of systemic risk embedded in the financial system in 2008 (F. S. Mishkin, 2011).

Timing the growing distrust in the banking system, a white paper was published in 2008 by the pseudonym author Satoshi Nakamoto to introduce Bitcoin, a crypto-currency, as a peer-to-peer payment system depicting an alternative to the traditional banking system (Nakamoto, 2019). As the price and popularity of Bitcoin increased, plenty of crypto-currencies based on the innovative blockchain technology were created. The super-exponential price growth was accompanied by bubbles and crashes as well as intensive debates about the intrinsic value of the currency. Some believed it would eventually drop to zero, while others believed in the enormous potential of the technology to replace our financial system and revolutionize our working lives. Wheatley et al. (2019) quantified the value of Bitcoin based on the number of users by applying the generalized Metcalfe's Law and found that the Bitcoin price heavily exceeded the calculated price on multiple occasions. Studies of the Bitcoin price and social dynamics such as web search queries or topic modeling of Bitcoin forum posts suggest that Bitcoin buyers have substantially been influenced by the fast increase of prices as well as news and social media (Gerlach et al., 2019). The Bitcoin bubble crashed in December 2017 from $20000 to $6000 in July 2018 (Wheatley et al., 2019). However, at the time of writing the Bitcoin price has re-build to $19000, approaching its all-time

high. Three years after the last crash the value of Bitcoin and other crypto-currencies is controversial and still debated.

The historical overview of financial bubbles shows that there are many common characteristics among bubbles. Typically, bubbles are initialized by a new opportunity such as an invention or development that sparks well-founded expectations for future profits. Early investments create a first price appreciation. Then, the investors develop an enthusiasm that becomes self-fulfilling (Sornette, 2003). They extrapolate the price growth into the future and expect it to continue to rise, imitate others who already realized profits and follow the crowd by reacting to news and social influence. The price growth accelerates due to the positive feedback until the asset becomes severely overvalued. In this unsustainable regime, it is only a question of time until the smallest trigger creates a crash. Any tiny negative event or news that would have been ignored in the upwards movement creates a downward spiral, in which investors become fearful and want to sell the asset quickly before the price can drop even further.

The realization that each bubble seems unique at the time when it is happening, but the underlying mechanisms among the bubbles are similar, brings us back to the question whether bubbles can be diagnosed in real time before they burst and whether the time of the crash can be forecasted significantly better than chance. In order to approach these questions, the Chair of Entrepreneurial Risk at ETH Zürich launched the Financial Crisis Observatory (FCO) in August 2008, which issues forecasts on bubbles and crashes for individual stocks as well as sectors, indices and commodities (`https://er.ethz.ch/financial-crisis-observatory.html`). The FCO uses the LPPLS model (Log-Periodic Power Law Singularity) to detect financial bubbles. The concept was introduced by Sornette et al. (1996) and elaborated in Johansen et al. (1999) among others. It describes that bubbles grow faster than exponentially due to positive feedback mechanisms of imitation and herding of investors. The price growth is accompanied by accelerating log-periodic oscillations describing the increased long-term volatility as the bubble grows, and the price increase results in a finite time singularity at which the bubble bursts. The track record of the FCO as well as several empirical publications suggests that bubbles can be identified before the crash (Zhou & Sornette, 2003; Jiang et al., 2010; Demirer et al., 2019; Forró et al., 2015). However, the FCO cannot answer the question what happens if investors exploit the knowledge about a crash in the near future to increase their wealth. There could be severe unintended consequences such as the amplification of bubbles or market turbulences. On the other hand, identifying a bubble in real-time provides the opportunity to intervene in the market to prevent the further built-up and crash of the bubble. However, these interventions could be ineffective or could threaten the long-term growth of the market. In order to test the market impact of different types of agents, it is necessary to describe the financial market by a model. The next section describes how this testing environment can be implemented.

1.3 Agent-based models in finance

During the financial crisis of 2007-2008, agent-based models (ABMs) became apparent as an essential tool to understand and describe financial markets. Traditional economic theory simplifies the complex structure of financial markets by modeling the market participants as if they were rational and had homogeneous beliefs. This concept can describe many market phenomena, but fails to describe systems close to bifurcation and phase transitions such as financial bubbles and crashes (Sornette, 2003). Bouchaud (2008) argues that a more realistic representation of financial markets, which takes into account behavioral aspects and empirical data, needs to replace the focus on axioms such as market efficiency, because these have no framework to explain "wild" markets. This framework can be provided by agent-based models (ABMs) which are a "bottom-up" approach to describe out-of-equilibrium phenomena taking heterogeneity and behavioral aspects into account. They belong to the broader class of computational economic models and simulate the dynamic interactions and individual characteristics of multiple agents. Thus, they can incorporate heteroge-

neous beliefs and non perfectly rational agents without being constrained to stationarity conditions. Consequently, they are well suited to describe out-of-equilibrium market regimes, such as financial bubbles. The agents' interactions at the micro level can naturally lead to a complex structure and emerging phenomena at the macro level including phase transitions and tipping points. With these features, ABMs are not only useful to model financial markets. As Axelrod (2006) states, they address problems that are fundamental to many disciplines and especially useful when the math describing the interaction of many agents is otherwise intractable. Regarding financial markets, Lux and Westerhoff (2009) emphasize the importance of modeling economic interactions, because most economic problems of the globalized financial markets are emerging phenomena of complex societies. They further argue that microeconomic regularities observed in behavioral economics and strong interactions should build the foundation of more methodological flexible models for a systemic perspective on globalized economic systems. Furthermore, J. D. Farmer and Foley (2009) highlight that ABMs are necessary to provide a quantitative description of how the economy is likely to react to government policies under different scenarios. Thus, they can guide economic decision makers by testing policies in controlled and repeated simulations.

Historically, the development of ABMs is strongly linked to the development of computers and became broadly adapted in the 1990s when the widespread availability of computers allowed to run complex simulations. However, already in the late 1940s, Sakoda (1949) developed a model of racial segregation in his PhD thesis that consists of agents located in cells on a grid using coins. The model was later adopted in Sakoda (1971) and Schelling (1971) as described by Hegselmann (2017). They showed that a small preference to be surrounded by neighbours of the same type leads to a highly segregated society over time. One of the first simulations of financial markets was developed by Stigler (1964). He tested the impact of regulating the Security Exchange Commission on the American stock market based on empirical data and an artificial order book, in which the orders are random variables. Later, Kyle (1985) and Black (1986) developed ABMs of financial markets with traders that focus on trends and price patterns instead of the fundamental value of an asset and coined the term noise trader.

The first ABMs describing the interplay between rational fundamentalists and boundedly rational chartists were developed to understand specific market phenomena such as bubbles and crashes. Fundamentalists aim at inferring the fundamental value of the asset while chartists look for patterns in the price path or are influenced by the opinion of other traders. An early example is the ABM of Frankel and Froot (1986), which was developed to understand the dollar bubble in the early 80s. Motivated by the U.S. stock market crash in 1987, Kim and Markowitz (1989) developed a microstructure market model in which portfolio insurers and rebalancers invest in a stock or in cash. They demonstrate that the portfolio insurance strategies can destabilize the stock market and thereby provide a framework to explain the crash in 1987. However, later time-series analyses showed that the resulting time-series of these models cannot reproduce typical characteristics of financial markets. In a more elaborate version of the model by Frankel and Froot (1986), De Grauwe et al. (1995) showed that the interaction of fundamentalists and noise traders can lead to chaotic behavior of exchange markets.

Simultaneously to these first ABMs of financial markets, the closely related field of minority games evolved. The idea is based on the El Farol Bar Problem by Arthur (1994) in which each individual enjoys to go to a bar when there are less than 60% of the population otherwise they prefer to stay at home. However, they all have to decide simultaneously without knowing how many people will be at the bar. More generally, the minority game was formulated by Challet and Zhang (1997). Each player makes a binary choice and players who choose the same as the minority of players win. The agents have access to a shared memory of the outcome of the last few time-steps and build their own inductive strategies based on that information. The basic minority game was generalized to include Darwinian selection, in which inferior strategies are replaced by new ones over time (Challet & Zhang, 1997, 1998), and evolution including random mutations (Li et al., 2000a, 2000b). Satinover and Sornette (2007) use minority games to demonstrate the "Illusion of

control", which means that agents who optimize their strategy based on the available information may perform worse than random strategies . A thorough mathematical analysis of various strategy types for different extensions of the minority game can be found in Coolen (2005). Minority games were developed in order to reproduce stylized facts of financial markets based on simple rules. The class of models based on minority games can be granularly adopted to imitate the behavior of real markets or other complex systems. However, it turned out that they fail to reproduce some important stylized facts (Chakraborti et al., 2011). On the other hand, they became a useful tool to understand the physics of disordered systems (Cavagna et al., 1999) or crowd versus anti-crowd movements (Johnson et al., 1999).

A famous ABM in which traders forecast future returns using a combination of classifiers is the Santa Fe stock market by Arthur et al. (1996). The classifiers are subject to genetic modifications such as mutations and over time the successful strategies are maintained whereas the unsuccessful ones are replaced. The related research on "adaptive rational equilibrium dynamics" describes in a broader sense the interaction of agents with heterogeneous beliefs about the future asset price. In Brock and Hommes (1997) traders switch between predictors of the future return based on the past performance of these predictors.

De Long et al. (1990, 2007) describe one of the first ABMs in which financial bubbles are created by the noise traders' positive feedback on the asset price. They show that the presence of noise traders can explain several financial anomalies such as excess volatility. This model was later adopted by Duffy and Ünver (2006) to generate asset prices with bubbles and crashes reproducing the features of laboratory asset markets as the one described in V. L. Smith et al. (1988). It was adopted by Haruvy and Noussair (2006) to reproduce and explain the effect of short selling on bubbles and crashes observed in experimental asset markets.

Lux and Marchesi (2000) propose another ABM in which noise traders and fundamentalists coexist. However, these noise traders switch their opinion between an optimistic and pessimistic view of the market influenced by the other traders' opinions. Furthermore, the traders switch between being fundamentalists and noise traders by comparing the expected returns of the strategies. Due to the noise traders' herding behavior and the switching between strategies, the model can reproduce some of the stylized facts of financial markets such as excess kurtosis and volatility clustering. However, critics claim that the large number of parameters and complicated transition rules make the identification of the sources of the phenomena intractable (Chakraborti et al., 2011). The interaction of fundamentalists and this type of noise traders is further analysed by Wyart and Bouchaud (2007), who observe that fundamentalists tend to stabilize the market, while noise traders destabilize it. The binary choice of the noise traders builds on the analogy with spins in an Ising model that switch their state influenced by their neighbors' states and an external field. Ising models, originally introduced as a mathematical description of ferromagnetism, have a widespread application in social sciences. The simple representation of interacting elements with a finite number of possible states, is useful to represent social interactions and especially opinion dynamics. In the 1970s, Weidlich (1971) described decision making within social groups and the polarization of their opinion as a physical ensemble of interacting spins. Another famous Ising-based decision model is the voter model by Holley and Liggett (1975), in which the opinion of an agent is a binary variable stochastically changing under the influence of its neighbors' opinions. The Ising model was further used for example to represent the opinion dynamics of a strike process in a plant containing satisfied and dissatisfied workers (Galam et al., 1982). More advanced versions were developed to take into account for example the spacial location of individuals on a complex network (Grabowski & Kosiński, 2006). Reviews of statistical physics as a tool to represent social dynamics can be found in Castellano et al. (2009) and specifically for finance in Sornette (2014).

More recently, Kaizoji et al. (2015) developed an ABM with fundamentalists and noise traders that is prone to develop bubbles that grow super-exponentially. Furthermore, the ABM can reproduce the most important stylized facts of financial markets such as volatility clustering and fat tails of returns. The fundamentalists invest such that they maximise the expected utility of their future

wealth. The noise traders are influenced by social imitation and momentum following. Similar ABMs with fundamentalists who optimize an expected utility have been developed for example by Xu et al. (2014) or in a multi-asset setting by Chiarella et al. (2009). However, the model by Kaizoji et al. (2015) differs in the definition of the noise traders. The noise traders' opinions towards the risky or risk-free asset are represented by spins in an Ising model. This allows one to have realistic dynamics of opinion formation as well as a mathematical description of the emerging phenomena.

1.4 Detailed overview of the thesis

This cumulative thesis is based on published work and working papers. All of the described projects are joint work with Prof. Didier Sornette. Chapter 2 is based on the first-authored publication Westphal and Sornette (2020b) and chapter 3 is based on the first-authored preprint Westphal and Sornette (2020a). Chapter 4 is a compilation of working papers. The first one is second-authored with Antoine Kopp as first author. The second one is realized as joint first author with Davide Cividino. The two projects were initialized as master theses. My contributions are the guidance throughout the development of the models and analyses, providing the primary code for the model implementation as well as writing substantial parts of the working papers and creating some of the figures.

1.4.1 Market impact and performance of arbitrageurs of financial bubbles in an agent-based model

This chapter analyses the consequences of predicting and exploiting financial bubbles in an ABM. The ABM consists of a risky and a risk-free asset as well as three different trader types: fundamentalists, noise traders and "dragon riders" (DRs). The DRs exploit their ability to diagnose financial bubbles from the endogenous price history to determine optimal entry and exit trading times. We study the DRs' market impact as a function of their wealth fraction. With a proportion of up to 10%, DRs are found to have a beneficial effect, reducing the volatility, value-at-risk and average bubble peak amplitudes. They thus reduce inefficiencies and stabilise the market by arbitraging the bubbles. At larger proportions, DRs tend to destabilise prices, as their diagnostics of bubbles become increasingly self-referencing, leading to volatility amplification by the noise traders, which destroy the bubble characteristics that would have allowed them to predict bubbles at lower fraction of wealth. Concomitantly, bubble-based arbitrage opportunities disappear with a large fraction of DRs in the population of traders.

1.4.2 How market intervention can prevent bubbles and crashes

Using the previously validated ABM with fundamentalists and chartists, this chapter investigates the usefulness and impact of direct market intervention. The policy maker diagnoses bubbles by forming an expectation of the future returns, then invests in burgeoning bubbles and sells countercyclically the overpriced asset to fight market exuberance. Preventing bubbles and crashes, this market intervention improves all analysed market return metrics, volatility, skewness, kurtosis and Value-at-Risk (VaR), without affecting long-term growth. This increases the Sharpe ratios of noise traders and of fundamentalists by approximately 28% and 45% respectively. The results are robust even for substantially miscalibrated long-term expected returns.

1.4.3 Agent-based models of multi-asset markets

This chapter derives two different generalizations of the original ABM with fundamentalists and noise traders. First, it derives an extension that is modified to describe a Fixed Income market containing bonds of different maturities and a stock treated as perpetual bond. Second, a market model

with one risk-free and multiple risky assets is derived. The noise traders' opinion dynamics are implemented as an $O(n)$ model and the extended model can explain the emergence of synchronization among financial bubbles.

Agent-based model generating stylized facts of Fixed Income Markets

This section develops an agent-based model (ABM) of a financial market with multiple assets belonging either to the Fixed Income or to the Equity asset class. The aim is to reproduce main stylized facts of Fixed Income markets with regards to the emerging dynamics of the yield curves. The research is rooted in the market model proposed by Kaizoji et al. (2015) considering two types of traders: the rational and risk-averse investor referred to as the fundamentalist and the noise traders who invest under the influence of social imitation and price momentum. The investors involved in the present market model diversify their investments between a preferred stock equivalent to a perpetual bond and multiple bonds of selected maturities. Among these assets, a zero-coupon bond provides a constant rate of return, while the remaining coupon-paying bonds' prices are determined at each time step by establishing the equilibrium between the investors' demands and supplies. As a result, the market model provides an evolving yield curve impacted by the investments of the aggregated traders of each type. It is moreover capable of reproducing transient turbulent periods in the prices' time series and a humped term structure of volatility. Ultimately, the comparison of the emerging dynamics arising from different processes governing the evolution of the risk-free rate with those of the historical U.S. treasury market enables to distinguish the capacity of the setup implementing Vasicek's model of interest rates to reproduce the surface of autocorrelation of the individual bonds' yields' volatilities.

Multi-asset financial bubbles in an agent-based model with noise traders' herding described by an n-vector Ising model

This section introduces a multi-asset model of a financial market in which rational fundamentalists and trend following noise traders co-exist. The interactions and opinion formation of the noise traders are described by an $O(n)$ model. The n components of the generalized Ising model represent the different assets in which the traders can invest. Rejection-free transition probabilities are derived to describe realistic investment decisions at the "micro" level. The agent-based model is validated by testing for several characteristics of financial markets such as volatility clustering and fat-tails of the distribution of returns. Furthermore, the model is prone to develop bubbles and crashes. We distinguish three different regimes depending on the inverse temperature that controls the traders' propensity to herd and imitate others. In the subcritical regime, the traders' opinions are idiosyncratic and no bubbles emerge. Around the critical value, partially synchronised bubbles emerge, triggered by an actual phase transition of the underlying $O(n)$ model. Above the critical value, synchronous bubbles emerge controlled by the time-varying external field representing the price momenta of the risky assets.

Chapter 2

Market impact and performance of arbitrageurs of financial bubbles in an agent-based model

This chapter analyses the consequences of predicting and exploiting financial bubbles in an agent-based model, with a risky and a risk-free asset and three different trader types: fundamentalists, noise traders and "dragon riders" (DR). The DR exploit their ability to diagnose financial bubbles from the endogenous price history to determine optimal entry and exit trading times. We study the DR market impact as a function of their wealth fraction. With a proportion of up to 10%, DR are found to have a beneficial effect, reducing the volatility, value-at-risk and average bubble peak amplitudes. They thus reduce inefficiencies and stabilise the market by arbitraging the bubbles. At larger proportions, DR tend to destabilise prices, as their diagnostics of bubbles become increasingly self-referencing, leading to volatility amplification by the noise traders, which destroy the bubble characteristics that would have allowed them to predict bubbles at lower fraction of wealth. Concomitantly, bubble-based arbitrage opportunities disappear with large fraction of DR in the population of traders.

Based on Westphal, R., & Sornette, D. (2020). Market impact and performance of arbitrageurs of financial bubbles in an agent-based model. *Journal of Economic Behavior & Organization, 171,* 1-23.

2.1 Introduction

What if a group of investors could develop reliable diagnostic tools to identify financial bubbles in the making and predict the likely time for their blow up? If successful, the obtained investment strategies would certainly flourish and could become progressively prevalent. How would this affect the financial markets? Would this lead to more efficient markets, with no more exuberant and breakdown phases? Or on the contrary, could this make markets more unstable? We investigate these questions with an agent-based model designed to generate endogenous bubbles and crashes from the interplay of fundamental value investors and noise traders who herd and follow price momentum. We introduce a third type of investors, called "dragon riders", who are endowed with ex-ante superior abilities to diagnose and to time bubbles and their bursts. By varying their market impact, and controlling for all properties of investors and of the price formation process, we run counterfactual experiments to answer these questions.

These questions are motivated by the research of the second author since 1996 (Sornette et al., 1996) and his group (see e.g. Sornette and Johansen (2001); Johansen and Sornette (2010);

Gerlach et al. (2019)), who have developed a research agenda to systematically test ex-ante the hypothesis that financial bubbles may be predictable (`http://www.er.ethz.ch/financial-crisis-observatory.html`). Given the interest from practitioners (obtained from private communications at professional financial investment conferences and workshops over the years), a natural question arises as to what could be the gains as well as unintended consequences on market properties of a wider adoption of bubble diagnostic techniques.

In an attempt to mimic real financial markets, we endow our dragon riders with an algorithm that tests for the presence of bubbles using a method called "log-periodic power law singularity" (LPPLS) (see e.g. Sornette (2003); Sornette and Cauwels (2015a) for pedagogical presentations), which has been widely tested on empirical data as mentioned above. The name "dragon rider" (DR) is based on the empirical observation that crashes that follow bubbles are exceptional events, outliers of strong significance (Johansen & Sornette, 2002, 2010), that the second author has proposed to name "dragon-kings" to emphasise their special status and specific amplifying mechanisms (Sornette, 2009) (for a pedagogical introduction, see `https://en.wikipedia.org/wiki/Dragon_King_Theory`). Thus, using calibrations of prices with a standard LPPLS methodology, the DR can sometimes obtain a diagnostic of the presence of a bubble and use this information to "ride" the ascending price bubble and they also exit the risky asset when they assess that the burst is close. In this way, they strive to maximize their risk-adjusted return by exploiting financial bubbles.

This present paper can be embedded in a broader literature, motivated by the global financial crisis in 2007-2008. In its wake, many financial professionals and researchers have realized the need to anticipate crises and have been working on the development of early warning signals. There are now a number of papers proposing early warning indicators of financial crisis. For instance, Drehmann and Juselius (2014) evaluate early warning indicators of banking crisis on a national level to define time-varying policy measures. Dawood et al. (2017) define an econometric model to predict sovereign debt crisis, and Shin (2013) compares early warning indicators based on market prices, normalized measure of total credit, and based on liabilities of financial intermediaries. However, while the focus of this research trend is on finding optimal warning signals, it remains an open question how trading and hedging according to additional knowledge, and exploiting bubbles and crashes impacts the occurrence and predictability of bubbles and crashes, and the stability of the market. This question is especially important when a large number of investors are using these early warning signals.

Traditional economic theory describes economic agents as if they were rational and had homogeneous beliefs. While this approach is able to describe many phenomena, it cannot describe systems close to bifurcation and phase transitions (Sornette, 2003). These phenomena can be described using agent-based models (ABMs). ABMs are computational economic models that simulate the individual operations and interactions of multiple agents. The first ABMs were developed by Sakoda (1949, 1971) and Schelling (1971) describing racial segregation of agents located in cells on a grid (Hegselmann, 2017). ABMs can incorporate heterogeneous beliefs and non perfectly rational agents and have a long tradition in economics, see Kirman (2012) for an overview. Thus, ABMs are not constrained to stationarity conditions and are well suited to describe out-of-equilibrium market regimes, such as financial bubbles. Simple interactions at the micro level can generate a complex structure at the macro level, as observed in real financial markets. The interaction of agents can naturally lead to regime shifts or tipping points, which makes this approach suitable for modeling extreme events in financial markets such as bubbles and crashes. Early ABMs of financial markets were discussed by Kyle (1985) and Black (1986). They describe agent-based models in finance that try to capture the behavior of traders who are more influenced by trends and price patterns than by the fundamental value of assets. The first model that uses the positive feedback caused by noise traders to create a market with financial bubbles was developed by De Long et al. (1990, 2007). Duffy and Ünver (2006) adopted the model by De Long et al. (1990) to generate asset prices with bubbles and crashes with the same features as in laboratory asset markets such as the one described in V. L. Smith et al. (1988). Haruvy and Noussair (2006) use a similar ABM to reproduce and

explain the effect of short selling on bubbles and crashes observed in experimental asset markets. Baghestanian et al. (2015) developed an ABM to identify different types of traders in experimental data and explain how the interplay of fundamental, speculative and noise traders creates bubbles in simulated and experimental data. Samanidou et al. (2007), Dieci and He (2018), and C. Hommes and LeBaron (2018) provide reviews of ABMs of financial markets. Recently, Kaizoji et al. (2015) developed an ABM consisting of both noise traders and fundamentalists, which is able to create a market with bubbles and also reproduces the main stylized facts of financial markets. The present paper is based on a modified version of this ABM, which is in particular extended to three different trader types, with the goal to analyse the impact that predictions of bubbles and crashes can have on financial markets.

Starting from the general framework of irrational bubbles (Schatz & Sornette, 2019), we study the market impact of bubble-based trading strategies. A specific model within this framework uses the concept of LPPLS (log-periodic power law singularities) as introduced by Sornette et al. (1996) and elaborated into a rational expectation model of bubbles in Johansen et al. (1999). Here, the expectation of the log-price is described as a log-periodic power law oscillations during a bubble. The pattern of log-periodic power law oscillations is taken as a clear indication of a bubble when it can be found with sufficiently strong statistical confidence (Sornette & Cauwels, 2015a). This framework focuses on price dynamics instead of financial ratios and provides the possibility of analyzing the consequences of trading with additional knowledge about bubbles and crashes. As already mentioned, the predictive power of LPPLS fits has been shown in a number of successful real time diagnostics of bubbles, such as the UK real estate bubble (Zhou & Sornette, 2003) and US real estate bubble (Zhou & Sornette, 2006), and the Chinese stock market bubble in 2005-2007 and 2008-2009 (Jiang et al., 2010). Furthermore, ex-post analyses of S&P500 index finds that LPPLS framework is able to capture some of the prominent bubbles, such as the Black Monday, Dot-com, and Subprime Crisis periods across different time scales (Demirer et al., 2019). A first approach to verify the predictive power of LPPLS fits in trading showed that LPPLS trading strategies outperform random strategies (Forró et al., 2015).

The next section defines the market model and trader types. Section 3 presents an evaluation of the comparative performance of the different trader types based on thousands of realisations of long time series of the price and traders' wealth. Section 4 analyses in details the market impact of the trading strategy of DR that predicts and exploits financial bubbles. We determine how the DR strategy affects the structure of the market in terms of the frequency and sizes of bubbles as well as in terms of different risk measures. Section 5 concludes. Three appendices complement the main text.

2.2 The market model

In order to analyse the impact of exploiting financial bubbles, the agent-based model (ABM) we use consists of two assets and three types of traders. The assets are a risky asset that pays dividends and a risk-free asset that pays a fixed interest rate. The model is an extension to three traders of a modified version of the ABM developed by Kaizoji et al. (2015). The model can reproduce stylized facts of financial markets such as fat-tails of the distribution of absolute returns, fast decay of the auto-correlation of signed returns, and slow decay of the auto-correlation function of absolute returns. Furthermore, it creates endogenous super-exponential bubbles (Kaizoji et al., 2015). We thus focus on this last characteristic for the present study and do not revisit the stylized facts.

The three trader types are noise traders (n), fundamentalists (f) and a novel third trader type that we refer to as "dragon riders" (DR) for reasons that become clear below. The noise traders make their investment decisions based on price momentum and social imitation. The fundamentalists are risk-averse rational investors, who maximize their constant relative risk aversion (CRRA) utility function. The dragon riders are fundamentalists equipped with some additional knowledge about bubbles, obtained using a methodology (LPPLS calibration) aimed at detecting local price accel-

eration, assumed to reflect transient unsustainable positive feedback loops. This methodology has been tested abundantly on empirical financial bubbles and is here applied for the first time on an ABM to study its market impact.

2.2.1 Set-up of the investment universe (assets)

Following Kaizoji et al. (2015), the investment universe available to the traders is made of a risk-free and a risky asset. In each time step the traders decide individually how they allocate their wealth between the two assets. The risk-free asset, which can be interpreted as a risk-free government bond or bank account with perfectly elastic supply, pays a fixed interest rate R_f in each time-step. The risky asset represents an index or risky investment fund, which pays a dividend d_t in each time-step. Its price is defined by the matching of demand and supply, as described below.

The dividend process d_t is a standard stochastic growth process, which reads

$$d_t = d_{t-1}\left(1 + r_t^d\right) \ , \tag{2.1}$$

where the growth rate r_t^d is a Gaussian process with mean value $r_d > 0$ and variance σ_d^2,

$$r_t^d = r_d + \sigma_d u_t, \tag{2.2}$$

with $u_t \overset{iid}{\sim} \mathcal{N}(0,1)$. This specification is different from that of (Kaizoji et al., 2015), in which the dividend price ratio was fixed, leading to a dependence of the dividend on the asset price, and the fundamentalists investing a constant fraction of their wealth in the risky asset.

2.2.2 Fundamentalists

The fundamentalists invest the fraction x_t^f of their wealth W_t^f in the risky asset and invest the remaining part of their wealth in the risk-free asset in order to maximize their expected utility in each time-step. This follows (Chiarella et al., 2006), C. H. Hommes and Wagener (2009), and Kaizoji et al. (2015). The fundamentalists are endowed with a constant relative risk aversion (CRRA) utility function with risk aversion parameter γ. This utility function is defined as (Ljungqvist & Sargent, 2018):

$$U(W) = \begin{cases} \log(W) & \text{for } \gamma = 1 \\ \frac{W^{1-\gamma}}{1-\gamma} & \text{for } \gamma \neq 1 \end{cases} \tag{2.3}$$

The difference between the return of the risky asset and the risk-free asset is the excess return. It is defined as the sum of the capital return $R_t := \frac{P_t}{P_{t-1}} - 1$ where P_t is the price of the risky asset at time t and of the return from the dividend d_t of the risky asset minus the risk-free rate R_f:

$$R_{excess,t} = R_t + \frac{d_{t-1} \cdot (1 + r_t^d)}{P_{t-1}} - R_f, \quad \text{with } R_t := \frac{P_t}{P_{t-1}} - 1 \ . \tag{2.4}$$

We obtain

$$E_{t-1}[R_{excess,t}] = E_{R_t} + \frac{d_{t-1} \cdot (1 + r_d)}{P_{t-1}} - R_f \ , \tag{2.5}$$

where

$$E_{R_t} := E_{t-1}[R_t] \tag{2.6}$$

is the fundamentalists' expected return of the risky asset. We take E_{R_t} as an exogenous variable, considered to derive from a secular view formed, for instance, from a long-term historical perspective on the long term growth of stock markets and of the underlying economy (Lera & Sornette,

2017). Thus, rationally, E_{R_t} should be equal to the long-term growth rate of prices. Consequently, the Fundamentalists' estimation of the fundamental value of the risky asset at time t is

$$P_t^f = P_0 \cdot (1 + E_{R_t})^t \ . \tag{2.7}$$

In our simulations, we verify that the long-term growth rate of prices is equal to the average growth rate r_d of dividends, as defined in expression (2.2). Figure 2.1 shows the long-term growth rate of the risky asset calculated as $frac1T - t \log \frac{P_T}{P_t}$ for values of r_d between 2% per year and 10% per year. Each value is calculated as the average over 1000 simulations with different random seeds and for two different processes for the parameter κ characterising the dynamics of noise traders: see below equation (2.16) for the definition of κ and expression (2.17) of the definition of the Ornstein-Uhlenbeck (OU) process for κ. The linear fit of the growth rate as a function of r_d has the slope 0.999 ± 0.002 for both OU κ, and constant κ. This confirms that the risky asset grows with the same rate as the dividend. In the sequel, E_{R_t} will thus be equated to r_d, which is a parameter of the model.

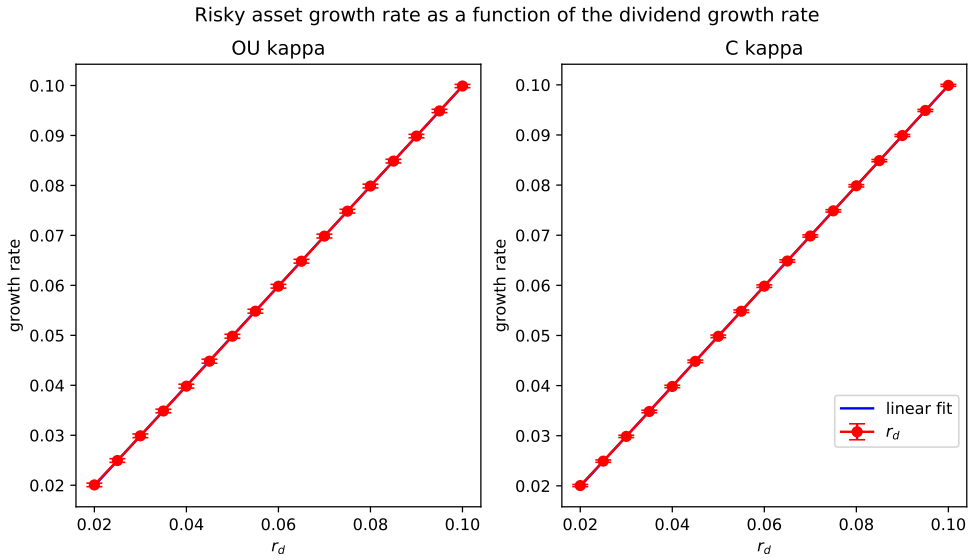

Figure 2.1: Growth rate of the risky asset as a function of the dividend growth rate r_d. The values are calculated over 240000 time-steps and averaged over 1000 simulations with different random seeds. The error bars indicate one standard deviation. The linear fit has slope 0.999 ± 0.002 for OU κ and constant κ (see below equation (2.16) for the definition of κ and expression (2.17) of the definition of the Ornstein-Uhlenbeck (OU) process for κ).

Fundamentalists are assumed to be identical, so that their cumulative behavior can be considered as one representative fundamental trader. The fraction of wealth invested in the risky asset with CRRA coefficient γ is in first order approximation and assuming $d_t \ll P_t$

$$x_{t-1}^f = \frac{1}{\gamma} \frac{E_{t-1}[R_{excess,t}]}{Var_{t-1}[R_{excess,t}]} = \frac{E_{R_t} + \frac{d_{t-1}}{P_{t-1}}(1 + r_d) - R_f}{\gamma(\sigma^2 + \frac{d_{t-1}^2 \cdot \sigma_r^2}{P_{t-1}^2})} \approx \frac{E_{R_t} + \frac{d_{t-1}}{P_{t-1}}(1 + r_d) - R_f}{\gamma \sigma^2} \tag{2.8}$$

The wealth gain from time $t-1$ to t by investing x_{t-1}^f in the risky and $1 - x_{t-1}^f$ in the risk-free asset is

$$W_t^f - W_{t-1}^f = \left(R_t + \frac{d_t}{P_{t-1}} \right) x_{t-1}^f W_{t-1}^f + (1 - x_{t-1}^f) W_{t-1}^f R_f. \tag{2.9}$$

Therefore, the wealth in time-step t can be written as

$$W_t^f = \left(R_t + \frac{d_t}{P_{t-1}} - R_f\right) x_{t-1}^f W_{t-1}^f + W_{t-1}^f (1 + R_f) \,. \tag{2.10}$$

Using the previous definitions and denoting the number of shares invested in the risky asset by $n_t^f := \frac{x_t^f W_t^f}{P_t}$, the excess demand of the fundamentalists for the risky asset is described by the following equation:

$$\begin{aligned} \Delta D_{t-1\to t} :=& n_t^f P_t - n_{t-1}^f P_t \\ =& x_t^f W_t^f - x_{t-1}^f W_{t-1}^f \frac{P_t}{P_{t-1}} \\ =& W_{t-1}^f \left(x_t^f \left[1 + R_f + x_{t-1}^f \left(R_t + \frac{d_t}{P_{t-1}} - R_f\right)\right] - x_{t-1}^f \frac{P_t}{P_{t-1}}\right) \end{aligned} \tag{2.11}$$

together with (2.6) and (2.8).

2.2.3 Noise traders

The noise traders define their strategy based on the opinion of other noise traders and the analysis of the price path. Each individual noise trader is polarized, in the sense that they do not diversify their portfolio and are fully invested in either the risky or the risk-free asset. Noise traders decide independently and probabilistically on their investment strategy. The probability of investing for each trader is influenced by the opinion of the other traders and the price momentum. With the number of traders N_t^+ invested in the risky asset and the number of traders N_t^- invested in the risk-free asset, the opinion index s_t is defined as

$$s_t = \frac{N_t^+ - N_t^-}{N_t^+ + N_t^-} \quad \in [-1, 1]. \tag{2.12}$$

The price momentum H_t is defined as the exponential moving average of the historical price returns defined in (2.4)

$$H_t = \theta H_{t-1} + (1 - \theta) R_t, \tag{2.13}$$

where $0 \leq \theta \leq 1$ controls the duration $\sim 1/(1-\theta)$ of the noise traders' memory.

Aggregating the independent investment decisions over all noise traders amounts to an equivalent representative noise trader who decides on the fraction x_t^n of his wealth invested in the risky asset given by

$$x_t^n = \frac{N_t^+}{N_t^+ + N_t^-} \quad \in [0, 1]. \tag{2.14}$$

Thus, putting together (2.12) and (2.14), we have

$$s_t = 2x_t^n - 1. \tag{2.15}$$

The probability at time-step t for a noise trader who is invested in the risky asset to switch his position to the risk-free asset is denoted as p_t^+. Respectively, the switching probability of a trader, who is invested in the risk-free asset, is denoted as p_t^-. These probabilities are given by

$$p_t^\pm = \frac{p_\pm}{2} \left(1 \mp \kappa_t (s_t + H_t)\right), \tag{2.16}$$

where the constants p_+ and p_- control the average holding time of each asset type. The time-dependent parameter κ_t determines the strength of social imitation and momentum following, which are assumed here to have the same quantitative impact.

One of the most important properties of financial markets is the existence of distinct regimes and of the switches between them. For instance, momentum-based investment strategies are known to perform very well in some periods of time and then to strongly underperform, and so on, in an irregular sequence (Lempérière et al., 2014; Taylor, 2014; Hwang & Rubesam, 2015). Similarly, financial markets are characterised by periods of exuberance alternating with periods of pessimistic mood (R. J. Shiller, 2006; Sornette, 2003). We propose to account for the existence of different regimes and for the random switches between them by allowing the noise traders to shift between periods when they have a large tendency to herd and when they have more heterogenous opinions. Thus, their susceptibility to herding is regime dependent. This incorporates the influence of exogenous factors such as economic and geopolitical regimes in the model. We account for these characteristics by allowing the coupling strength κ_t to be time-dependent according to a discretized Ornstein-Uhlenbeck process

$$\kappa_t = \kappa_{t-1} + \eta_\kappa(\mu_\kappa - \kappa_{t-1}) + \sigma_\kappa v_t. \tag{2.17}$$

where η_κ is the strength of mean reversion that controls the persistence time $\sim 1/\eta_\kappa$ of deviations from the mean μ_κ, driven by fluctuations with standard deviation σ_κ with $v_t \overset{iid}{\sim} \mathcal{N}(0,1)$. The expected value of the Ornstein-Uhlenbeck process κ_t starting at an initial value κ_0 is

$$E[\kappa_t] = \kappa_0 \cdot \exp(-\eta_\kappa \cdot t) + \mu_\kappa(1 - \exp(-\eta_\kappa \cdot t)). \tag{2.18}$$

Thus, the estimated time to revert from a value $\kappa_0 > \kappa_c$ to a value $\kappa_c > \mu_\kappa$ is

$$\Delta T = \frac{1}{\eta} \log\left(\frac{\kappa_0 - \mu_\kappa}{\kappa_c - \mu_\kappa}\right). \tag{2.19}$$

The aggregate wealth fraction of the noise traders invested in the risky asset then evolves as

$$x_t^n = \frac{1}{N_t^+ + N_t^-} \sum_{k=1}^{N_{t-1}^+} \left(1 - \xi_k(p_{t-1}^+)\right) + \frac{1}{N_t^+ + N_t^-} \sum_{l=1}^{N_{t-1}^-} \xi_l(p_{t-1}^-), \tag{2.20}$$

where $\xi(p)$ are Bernoulli random numbers. The corresponding aggregated wealth equation has the same structure as (2.10):

$$W_t^n = \left(R_t + \frac{d_t}{P_{t-1}} - R_f\right) x_{t-1}^n W_{t-1}^n + W_{t-1}^n \left(1 + R_f\right). \tag{2.21}$$

Combining the previous equations, the resulting aggregated excess demand from the noise traders for the risky asset is described by the following equation:

$$\Delta D_{t-1\to t}^n = W_{t-1}^n \left(x_t^n \left[x_{t-1}^n \left(R_t + \frac{d_t}{P_{t-1}} - R_f\right) + R_f + 1\right] - x_{t-1}^n \frac{P_t}{P_{t-1}}\right) \tag{2.22}$$

together with equations (2.20) and (2.16) with (2.15) and (2.17).

2.2.4 Dragon riders (DR)

The above description of the demand functions of the fundamentalists and noise traders is essentially just a slightly modified version of the ABM developed by (Kaizoji et al., 2015), with only an improved specification of the dividend process, as described above. The main novelty of the present paper is the addition of a third type of investors, the dragon riders, who are a kind of improved fundamentalists. Starting from their anchoring on the demand function (2.11) of the fundamentalists, the dragon riders attempt to obtain some additional knowledge about the market, by using bubble diagnostic tools. Because transient bubble-like behaviour has been documented in this ABM (Kaizoji et al., 2015) that mirrors quite well empirical financial bubbles, the ability to identify incipient bubbles with some success may provide a significant investment advantage, a kind of statistical

arbitrage, which in turn may modify the properties of the market as a result of the market impact of these investment strategies.

Specifically, our DR behave as fundamentalists as long as they do not have a good enough diagnostic or when they diagnose that the market is not in an exploitable bubble. Only when they infer that the market is in the early stage of a bubble, do they deviate from the fundamentalist strategy by buying more of the risky asset in order to "ride" the bubble. When they conclude that the bubble is close to burst, they sell the risky asset (no short positions are allowed). Quantitatively, the DR perform LPPLS calibrations at each time step and decide on how much to allocate to the risky asset depending on the result of the calibrations. The specific procedure controlling their investment decision is summarised in the flowchart shown in figure 2.2. and the following four steps correspond to the right branch of the flowchart ($t > t_{min}$).

Step 1 (LPPLS-fits): At any time-step t, the DR perform LPPLS fits over time windows $[t_1, t]$ where the largest fitting window is 500 time steps and the smallest window size is 20 steps and every i-th window size is considered. Afterwards, a filter $\mathcal{F}_1$, which is given in Table 2.3, is applied to the fit parameters to determine for each fit if it diagnoses a bubble. See 2.5.1 for a description of the LPPLS fits and the filtering.

Step 2 (Determining the beginning of the bubble): The beginning at time t_1^* of the bubble is determined using the Lagrange regularization approach introduced by (Demos & Sornette, 2019) and further tested by (Gerlach et al., 2019) on crypto-currency price time series. The goal is to find the fitting window with the most distinct LPPLS characteristics, while taking overfitting into account. This fitting window is determined by calculating the sum of squared residuals for each fit, performing a linear regression as a function of window size and subtracting the trend of this regression. The fit that minimizes this is the optimal fitting window, whose starting point is taken as being the beginning t_1^* of the bubble. The fits with fitting windows starting before the beginning t_1^* of the bubble are deleted.

Step 3 (Investment decision): The investment decision of DR is based on the fit starting at t_1^* and the confidence level of the other fits for shorter windows starting at later times until the smallest fitting window starting at $t - 20$. The confidence level is the fraction of fits that qualify a bubble. Depending on the fitted parameters, the DR determine if they are in the early state of a bubble, in the late state of a bubble, or not in a bubble. The confidence level was introduced by (Sornette et al., 2015) and elaborated in (Zhang, Sornette, et al., 2016) and (Zhang, Zhang, & Sornette, 2016). After the calibration, the DR decide on their desired position V_t at the time-step t. This means they decide if they just follow the fundamental strategy ($V_t = x_{t-1}^f$), buy ($V_t = 1$), or sell ($V_t = 0$) the risky asset. They want to buy the risky asset, if the fit starting at t_1^* qualifies a bubble, the confidence level is above the threshold $confThresh1$, and an additional parameter confirms the quality of the fit (the damping factor defined in Table 2.3 should be larger than $dampThresh1$, which measures the relative strength of the price acceleration with respect to a low-frequency measure of volatility). The DR sell the asset when the fit starting at t_1^* qualifies a bubble, the confidence level is above $confThresh2$, and the damping factor is above $dampThresh2$. In any other cases, they behave as fundamentalists. The confidence thresholds prevent the DR from investing to much based on spurious signals. Note that the damping threshold indicates the early and late stage of the bubbles.

Step 4 (Order execution): After a DR decides on his desired position $V_t \in \{x_{t-1}^f, 1, 0\}$, he executes his order linearly in time, with an execution rate $s(V_t)$ depending on the confidence level of the calibrations, and with a minimum order execution time τ so as to avoid too much market impact. When the DR diagnose a bubble, the rate $s(V_t)$ of the order execution is the previously described confidence level of the fits. This means that the DR execute their order faster when they are more confident about being in a bubble and execute their order slower when they are less certain. When they do not see a bubble, $s(V_t)$ is one minus the confidence level of the fits, because the confidence level indicates how certain the trader is to be in a bubble and in this case they want to execute their order faster when they are more certain not to be in a bubble. Thus, $s(V_t)$ describes the confidence that the DR have in their diagnostic.

In summary, the fraction of the wealth of DR invested in the risky asset at time t is given by

$$x_t^{DR} = x_{t-1}^{DR} + \frac{V_t - x_{t-1}^{DR}}{\tau} \cdot s(V_t) \, . \tag{2.23}$$

2.2.5 Price setting via equilibrium of supply and demand

The market clearing condition equilibrates supply and demand, which means that the market price is set according to Walras' theory of general equilibrium (Walras, 1954). Each trader formulates their excess demand for the next time-step and the price is calculated as the equilibrium in which supply equals demand. This is formulated as:

$$0 = \sum_i \Delta D_{t-1\to t}^{DR^i} + \Delta D_{t-1\to t}^{n} + \Delta D_{t-1\to t}^{f} \tag{2.24}$$

where $\sum_i \Delta D_{t-1\to t}^{DR^i}$ is the sum over the excess demand of the types i of the DR strategy that are included in the market. Next, the explicit formulations of the demand from (2.11), (2.22) and the demand of the DR obtained from the decision tree shown in figure 2.2 are inserted into the market clearing condition (2.24). This yields a quadratic equation for the price P_t. The positive solution of this equation is the price in the next time-step. The derivation and resulting price equation is given in 2.6.

2.2.6 Parameters and initial values

The parameters and initial values used in the simulations are derived in (Kaizoji et al., 2015) and reported in Table 2.4. The parameters are in particular set such that one time-step in the ABM represents approximatively one trading day. By this, we mean that the volatility (standard deviation of one time-step returns) matches approximately the daily volatility of typical developed financial markets, namely 1-2%. The details of this correspondence are given in (Kaizoji et al., 2015), which serves at the basis of the present work.

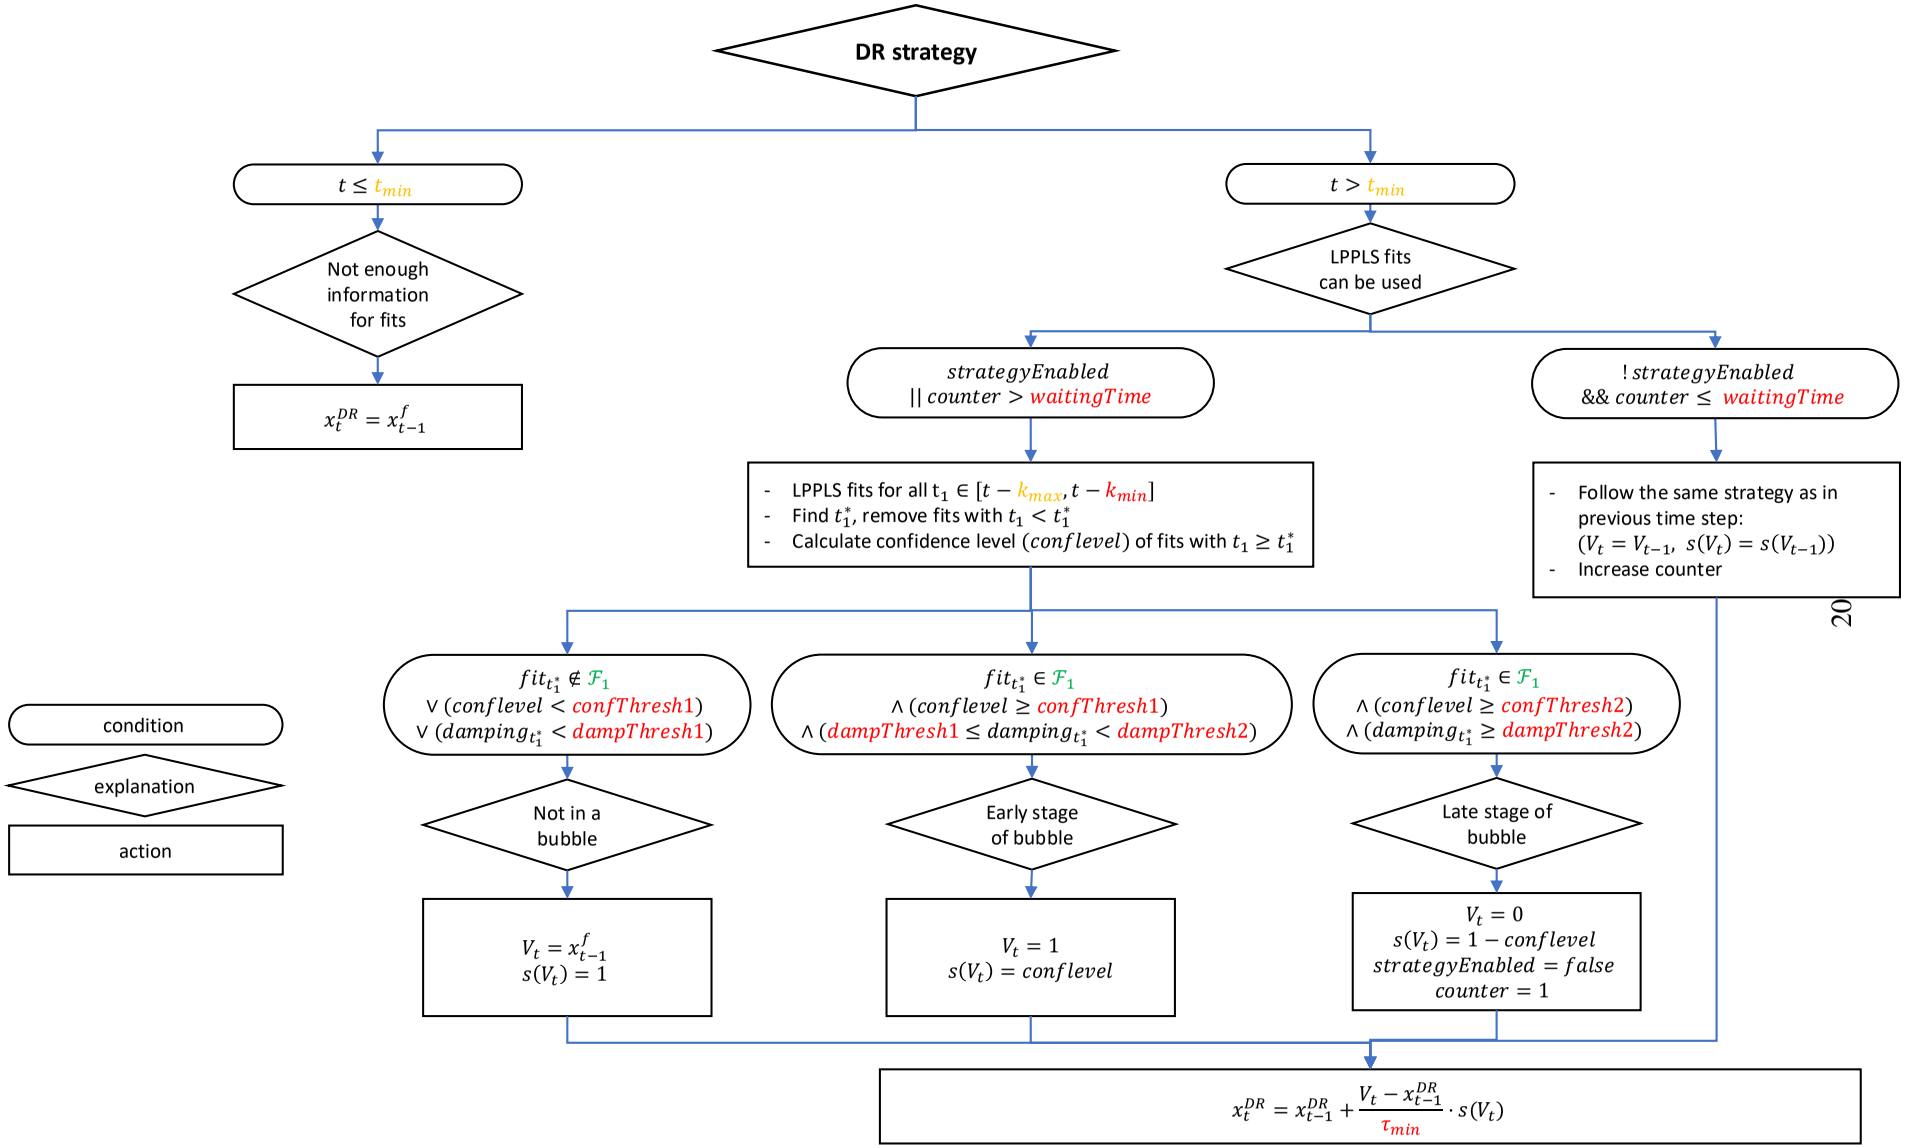

Figure 2.2: Flowchart of the DR's decision process to invest in a risky fraction x_t^{DR} at any time-step t.

2.3 Properties of the market and performance of fundamentalists, noise traders and dragon riders

2.3.1 Properties of the price and wealth dynamics

Using the parameters given in table 2.4 in Appendix A3, figures 2.3 and 2.4 show typical realisations of the price process, presenting the time evolution of the risky asset price (top panels), the fraction of wealth invested in the risky asset by the three types of traders (middle panels) and their wealths, this for two types of processes controlling the social imitation strength κ defined in equation (2.16): constant κ (figure 2.3) and varying $\kappa(t)$ according to an Ornstein-Uhlenbeck process such that the mean of $\kappa(t)$ is equal to the constant value of the first case (figure 2.4). 2000 time-steps are shown, which are sufficiently far in time after the burning period so that possible transient effects in the dynamics have decayed. Recall that the chosen parameters are such that one time-step corresponds approximately to one calendar day for a real market.

As shown in (Harras et al., 2012) and (Kaizoji et al., 2015), the dynamics of noise traders of the type presented in section 2.2.3 corresponds to a generalisation of standard Ising models, whose main characteristics is to present a phase transition or bifurcation that separates a disordered phase, where noise traders exhibit a large heterogeneity in their decisions to buy or to sell, leading to an average vanishing demand, from an ordered regime in which a majority opinion dominates, which can be either a strong net excess buy or sell demand. The presence of an Ising phase transition is important because it embodies the main property resulting from the spontaneous collective organisation of individuals who interact repeatedly. For instance, in these models, excess volatility results from the existence and proximity to such a critical phase transition (Harras et al., 2012).

Denoting κ_c, the critical value corresponding to the underlying Ising critical point of the noise trader dynamics, figure 2.3 has been generated with a constant $\kappa = 0.98\kappa_c$, which is sufficiently close to the critical point so that the collective decisions of noise traders can lead to bursts of volatility. Figure 2.4 had been constructed with a $\kappa(t)$ that fluctuates according to an Ornstein-Uhlenbeck process (AR(1) process in discrete time), around the same value $0.98\kappa_c$, with a standard deviation of the residuals sufficiently large so that $\kappa(t)$ wanders above κ_c in the organised phase of strong collective majority opinions among the noise traders. The parameter values are given in table 2.4.

In both cases, one can observe in the top panels of the two figures the development of bursts of volatility in the return time series. However, there is a striking difference between the two price trajectories, with structures, only present in figure 2.4, that are reminiscent of bubbles, with transient accelerated price growth followed by sharp drawdowns. The amplitude of these bubbles is typically of the order of 100%, with over-pricing and their correction developing over 100 time steps or less for the parameters used in this simulation. The development of these bubbles can be attributed to the emergence of a collective organisation of the noise traders that tend to buy or sell together, when the imitation strength becomes critical or super-critical for the generalised Ising model underlying the demand function of the noise traders (Kaizoji et al., 2015). For certain parameters, the model gives rise to bimodality in the distribution of the mispricing measured by the logarithm of the ratio of price to fundamental value. Similar bimodality has been observed in real data (Schmitt & Westerhoff, 2017; Majewski et al., 2020).

Comparing the middle panels, the noise traders are much more willing to buy the risky asset than fundamentalists, in both figures. The main difference lies in the exposure of the DR to the risky asset shown in figure 2.4, which is mostly in synchrony with the development of the bubbles. In contrast, in figure 2.3, the excess DR exposure to the risky asset above the value chosen by fundamentalists is either noisy or absent.

In the bottom panel of figure 2.3, the wealths of the DR and fundamentalists are growing almost equally, with small variations in the DR wealth around that of the fundamentalists. In contrast, the noise traders are over-reacting and either over-perform quite substantially or under-perform badly.

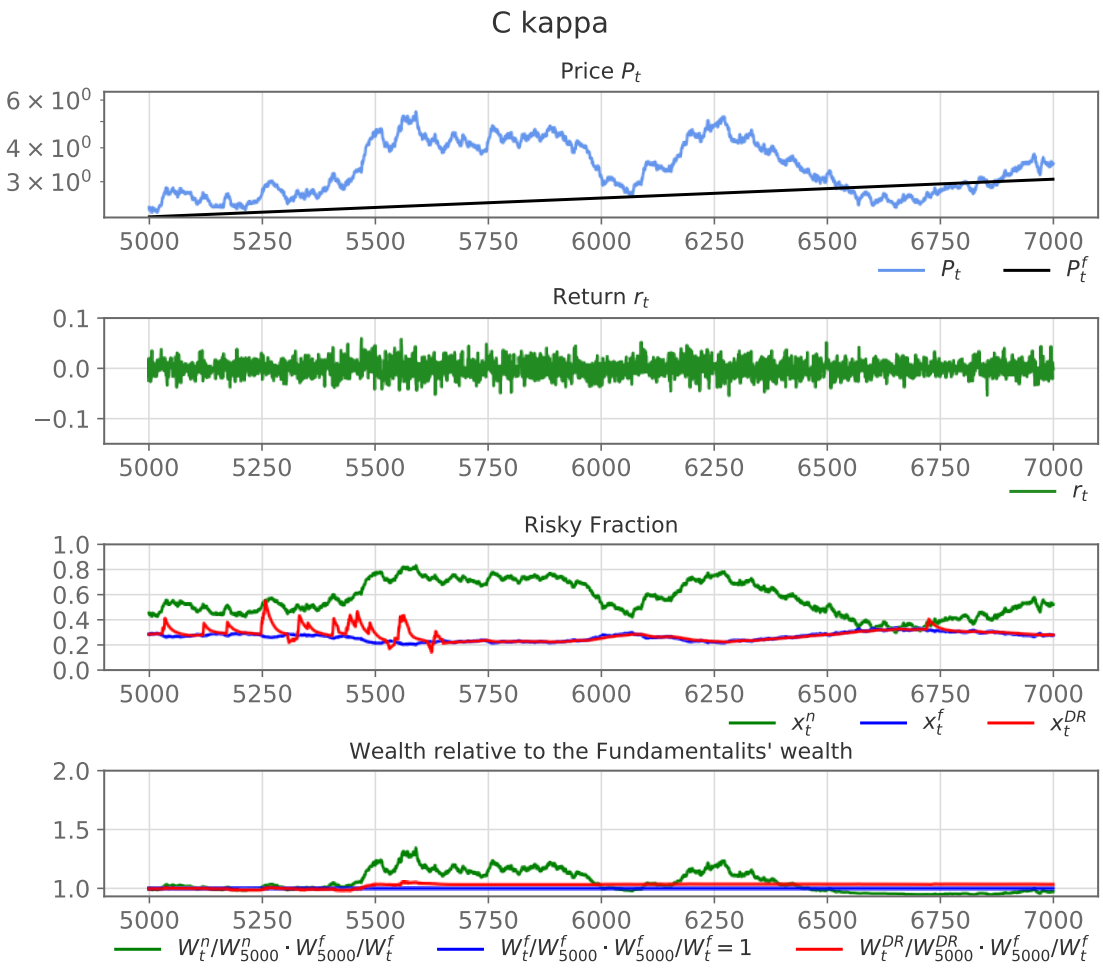

Figure 2.3: The simulation shown here has been generated with a constant $\kappa = 0.98\kappa_c$ close but below the critical value κ_c at which the interactions between noise traders lead to an Ising phase transition, as explained in the text. The upper panel shows a price path over 2000 time-steps, each time-step corresponding approximately to one trading day. The black continuous line shows the fundamental value of the risky asset given by expression (2.7) with $E_{R_t} = r_d$. The middle panel shows the fraction of their wealth that each trader type invests in the risky asset. The lower plot shows the wealth of each trader normalized to their value at the beginning of the time-interval and relative to the fundamentalists' wealth at each time-step t. The wealth given to the DR is very small so that their market impact is negligible.

At the end of the 2000 time steps, corresponding approximately to 8 years of trading, the noise traders find their wealth slightly below the fundamentalists and DR, notwithstanding having lived through periods with strong over-performance, which were followed by large drawdowns in their strategy. Their imitation and momentum-based investment generates a large wealth volatility. The lower panel of figure 2.4 demonstrates the existence of wild swings in the wealth of the noise traders, essentially in synchrony with the bubbles and crashes. In contrast, the wealth of fundamentalists is much smoother due to the more moderate exposition to the risky asset. Given their ability to diagnose the birth and death of bubbles, DR have a much better performance, with an almost steady and substantial growth compared to the fundamentalists. And they over-perform the noise traders significantly while exhibiting a much smaller volatility.

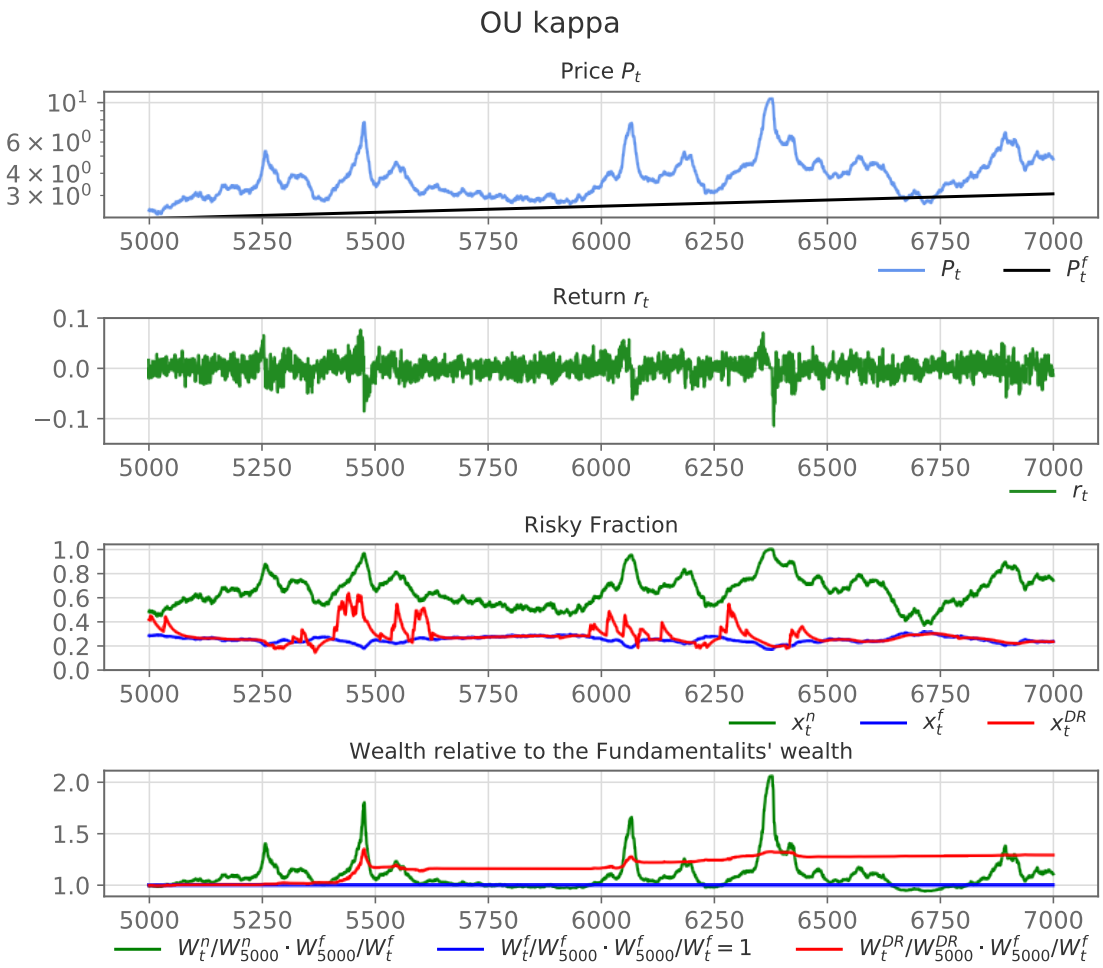

Figure 2.4: Same plots as figure 2.3 generated for $\kappa(t)$ that now depends on time according to an Ornstein-Uhlenbeck process fluctuating around $0.98\kappa_c$ with standard deviation of the residuals equal to $\sigma_\kappa = 0.01$ and the mean reversion of the social coupling strength $\eta_\kappa = 0.11$. Following expression (2.19), this process reverts within $\Delta T = 20$ time-steps from two standard deviation above μ_κ to the subcritical regime.

2.3.2 Performance of fundamentalists, noise traders and dragon riders

Performance when DR have negligible market impact

The performance of the three types of traders is quantified by the Sharpe ratio of their investment strategies, both in mean and in distribution over a large number of statistical realisations, and for different sets of parameters of the model. We study how these performances change (and degrade for the DR, notably) as the fraction of wealth held by the DR increases so that their market impact becomes important. We also compare the wealth dynamics of the fundamentalists and of the DR, both in terms of the mean and standard deviation over a large number of statistical realisations, at fixed model parameters.

The Sharpe ratio of an investment strategy, which is an average risk-adjusted return, is defined as (Sharpe, 1994)

$$S^i = \frac{\frac{1}{T}\sum_t r_t^i - r_f}{\sqrt{\text{var}(r_t^i - r_f)}}\,, \tag{2.25}$$

where r_t^i is the return of the trader type i from time $t-1$ to t, r_f is the constant risk-free rate, and T

is the length of the time interval over which the Sharpe ratio is calculated. The Sharpe ratios in the analysis are calculated from $t = 500$ to $t = 50000$. The first time-steps are not included to remove the effect of the transient behavior at the beginning of the time-series. For each set of parameters and conditions, we generate a statistical ensemble of 1000 time series from 1000 different random seeds.

Figure 2.5 shows the Sharpe ratio of strategies of the DR for constant social coupling $\kappa = 0.98\kappa_c$ and for the OU social coupling strength $\kappa(t)$, as a function of six parameters of the DR strategy. In each panel, one parameter is varied while the other parameters are fixed to default values given in Table 2.1. These Sharpe ratios for DR are compared to the Sharpe ratio of the fundamentalists. The noise traders have much smaller Sharpe ratios and are not shown.

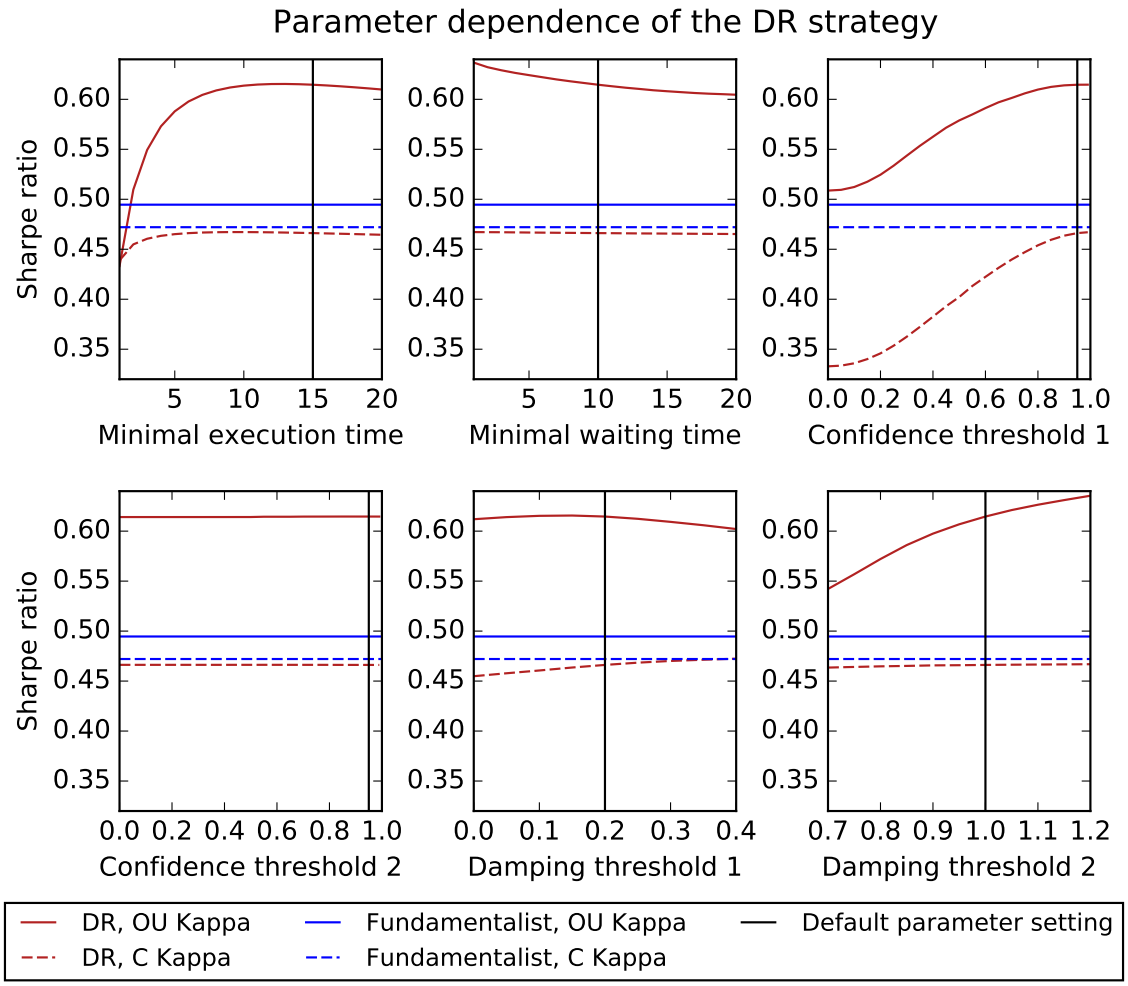

Figure 2.5: Sharpe ratio of strategies of the fundamentalists (blue) and of the DR (red) for constant social coupling $\kappa = 0.98\kappa_c$ (dashed lines) and for the OU social coupling strength $\kappa(t)$ (solid line), as a function of six parameters of the DR strategy. We do not show the Sharpe ratio of the noise traders, as it is much smaller: mean value of 0.241 for OU κ and 0.265 for constant κ. The Sharpe ratios shown here are obtained by averaging over 1000 simulations with different random seeds calculated over approximately 50000 time steps each. In these simulations, the initial wealth of the DR is an infinitesimal fraction of the total initial wealth, so that the market impact of the DR demand function is negligible. In each panel, only one parameter is varied, while the other parameters are fixed at the default values given in Table 2.1.

For constant social coupling $\kappa = 0.98\kappa_c$, since there are no bubbles and crashes to exploit, the DR are better off not trying to be smarter than the fundamentalists, as they can be misled in

wrongly diagnosing a bubble and then being hurt by over-weighting their portfolio on the risky asset. Hence, all the panels confirm that the DR always underperform the fundamentalists, and asymptotically reach the same performance when using parameters that amount to prevent them from diagnosing bubbles. This corresponds to have short execution time and large thresholds.

For OU social coupling strength $\kappa(t)$, the performance of the DR depends significantly on the minimal execution time, confidence threshold 1, and damping threshold 2. The Sharpe ratio is quite stable with respect to the other three parameters. In all analyzed parameter sets, except for a minimal execution time smaller than 2 time steps, the DR strategy significantly outperforms the fundamentalists. This clearly demonstrates that the DR strategy can exploit bubbles and crashes over a wide range of parameters and the specific strategy that they use is not at all essential.

Parameter name	**Explanation**	**Value**
t_{wait}	Minimum waiting time after exiting the strategy	10
τ	Minimal time for order execution	15
$confThresh1$	Confidence threshold for investing	0.95
$confThresh2$	Confidence threshold for late phase of bubble	0.95
$dampThresh1$	Damping threshold for early phase of bubble	0.2
$dampThresh2$	Damping threshold for late phase of bubble	1.0

Table 2.1: Default parameters used to defined the DR strategy, as described in section 2.2.4 and figure 2.2.

Figure 2.6 complements figure 2.5 by showing the distribution of Sharpe ratios of the three agent types in the two set-ups of a constant κ and OU $\kappa(t)$. The parameters defining the DR strategy are set at their default values given in Table 2.1. For constant κ, the fundamentalists are the winners, just slightly better than the DR, and well above the noise traders. In the absence of bubbles, the slight underperformance of the DR with respect to the fundamentalists stems from the few false diagnostics of bubbles, leading the DR in these instances to mostly overweight the risky asset in their portfolio. For the OU $\kappa(t)$ case, for which bubbles and large drawdowns often flourish, the DR perform well above the fundamentalists (and even more above the noise traders), both in distribution and for each specific realisation: in each of the 1000 simulations, the DR outperform the fundamentalists (and the noise traders of course) in terms of Sharpe ratios. This confirms the previous conclusion that the DR strategy can exploit bubbles and crashes, adding significant performance over the fundamentalist strategy.

In order to obtain credible asymptotic results, the Sharpe ratios in figures 2.5 and 2.6 are calculated over 49500 time steps, which corresponds to approximately 200 years. It is interesting to also compare the performance of the strategies of the DR and fundamentalists over time scales more in line with the typical time horizons used by investors to assess the quality of their investment. For this, figure 2.7 shows the time dependence of the mean value and standard deviation of the cumulative excess return of the DR with respect to the fundamentalists, calculated over 1000 realisations with parameters given in Table 2.1. This figure shows that, even for constant κ, the DR tend to overperform the fundamentalists, in terms of their cumulative return. This is because they sometimes misdiagnose a short-lived upward price fluctuation for a burgeoning bubble, which makes them invest more into the risky asset in the hope of "surfing" the bubble during its ascent. Notwithstanding the absence of bubble, the risky asset offers a higher average return than the risk-free rate, by construction. Hence, buying more of it gives over the long term a higher return, equal to the long-term rate r_d defined by equation (2.2). This comes however at the cost of a much larger volatility, hence the slight underperformance in terms of risk-adjusted return measured by the Sharpe ratio of the DR with respect to the fundamentalists for the prices generated with a constant κ, as shown previously in figures 2.5 and 2.6.

Figure 2.7 provides a complementary method to compare DR and fundamentalists. Indeed, as

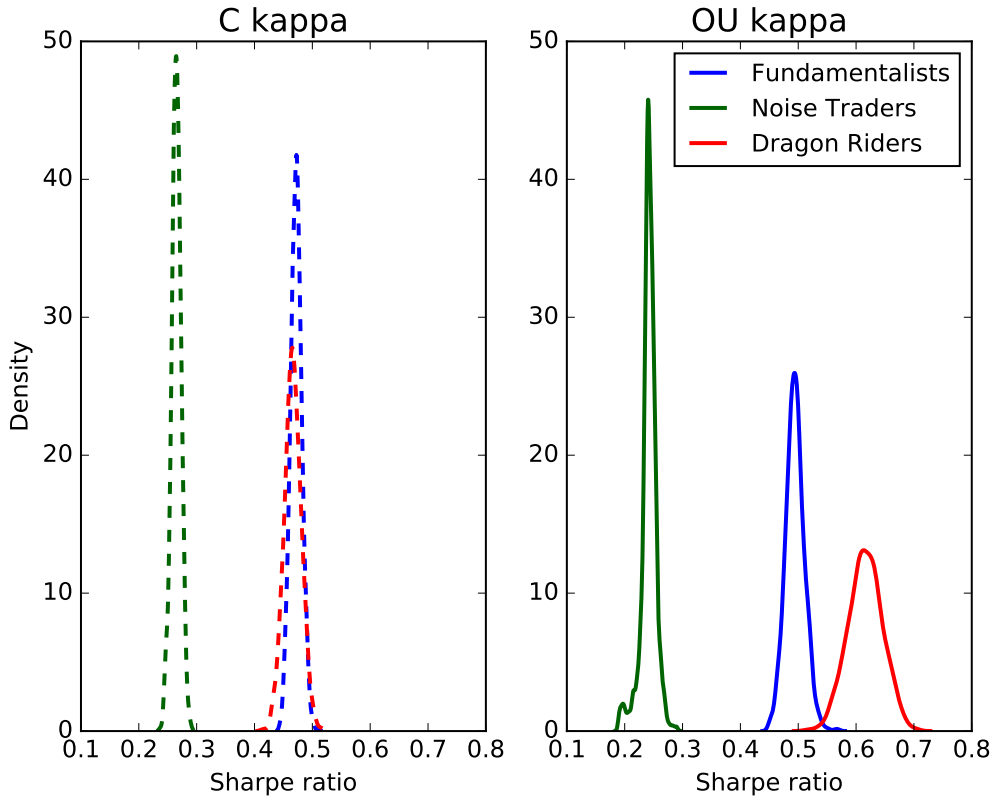

Figure 2.6: Distribution of the Sharpe ratios obtained from 1000 simulations with constant (left) and OU (right) social coupling strength κ and different random seeds for the three trader types. The mean Sharpe ratios for constant κ are respectively 0.265 for the noise traders, 0.472 for the fundamentalists and 0.466 for the DR. The mean Sharpe ratios for OU κ are respectively 0.241 for the noise trader, 0.495 for the fundamentalists, and the 0.616 for the DR.

reviewed in details in (Sornette et al., 2019), the value of an investment characterised by its mean return μ and standard deviation σ per time step can be assessed via the typical time $T_c \approx (\sigma/\mu)^2$ needed to decide whether an observed positive cumulative return is due to the positive systematic drift term μ (a proxy for skill) or to a lucky set of innovations for the random component with standard deviation σ. For constant κ, we find $T_c = 25733$, corresponding to about 100 years. This very long decision time confirms that the DR strategy for constant κ is more risky, and provides inferior value for an investor, on a risk-adjusted basis. For OU κ for which bubbles and crashes are defining attributes of the dynamics of prices, the cumulative excess return of DR with respect to fundamentalists grows much faster and $T_c = 1796$, corresponding to about 7 years. While still rather long, this time is typical of the time scale over which skill can be clearly distinguished from luck in real financial markets (we refer to (Sornette et al., 2019) and references therein for a detailed exposition and the relevant references).

Performance when DR have increasing market impact

All previous results have been obtained with a total wealth held by DR that is negligible compared to that of fundamentalists and noise traders. Hence, the results presented above have been obtained under the condition that the DR have no impact on the price. In other words, their decision to exploit bubbles and to sell the risky asset when they think a crash is coming does not have any impact on the price. This is realistic as long as the DR strategy is adopted only by a minority of

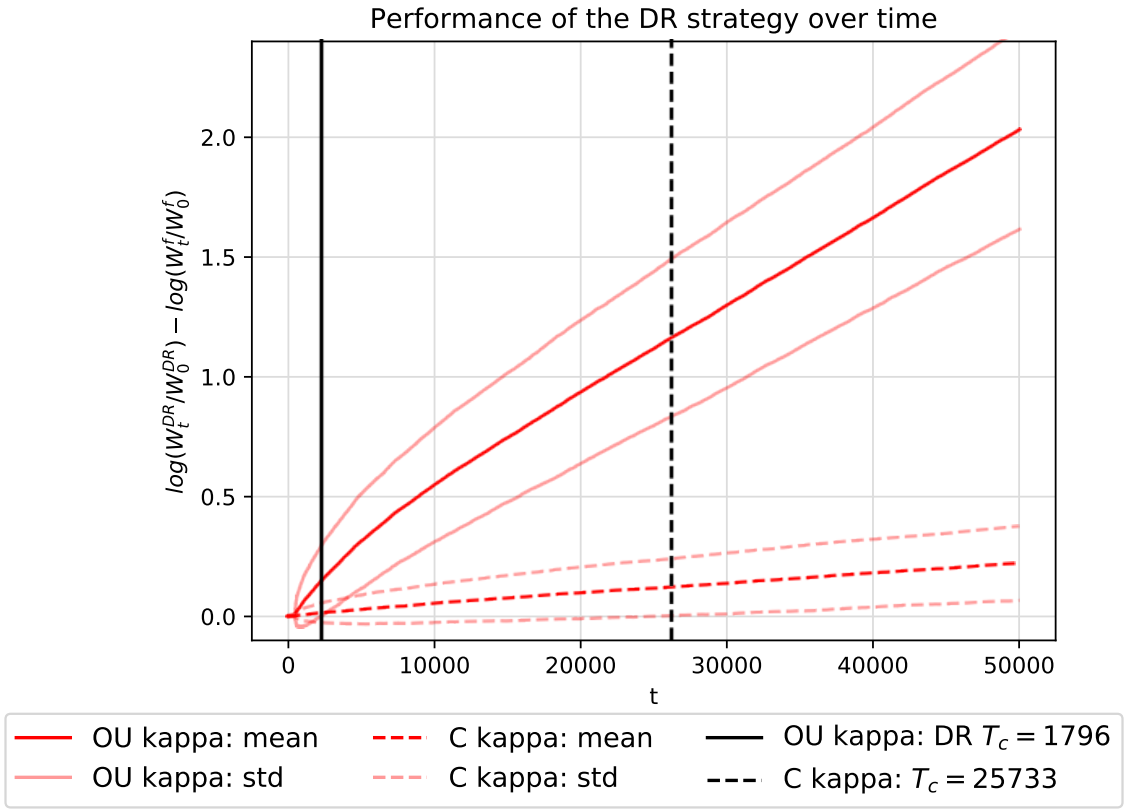

Figure 2.7: Mean value and standard deviation of the cumulative excess return of the DR with respect to the fundamentalists as a function of time, over 1000 realisations with parameters given in Table 2.1. The mean values and standard deviations are calculated over the set of 1000 simulations with OU and constant $\kappa = 0.98\kappa_c$, respectively. The vertical lines indicate the time T_c beyond which the mean of the cumulative return becomes larger than its standard deviation. T_c is evaluated by measuring the time when mean and standard deviation are equal, and subtracting to it 480 days, which is the length of the largest LPPLS fitting window that DR use in their algorithm to diagnose bubbles. Thus, from the measurement that mean and standard deviations become equal at $t = 26213$ for a constant κ, we obtain $T_c = 26213 - 480 = 25733$. i.e. about 100 years. Similarly, for OU κ, $t = 2276$ leads to $T_c = 2276 - 480 = 1796$, i.e about 7 years.

the fundamentalists. But, as the DR strategy is quite successful and over-performs significantly the fundamentalists (and even more the noise traders) for the OU κ case, it is likely that, in a real market populated by learning and adaptive agents, it would be progressively adopted. This is always the concern of those who invent a novel winning strategy: initially, it is known and used by a marginal number of investors; as its success becomes more and more visible, more and more traders adopt it, and the total size of buy and sell orders associated with the strategy become substantial, thus increasingly influencing the price dynamics. In general, this tends to decrease the value of the strategy, eventually resulting in its demise (J. Farmer, 2002). It is common lore in financial markets that the superiority of new trading strategies typically last only a few years at best, until the widespread adoption leads to the arbitrage of the anomalies that they exploit. This arbitrage mechanism is at the core of the efficient market hypothesis (Fama, 1970, 1991), which pushes to the extreme the logic that arbitrageurs remove exploitable features by using them in their strategies, which then change the nature of the price dynamics. This opens up two questions that we now investigate in turn in this subsection and in the next section: (i) how does the performance of DR change when their wealth fraction increases? (ii) how does the structure of the market evolve as the DR wealth fraction increases? In other words, do DR tend to decrease or increase the frequency

and/or size of bubbles and crashes?

To be concrete, we endow fundamentalists and noise traders each with an initial wealth of 10^9 monetary units (MU). With an initial price of $P_0 = 1.0$ MU and an initial fraction of wealth of 30% invested in the risky asset, this implies that the total number of shares is $6 \cdot 10^8$. The total number of shares remains constant throughout the simulation. With a fraction of wealth invested in the risky asset about twice as large as the fundamentalists risky fraction, the noise traders typically own $4 \cdot 10^8$ shares, while the fundamentalists typically own $2 \cdot 10^8$ shares.

In the first treatment, referred to as the 0% DR case, we provide the DR with an initial wealth of just 100 MU. Hence, they can buy at most 50 shares when the price is 2, which is completely negligible compared to the size of the portfolios of the fundamentalists and noise traders. In the following treatments, we keep the wealth fraction of noise traders constant at 50%, i.e. their initial wealth remains 10^9 MU, while the other 10^9 MU are divided among the fundamentalists and the DR. This is logically motivated by the fact that the DR use the fundamentalists' strategy by default when no bubble or crash is diagnosed. We vary the initial wealth of DR from 100 MU to $10^9 - 100$ MU (and accordingly the fundamentalists have from $10^9 - 100$ MU to 100 MU).

The dependence of the Sharpe ratios of the three trader types as a function of the fraction of wealth held by DR is shown in figure 2.8, both for the constant κ and OU κ cases. As expected, the performance of the DR decreases with their market size, in both cases. In the constant κ case, the Sharpe ratio of DR decreases steadily from just below that of the fundamentalists at 0% fraction to just above that of the noise traders for their maximum fraction at 50%. In the OU κ case, the DR Sharpe ratio remains above that of the fundamentalists up to a wealth fraction approximately equal to 7%, beyond which they underperform, while always remaining better than the noise traders. The Sharpe ratio of DR in the OU κ case parallels the Sharpe ratio of DR in the constant κ case for more than 20% wealth fraction given to the DR. In contrast, the average Sharpe ratios of fundamentalists' and noise traders' strategies do not change much over the full range of DR wealth fraction, varying between 0.48 and 0.52 for the fundamentalists and between 0.23 and 0.25 for the noise traders.

In sum, as expected, the DR perform best when they have a negligible fraction of the total initial wealth and therefore no market impact. They profit from diagnosing the price patterns associated to bubbles and crashes, which are created by the noise traders. But, when their market size becomes a significant fraction of the total market, their market impact modifies the very patterns that they are supposed to be good at detecting, leading to strong price distortions and progressive removal of the useful (for them) bubble signals.

To clarify the nature of the market impact of the DR, the next section analyses in details the characteristics of the alterations of the price patterns induced by the DR, for different values of their initial wealth fraction.

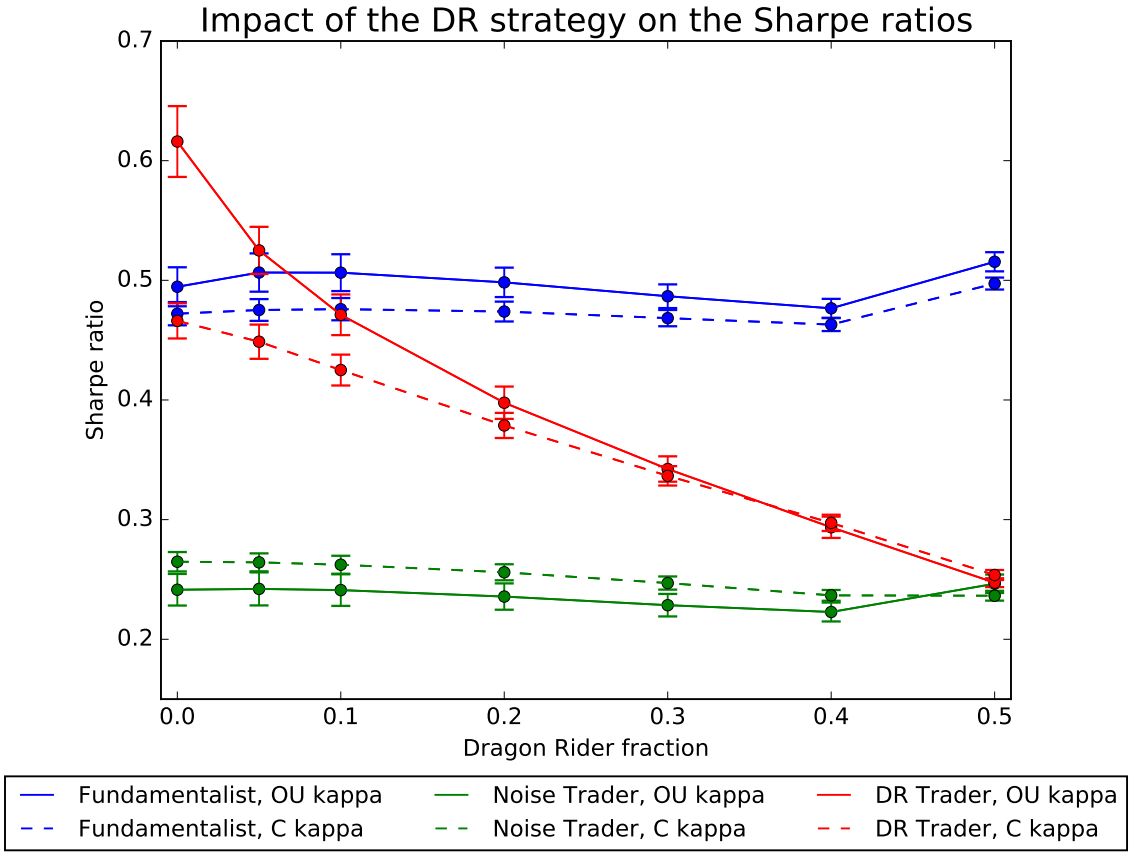

Figure 2.8: Sharpe ratios of the three trader types with an increasing fraction of DR for OU κ (solid lines) and constant κ (dashed lines). The error bars indicate one standard deviation calculated over 1000 realisations with different random seeds and over approximately 50000 time steps. The distributions of Sharpe ratios for negligible DR fraction (referred to as 0%) is shown in Figure 2.6.

2.4 Impact of Dragon Rider strategy on the market price properties

In this section, we analyse the impact of DR from the view point of a regulator or an external investor, who are interested in whether the addition of DR strategies is beneficial or harmful for the financial market. This question is investigated in two ways. The first one focuses on the effect of DR on the frequency and sizes of bubbles (subsection 2.4.1). The second one characterises the modifications brought to the markets by the DR in terms of expected return and standard deviation of the risky asset (subsection 2.4.2). As in subsection 2.3.2, we vary the wealth fraction of DR from 0% to 50% of the total wealth at inception of the market, with the fundamentalists' wealth fraction correspondingly decreasing from 50% to 0%, so that the sum of the two remains constant, ensuring also a constant 50% share of the initial wealth held by the noise traders.

2.4.1 Impact on frequency and size of bubbles

To study the impact of the DR on the frequency and sizes of bubbles, we first describe our methodology to diagnose bubbles ex-post and measure their sizes.

Identification of peaks and valleys and calculation of drawdowns

Different from the DR who attempt to diagnose bubbles in real time, here we have a much simpler task in identifying bubbles, because we can use the whole time series, so that a bubble can be

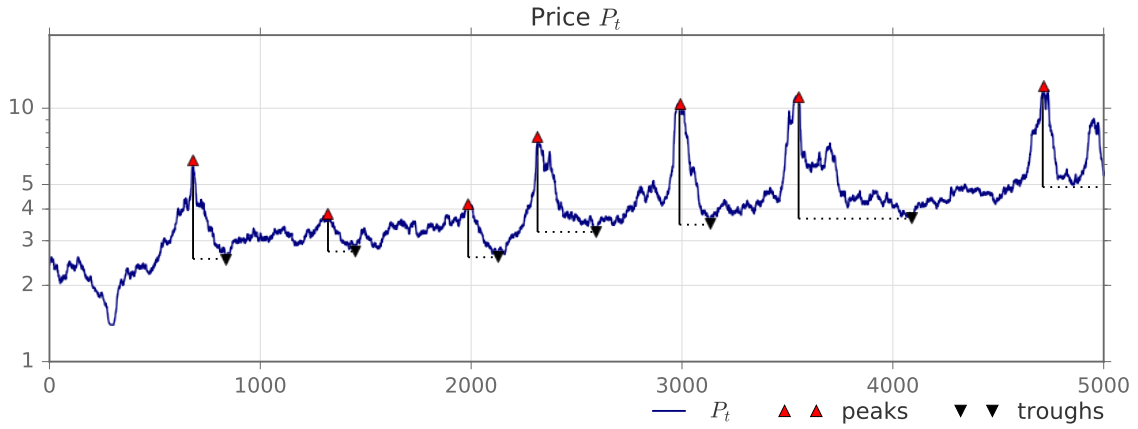

Figure 2.9: Example of a typical price path for the OU κ case over 5000 trading time steps (corresponding to 20 calendar years) with peaks (red filled triangles) and troughs (black inverted triangles). The minimum distance between peaks is chosen equal to $k = 250$. The lengths of the vertical line segments starting from the peaks represent the amplitude of the drawdowns, i.e., losses from peaks to valleys. The algorithm used to identify the peaks and troughs is described in subsection 2.4.1.

identified from both its abnormal rise and subsequent extreme drawdown.

The first step is to detect peaks and troughs in the price time series. A peak at time scale k in the price path occurs at time-step t_i if

$$P_{t_i} \geq P_{t_j} \quad \forall t_j \in [t_i - k, t_i + k], \tag{2.26}$$

where P_t is the price at time t and k is the minimum distance between two peaks. Thus, a peak occurs at a given time if the price at this time is larger than the price at the k previous and consecutive times. In this analysis, the minimum distance between to time-steps is chosen to be $k = 250$ trading days, which corresponds to approximately one calendar year.

The valley is defined as the time at which the price is minimal between two consecutive peaks. The size of a drawdown d_{t_i} is defined as the difference between the log-price at the time of the peak (t_i) and the log-price at the time of the consecutive valley

$$d_{t_i} = \log(P_{t_{peak}=t_i}) - \log(P_{t_{valley}}). \tag{2.27}$$

The detected peaks, valleys and drawdowns for a realisation of the price path with OU κ are shown in Figure 2.9. In this figure, we choose $k = 250$ trading days, which corresponds approximately to one calendar year. For this value $k = 250$ of the minimum distance between two peaks, all peak-to-valley drawdowns are at least 40% and often much larger, which makes them reasonable candidates for the typical sizes of corrections / crashes after bubbles. Hence, we take as a measure of frequency of bubbles simply the frequency of the peaks defined by (2.26) with $k = 250$. And the drawdowns are quantified as equal to the peak-to-valley amplitudes following these peaks, as shown by the vertical segments starting from the peaks.

Dependence of frequency of bubbles and drawdown sizes on DR's wealth fraction

The impact of the DR strategy on the average number of peaks for OU and constant kappa is shown in Figure 2.10.

For constant κ, apart from a slight decrease of the average peak-to-valley drawdown sizes from 0% to 5%, increasing the fraction of wealth held by DR has clearly an overall negative impact as both the numbers of peaks and their subsequent drawdowns, which proxy for bursts-and-busts and

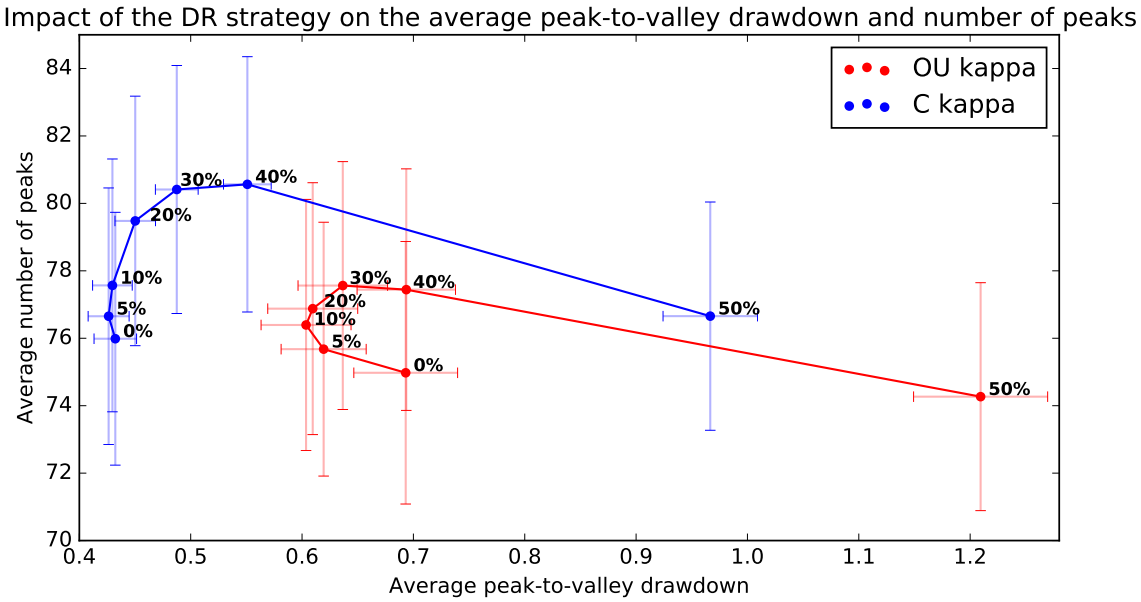

Figure 2.10: Number of peaks and average peak-to-valley drawdown in simulations with different DR wealth fractions. The DR wealth fractions vary between 0% and 50% as indicated on the curves, and each scenario is simulated 1000 times over T=50000 time-steps. The values are calculated for OU kappa (red) and constant kappa (blue). The error bars indicate the standard deviation of the number of peaks and peak-to-valley drawdown calculated over the 1000 realizations. With about 75 to 80 peaks occurring over 50000 time steps, this corresponds to one peak about every 2.5 years on average, quite larger than the minimum interval $k = 250$ time steps ($= 1$ calendar year) used to define the peaks, as shown in figure 2.9.

thus excess price movements, increase significantly. And the decrease in the number of peaks for the DR wealth fraction going from 40% to 50% is more than compensated by a huge increase in the average drawdown size from 55% to 95%. This is in line with the fact that the the DR misdiagnose the market fluctuations for the presence of bubbles, which lead them to overweight the risk asset and thus push its price even larger. Being non-sustainable, this is followed by a crash, and the process repeats.

In contrast, for the OU κ scenarios, for which real bubbles are created by sufficiently large herding and momentum following noise traders, the increasing weight of DR is beneficial for the market, at least up to about 10% wealth fraction, in decreasing significantly the amplitude of their subsequent drawdowns (up to 14%), however at the expense of a slight increase (2%) in the number of bubbles. When DR have more than 10% of the total wealth, their impact becomes detrimental in the sense of both increasing the number of bubbles and the amplitude of their subsequent drawdowns. When DR have completely replaced the fundamentalists (wealth fraction equal to 50%), this creates huge bubbles and crashes.

2.4.2 Impact of DR on the stability of the market

This subsection analyses how the initial wealth fraction of DR in the market influences the moments of the daily return (variance, skewness, excess kurtosis) as well as the value-at-risk (VaR), a measure of tail risks.

Figures 2.11 and 2.12 show typical realisations of the price process and the corresponding fractions of wealth invested in the risky asset for the three trader types, for different wealth fractions of DR from 0% to 50%. To be comparable, we use the same random number seed in all six shown simulations. Using the same random number seed allows us to perform counterfactual analyses of how a given market development would unfold in the presence of different fractions of DR, everything else, including the specific random numbers encoding the history of dividends as well as the decision of the noise traders, being exactly the same.

In the case of constant κ (Figure 2.11) with 0% DR, as already noted, the DR strategy only temporarily deviates from the fundamentalist strategy and returns quickly to the fundamentalist strategy after short-lived incorrect diagnostic of bubbles. As the fraction of DR increases, the typical length of the time intervals in which the DR strategy deviates from the fundamentalist strategy increases. The trading (buy and sell) of the DR in excess to that of fundamentalists leads to larger fluctuations in the price path, which has an amplification effect on the activity level of the DR strategy, through a positive feedback mechanism. The price trajectories become visually more volatile and bursty with larger DR wealth fractions.

Figure 2.12 shows the impact of the DR strategy in simulations with OU κ. Since the activity of the DR strategy is now more in tune with the detection of genuine bubbles, the impact of the DR strategy on the price trajectory can become beneficial. This is particularly visible for a DR wealth fraction of 20%, for which the large bubble and drawdown at time $t \sim 1200$ in the absence of DR impact has been essentially eliminated. For a DR wealth fraction of 10%, the amplitudes of the bubble and drawdown is decreased slightly but the DR create a second large bubble that did not exist in their absence. This reveals the stochastic nature of the interventions of DR, whose sometimes incorrect diagnostics or bad timing of bubbles may lead to create the very events they are supposed to arbitrage. For larger DR wealth fractions, the price dynamics develops some bubbles that do not have clear relations with the one in their absence. Eventually, for the largest fractions of 50% of DR, the price dynamics becomes very bursty as the positive feedback mechanism resulting from the impact of the actions of the DR completely dominates.

Table 2.2 quantifies the impact of the DR strategy with different metrics derived from the distribution of the returns in a market with OU and constant κ. We report (i) the variance, which is the second centered moment; (ii) the skewness, which is the third normalised centered moment and describes the asymmetry of the return distribution; (iii) the excess kurtosis, which is the normalised

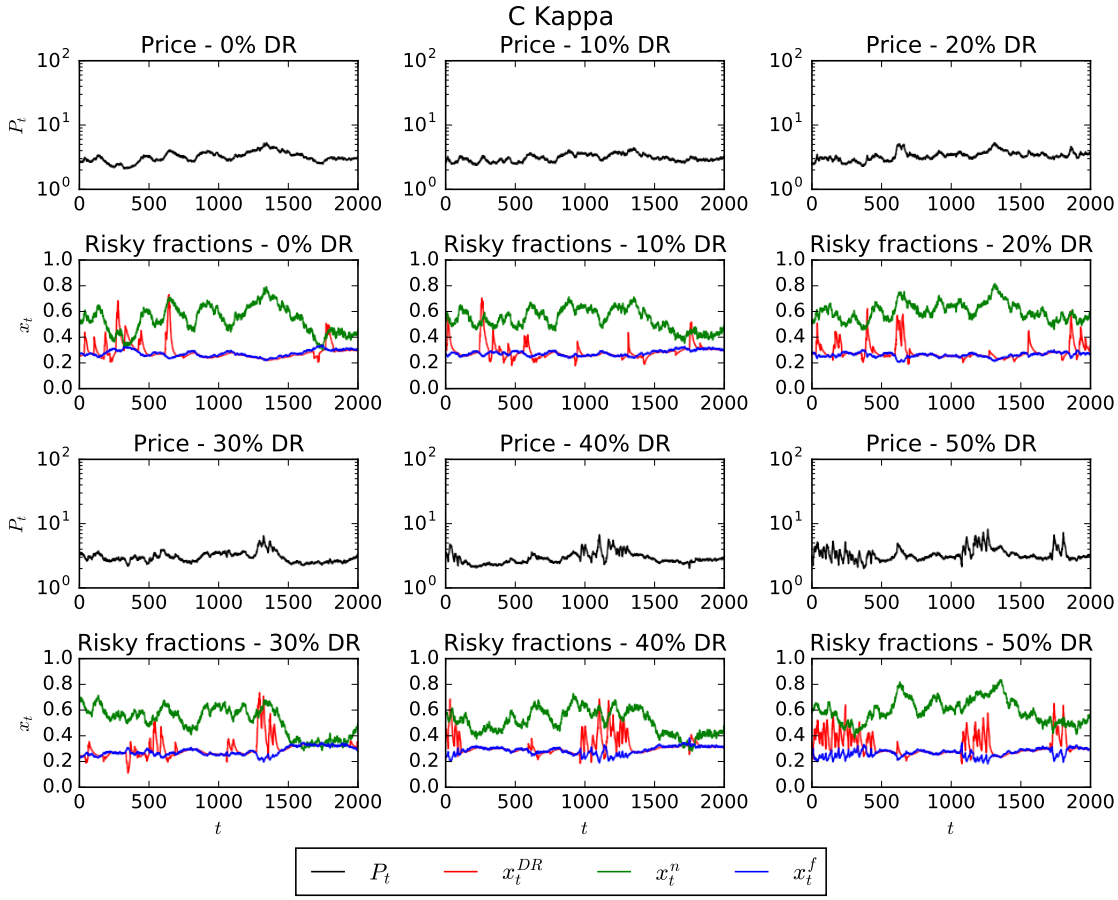

Figure 2.11: Section of a price path and the corresponding fractions invested in the risky asset for the three trader types. Each pair of panels shows a different DR wealth fraction in the market, from 0% (top left) to 50% (bottom right). The social coupling strength κ is constant.

	0% DR	5% DR	10% DR	20% DR	30% DR	40% DR	50% DR
Constant kappa							
Variance ($\times 10^{-4}$)	2.03($\pm$0.11)	**2.01**($\pm$0.11)	2.04($\pm$0.11)	2.25($\pm$0.12)	2.65($\pm$0.12)	3.44($\pm$0.13)	9.67($\pm$0.38)
Skewness	0.04($\pm$0.01)	0.05($\pm$0.01)	0.08($\pm$0.02)	0.23($\pm$0.04)	0.51($\pm$0.08)	0.92($\pm$0.12)	1.65($\pm$0.12)
Excess kurtosis	0.23($\pm$0.06)	0.26($\pm$0.07)	0.33($\pm$0.08)	0.84($\pm$0.17)	2.14($\pm$0.40)	4.47($\pm$0.73)	7.19($\pm$0.78)
1%-VaR	0.033($\pm$0.001)	**0.033**($\pm$0.001)	**0.033**($\pm$0.001)	0.034($\pm$0.001)	0.037($\pm$0.001)	0.042($\pm$0.001)	0.068($\pm$0.002)
OU kappa							
Variance ($\times 10^{-4}$)	2.07($\pm$0.18)	**1.84**($\pm$0.14)	**1.88**($\pm$0.15)	2.14($\pm$0.15)	2.61($\pm$0.17)	3.47($\pm$0.21)	11.76($\pm$0.59)
Skewness	$-$0.25($\pm$0.09)	$-$0.13($\pm$0.10)	0.05($\pm$0.12)	0.47($\pm$0.14)	0.99($\pm$0.19)	1.59($\pm$0.22)	2.31($\pm$0.14)
Excess kurtosis	4.58($\pm$1.78)	**4.56**($\pm$1.90)	**4.50**($\pm$1.87)	5.86($\pm$1.79)	9.01($\pm$2.33)	13.45($\pm$3.14)	13.60($\pm$2.09)
1%-VaR	0.037($\pm$0.002)	**0.034**($\pm$0.002)	**0.034**($\pm$0.002)	**0.035**($\pm$0.002)	0.038($\pm$0.002)	0.044($\pm$0.001)	0.080($\pm$0.002)

Table 2.2: Dependence as a function of DR wealth fractions of different metrics of the risky asset returns in markets with constant and Ornstein-Uhlenbeck social coupling strength κ. Each case is simulated 1000 times over 50000 time-steps. The values in brackets indicate the standard deviation over 1000 realisations. Bold values indicate values smaller than in the realisations with DR with 0% market impact.

fourth order cumulant minus 3 and is a measure of the distance from the normal distribution especially in the tails (a positive excess kurtosis is called leptokurtic and has fatter tails than a normal distribution); (iv) the value-at-risk (VaR), which is calculated in order to analyze the risk of loss in the distribution of the returns. The VaR describes the maximum possible loss, which is not exceeded with the probability $(1-\alpha)$, thus it is the $(1-\alpha)$-quantile of the distribution of returns.

Figure 2.12: Same as figure 2.11 for OU κ processes.

Here, we report the $\text{VaR}_{\alpha=1\%}$ of the daily returns. The mean values and the standard deviations (in parentheses) of these metrics are estimated over 1000 simulations for each market scenario.

The upper half of table 2.2 corresponds to the constant κ scenarios. The skew and excess kurtosis increase monotonically with the wealth fraction of DRs in the market. As the skew of the return increases, the distribution of returns becomes more right-skewed and the positive tails are getting fatter. Thus, there are more positive extremes than negative extremes. Without DR, the excess kurtosis is 0.23, which means the tails are only slightly fatter than the normal distribution. In contrast, with 50% DR, the excess kurtosis reaches 7.19. The variance and $\text{VaR}_{1\%}$ decrease slightly when a small fraction of DRs is included in the market and increase monotonically with the fraction of DRs for fractions of DR larger than 5%. Since both measures can be interpreted as quantifying the risks of the asset, a small fraction of DR slightly reduces the riskiness of the asset, while a larger fraction of DR increases the riskiness. This result is consistent with the observations in Section 2.4.1 that a small fraction of DR in the market reduces the average peak size, while a larger fraction of DR increases the average peak size.

For OU κ scenarios shown in the lower half of table 2.2, the averages of the variance, excess kurtosis, and $\text{VaR}_{1\%}$ calculated over 1000 simulations with 50000 time-steps decrease as the wealth fraction of DR in the market increases, up a wealth fraction of DRs 5% for the variance and 10% for the excess kurtosis and $\text{VaR}_{1\%}$. Beyond these wealth fraction of DR, these risk measures increase with larger fractions of DR. Thus, similarly to the simulations without bubbles for constant κ, there is an optimal fraction of DR in the market that reduces the riskiness of the asset in terms of the variance and $\text{VaR}_{1\%}$. However, when the fraction of DR in the market is larger than the optimum, the riskiness of the asset increases significantly. Since the Noise traders and DR often buy the risky

asset at a similar time, when the price increases, the presence of a large fraction of DR accelerates the price increase. Since the DR aim at selling shortly before the peak, they also trigger an earlier crash of the bubbles. Furthermore, the decreasing influence of the fundamentalist traders who have a stabilizing effect, amplifies the fluctuations in the price path. Thus, replacing too much of the fundamentalists by DR has a destabilizing impact on the market.

(Patzelt & Pawelzik, 2012) demonstrated in a simple ABM that the same mechanism that leads to market efficiency can also lead to market instabilities. Analogously, the action of the DR can have a stabilizing and a destabilizing effect on the market. A small fraction of DR in the markets and their actions to exploit the information about bubbles tend to stabilise the market. However, it becomes increasingly difficult for the DR to distinguish noise from information as the market gets closer to efficiency. This makes the DR to overreact on noise and destabilize the market with their actions. In the simulations with OU κ, there is initially more information to exploit. Therefore, a larger fraction of DR remains beneficial in stabilizing the market than in the simulations without bubbles for constant κ, where the information about bubbles that the DR receive from their LPPLS fits is mostly spurious.

2.5 Conclusion

We have built on a previous agent-based model (ABM) developed by (Kaizoji et al., 2015), which is characterised by the spontaneous formations of bubbles and drawdowns as the result of the interaction between fundamentalists and noise traders. The noise traders invest based on herding and momentum, according to a social imitation strength that can vary with time. We have extended this ABM in two ways: (i) we have proposed a more realistic dividend process driving the fundamental value of the risky asset; (ii) we have introduced a third type of agents, called "dragon riders" (DR), who are endowed with some ability in diagnosing incipient bubbles and in timing their end.

The goal of this work has been to provide a first analysis of the performance and impact of the DR. We have found that the DR strategy can successfully exploit the bubbles, which lead them to outperform the other traders on a risk-adjusted basis. The DR achieve a significantly higher average Sharpe ratio than the two other trader types in markets with bubbles created by noise traders in the presence of a social coupling strength $\kappa(t)$ following an Ornstein-Uhlenbeck process. For constant social coupling strength κ for which no clear bubble appears, the DR perform only slightly worse than the fundamentalists and much better than the noise traders. The DR can thus exploit the additional information about the emergence of bubbles and crashes using a methodology (called LPPLS calibration) designed to identify transient faster-than-exponential price growth, which has been amply documented in real markets. The DR can thus increase their wealth by predicting the time of the crash better than chance. We also find that, by exploiting the bubbles, the DR tend to destroy the characteristics of the bubbles and therefore progressively eliminate the arbitrage opportunity associated with the bubbles and crashes. In markets with a large fraction of DR, their strategy based on exploiting bubbles and crashes performs worse than the fundamentalist strategy, due to their larger number of false positive bubble detections.

In contrast, the presence of a small fraction of DR in the ABM market reduces the average bubble size, and can even in some cases completely suppress bubbles. This is due to the market impact associated with the early exit of the DR diagnosing a coming bubble burst. When the wealth fraction of DR increases too much, the average size of peaks and the fluctuations in the price path increase, as a result of false positive bubble diagnostics and the impact of their strategy on the risky asset. In turn, the noise traders react by herding more, further amplifying the price fluctuations, in a positive feedback loop increasingly destroying the bubble characteristics and thus worsening the bubble properties that the DR are skilled to detect. The presence of a large fraction of DR in the market also increases the variance of the return and the value-at-risk. Furthermore, the distribution of returns has a larger positive skew and fatter tails, the larger is the proportion of DR.

In sum, in our set-up, there is an optimal fraction of DR in the market that reduces the riskiness

of the asset, while a larger fraction of DR in the market increases the riskiness of the asset beyond the level in their absence. The optimal fraction of DR in markets without bubbles (constant κ) is approximately 5%, and it is 10% in the presence of bubbles (Ornstein-Uhlenbeck κ).

Acknowledgements: We are thankful to Ralf Kohrt for help in a first part of the research.

Appendix

2.5.1 LPPLS fits

Here, to endow the dragon riders (DR) with bubble detection skills, we follow the model of financial bubbles and crashes introduced by Johansen et al. (1999) (JLS model). This model describes the evolution of the log-price trajectory during bubbles as log-periodic power law singularities (LPPLS). The key idea is that, trend following and herding of traders lead to super-exponential growth (Sornette & Cauwels, 2015a), while social hierarchies of the traders (Zhou et al., 2005), as well as the combination of inertia, nonlinear trend following and nonlinear mean-reversal (Ide & Sornette, 2002), lead to log-periodic oscillations of the price with decreasing amplitudes. According to the JLS model, the expectation of the log-price can be written as follows (Johansen et al., 1999):

$$\mathbf{E}[\log P_t] = A + B(t_c - t)^m + C(t_c - t)^m \cdot \cos(\omega \log(t_c - t) - \phi) \tag{2.28}$$

The parameter A is the terminal log-price at t_c, B controls the amplitude of the power law acceleration, t_c is the critical time (the most probably time for the end of the bubble), m controls the convexity of the hyperbolic growth of the price, C is the amplitude of the log-periodic oscillations, ω is their angular log-frequency and ϕ is a phase constant, which embodies a characteristic time scale of reference.

The price structure described by (2.28) is a clear indication of a bubble when it can be found with sufficiently strong statistical confidence (Jiang et al., 2010; Johansen & Sornette, 2010; Sornette & Cauwels, 2015a). Therefore, the equation can be used to diagnose bubbles before the crash occurs. Thus, we endow our DR with the ability to fit (2.28) to the price path over different time-windows. The fitting procedure that our DR follow was introduced in Filimonov and Sornette (2013). The fitting corresponds to determining the vector of 7 parameters $\Phi^* = \{A, B, C, \phi, m, \omega, t_c\}$ that minimizes the sum of squared residuals between the price realization and expression (2.28).

After each fit, the obtained parameters are compared to parameter values that have been documented to be typically associated with a bubble regime, gathered from empirical investigation of previous bubbles (Zhou & Sornette, 2003; Sornette et al., 2015; Zhang, Sornette, et al., 2016). The fits that fulfil the criteria given in Table 2.3 are considered as valid fits, the other fits are discarded. The qualified fits are used by the DR to diagnose bubbles, as explained in the text.

Parameter name		**Filter Bounds**
Amplitude of the log-oscillations	$\lvert C\rvert$	< 1
Power law exponent	m	$[0.1, 0.9]$
Log-periodic Frequency	ω	$[2, 25]$
Damping	$D = \frac{m\lvert B\rvert}{\omega\lvert C\rvert}$	$\geq dampThresh1$

Table 2.3: Filtering conditions for the LPPLS parameters. $dampThresh1$ is discussed in the text.

2.6 The Equilibrium Market Price

The market clearing condition states an equilibrium between supply and demand. This means that the market price is defined following Walras' theory of general equilibrium (Walras, 1954). Each trader formulates their excess demand for the next time step and the price is calculated as the equilibrium in which supply equals demand. This is formulated as:

$$0 = \sum_i \Delta D^{DR^i}_{t-1\to t} + \Delta D^{n}_{t-1\to t} + \Delta D^{f}_{t-1\to t} \tag{2.29}$$

where $\sum_i \Delta D^{DR^i}_{t-1\to t}$ is the sum over the excess demand of the types i of the Dragon Rider (DR) strategy that are included in the market. Furthermore, the risky fraction of the DR is split up into a fraction k_t that is invested as fundamentalist (price maker) and a fraction l_t that is invested as a price taker. In the strategy used by DR in the present study, it is either $(k_t = 1, l_t = 0)$ or $(k_t = 0, l_t = 1)$. Next, the explicit formulations of the demand from (2.11), (2.22) and the demand of the DR obtained from a decision tree presented in the main text are inserted into the market clearing condition (2.24):

$$\begin{aligned}
0 =& \Delta D^{DR}_{t-1\to t} + \Delta D^{n}_{t-1\to t} + \Delta D^{f}_{t-1\to t} \\
=& \sum_{i=1}^{N} W^{DR^i}_{t} \left(x^{DR^i}_{t} \left[x^{DR^i}_{t-1} \left(R_t + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x^{DR^i}_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^{n}_{t-1} \left(x^{n}_{t} \left[x^{n}_{t-1} \left(R_t + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x^{n}_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^{f}_{t-1} \left(x_{t} \left[x^{f}_{t-1} \left(R_t + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x^{f}_{t-1} \frac{P_t}{P_{t-1}} \right) \\
=& \sum_{i=1}^{N} W^{DR^i}_{t-1} \left(\left(k^i_t \cdot \frac{E_{R_{t-1}} + \frac{d_t}{P_t}(1+r) - R_f}{\gamma \sigma^2} + l^i_t \right) \cdot \left[x^{DR^i}_{t-1} \left(R_t + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] \right. \\
&\left. - x^{DR^i}_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^{f}_{t-1} \left(\frac{E_{R_{t-1}} + \frac{d_t}{P_t}(1+r) - R_f}{\gamma \sigma^2} \left[x^{f}_{t-1} \left(R_t + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x^{f}_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^{n}_{t-1} \left(x^{n}_{t} \left[x^{n}_{t-1} \left(R_t + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x^{n}_{t-1} \frac{P_t}{P_{t-1}} \right)
\end{aligned}$$

Finally, the equation is multiplied by P_t, the return $R_t := \frac{P_t}{P_{t-1}} - 1$ is inserted and formula is organised in powers of P_t to obtain

$$
\begin{aligned}
0 =& \sum_{i=1}^{N} W_{t-1}^{DR^i} \Bigg(\left(k_t^i \cdot \frac{E_{R_{t-1}} P_t + d_t(1+r) - R_f P_t}{\gamma \sigma^2} + l_t^i P_t \right) \\
& \cdot \left[x_{t-1}^{DR^i} \left(\frac{P_t}{P_{t-1}} - 1 + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x_{t-1}^{DR^i} \frac{P_t^2}{P_{t-1}} \Bigg) \\
& + W_{t-1}^{f} \left(\frac{E_{R_{t-1}} P_t + d_t(1+r) - R_f P_t}{\gamma \sigma^2} \left[x_{t-1}^{f} \left(\frac{P_t}{P_{t-1}} - 1 + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x_{t-1}^{f} \frac{P_t^2}{P_{t-1}} \right) \\
& + W_{t-1}^{n} \left(x_t^n P_t \left[x_{t-1}^{n} \left(\frac{P_t}{P_{t-1}} - 1 + \frac{d_t}{P_{t-1}} - R_f \right) + R_f + 1 \right] - x_{t-1}^{n} \frac{P_t^2}{P_{t-1}} \right) \\
=& \frac{P_t^2}{P_{t-1}} \Bigg(\sum_{i=1}^{N} W_{t-1}^{DR^i} x_{t-1}^{DR^i} \left(k_t^i \frac{E_{R_{t-1}} - R_f}{\gamma \sigma^2} + l_t^i - 1 \right) + W_{t-1}^{f} x_{t-1}^{f} \left(\frac{E_{R_{t-1}} - R_f}{\gamma \sigma^2} - 1 \right) \\
& + W_{t-1}^{n} x_{t-1}^{n} \left(x_t^n - 1 \right) \Bigg) \\
& + P_t \Bigg(\sum_{i=1}^{N} W_{t-1}^{DR^i} \left(x_{t-1}^{DR^i} k_t^i \frac{d_t(1+r)}{\gamma \sigma^2 P_{t-1}} + \left(k_t^i \frac{E_{R_{t-1}} - R_f}{\gamma \sigma^2} + l_t^i \right) \cdot \left(x_{t-1}^{DR^i} \left(\frac{d_t}{P_{t-1}} - R_f - 1 \right) + R_f + 1 \right) \right) \\
& + W_{t-1}^{f} \left(x_{t-1}^{f} \frac{d_t(1+r)}{\gamma \sigma^2 P_{t-1}} + \frac{E_{R_{t-1}} - R_f}{\gamma \sigma^2} \cdot \left(x_{t-1}^{f} \left(\frac{d_t}{P_{t-1}} - R_f - 1 \right) + R_f + 1 \right) \right) \\
& + W_{t-1}^{n} x_t^n \left(x_{t-1}^{n} \left(\frac{d_t}{P_{t-1}} - R_f - 1 \right) + R_f + 1 \right) \Bigg) \\
& + \sum_{i=1}^{N} W_{t-1}^{DR^i} \left(k_t^i \frac{d_t(1+r)}{\gamma \sigma^2} \left(x_{t-1}^{DR^i} \left(\frac{d_t}{P_{t-1}} - R_f - 1 \right) + R_f + 1 \right) \right) \\
& + W_{t-1}^{f} \frac{d_t(1+r)}{\gamma \sigma^2} \left(x_{t-1}^{f} \left(\frac{d_t}{P_{t-1}} - R_f - 1 \right) + R_f + 1 \right)
\end{aligned}
$$

Next, some definitions are made to simplify the equation.

$$v_1 := \frac{E_{R_{t-1}} - R_f}{\gamma \sigma^2} \tag{2.30}$$

$$v_2 := \frac{d_t}{P_{t-1}} - R_f - 1 \tag{2.31}$$

$$v_3 := R_f + 1 \tag{2.32}$$

$$v_4 := \frac{d_t(1+r)}{\gamma \sigma^2} \tag{2.33}$$

The resulting equation is given in the following:

$$
\begin{aligned}
0 =& \frac{P_t^2}{P_{t-1}} \left(\sum_{i=1}^{N} W_{t-1}^{DR^i} x_{t-1}^{DR^i} \left(k_t^i v_1 + l_t^i - 1 \right) + W_{t-1}^f x_{t-1}^f \left(v_1 - 1 \right) + W_{t-1}^n x_{t-1}^n \left(x_t^n - 1 \right) \right) \\
&+ P_t \left(\sum_{i=1}^{N} W_{t-1}^{DR^i} \left(x_{t-1}^{DR^i} k_t^i \frac{v_4}{P_{t-1}} + \left(k_t^i v_1 + l_t^i \right) \cdot \left(x_{t-1}^{DR^i} v_2 + v_3 \right) \right) \right. \\
&\left. + W_{t-1}^f \left(x_{t-1}^f \frac{v_4}{P_{t-1}} + v_1 \cdot \left(x_{t-1}^f v_2 + v_3 \right) \right) + W_{t-1}^n x_t^n \left(x_{t-1}^n v_2 + v_3 \right) \right) \\
&+ \sum_{i=1}^{N} W_{t-1}^{DR^i} \left(k_t^i v_4 \left(x_{t-1}^{DR^i} v_2 + v_3 \right) \right) + W_{t-1}^f v_4 \left(x_{t-1}^f v_2 + v_3 \right)
\end{aligned}
\tag{2.34}
$$

Equation (2.34) is a quadratic equation of the form $a_t P_t^2 + b_t P_t + c_t = 0$ where

$$
\begin{aligned}
a_t =& \left(\sum_{i=1}^{N} W_{t-1}^{DR^i} x_{t-1}^{DR^i} \left(k_t^i v_1 + l_t^i - 1 \right) \right. && (2.35) \\
&\left. + W_{t-1}^f x_{t-1}^f \left(v_1 - 1 \right) + W_{t-1}^n x_{t-1}^n \left(x_t^n - 1 \right) \right) / P_{t-1} && (2.36) \\
b_t =& \sum_{i=1}^{N} W_{t-1}^{DR^i} \left(x_{t-1}^{DR^i} k_t^i \frac{v_4}{P_{t-1}} + \left(k_t^i v_1 + l_t^i \right) \cdot \left(x_{t-1}^{DR^i} v_2 + v_3 \right) \right) && (2.37) \\
&+ W_{t-1}^f \left(x_{t-1}^f \frac{v_4}{P_{t-1}} + v_1 \cdot \left(x_{t-1}^f v_2 + v_3 \right) \right) + W_{t-1}^n x_t^n \left(x_{t-1}^n v_2 + v_3 \right) && (2.38) \\
c_t =& \sum_{i=1}^{N} W_{t-1}^{DR^i} \left(k_t^i v_4 \left(x_{t-1}^{DR^i} v_2 + v_3 \right) \right) + W_{t-1}^f v_4 \left(x_{t-1}^f v_2 + v_3 \right) && (2.39)
\end{aligned}
$$

The equation has the following two solutions:

$$
P_t^{1,2} = \frac{-b_t \pm \sqrt{b_t^2 - 4 a_t c_t}}{2 a_t} \tag{2.40}
$$

The positive solution of this equation is the price in the next time-step.

2.7 Table with parameter values

Parameter name	Explanation	Value
Market		
T	Simulation length	50000
$seed$	Random seed	6 digit number
Assets		
r_f	Risk free interest rate	0.01/250=0.00004
d_0	Initial dividend	0.04/250=0.00016
r_d	Expected growth rate of the dividend	0.04/250=0.00016
σ_d	Expected standard deviation of the dividend growth rate	0.000016
P_0	Initial price of the risky asset	1
σ_r	Expected standard deviation of the risky asset price	$\sqrt{0.10/250} = 0.02$
N^r	Number of risky assets	1
Noise traders		
x_o^n	Initial fraction of the risky asset held by the noise traders	0.3
W_0^n	Initial wealth of the noise traders	10^9
c_H	Momentum weight	1
c_s	Opinion index weight	1
p_+	Switching probability when holding the risky asset	0.199375
p_-	Switching probability when not holding the risky asset	0.200625
θ	Memory parameter	0.95
H_0	Initial momentum	0.00016
N^n	Number of noise traders	1000
Fundamentalists		
x_0^f	Initial fraction of the risky asset held by the fundamentalists	0.3
W_0^f	Initial wealth of the fundamentalists	0 to 10^9
E_{R_t}	Expected return of the risky asset $= r_d$	0.00016
DR		
W_0^{DR}	Initial wealth of the DR	100 to 10^9
x_0^{DR}	Initial fraction of the risky asset held by the DR	0.3
E_{R_t}	Expected return of the risky asset $= r_d$	0.00016
LPPLS fits		
t_1	Starting time of LPPLS fits	480
t_2	End of LPPLS fits	1500 or 5000
k_{max}	Maximal size of the LPPLS fitting window	480
k_{min}	Minimal size of the LPPLS fitting window	10
Social coupling strength		
κ_0	Initial social coupling strength	$0.98 \cdot 0.199375$
μ_κ	Mean of the OU social coupling strength	$0.98 \cdot 0.199375$
η_κ	Mean reversion of the OU social coupling strength	0.11
σ_κ	Standard deviation of the OU social coupling strength	0.01

Table 2.4: Parameters characterising the model financial markets and used in the simulations. Dividends, interest rates and standard deviations are given as daily values (per time-step).

Chapter 3

How market intervention can prevent bubbles and crashes

Using a previously validated agent-based model with fundamentalists and chartists, this chapter investigates the usefulness and impact of direct market intervention. The policy maker diagnoses bubbles by forming an expectation of the future returns, then invests in burgeoning bubbles and sells countercyclically the overpriced asset to fight market exuberance. Preventing bubbles and crashes, this market intervention improves all analysed market return metrics, volatility, skewness, kurtosis and VaR, without affecting long-term growth. This increases the Sharpe ratios of noise traders and of fundamentalists by approximately 28% and 45% respectively. The results are robust even for substantially miscalibrated long-term expected returns.

Based on Westphal, R., & Sornette, D. (2020). How market intervention can prevent bubbles and crashes. *Swiss Finance Institute Research Paper*, (12-74).

3.1 Introduction

Policy makers aim for well functioning financial markets, because of the key roles financial markets have for a society. They are essential financing tools catalysing the development of enterprises and fostering innovation. Price setting by the equilibrium between supply and demand of multiple investors with varied sources of information usually ensures the "correct" valuation, allowing an efficient and rational allocation of resources to the different sectors of the economy. Because investors allocate funds to firms with promising future standing and/or growth prospects, financial markets are inherently forward looking. In other terms, they are predictors of future economic growth. Financial markets also provide storage of value. The wealth effect of investors feeling richer upon stock market appreciation is well documented to boost consumption in a virtuous circle of economic expansion.

This ideal description flies in the face of a more complex reality in which excess volatility phases (R. Shiller, 1981) are the norm more than the exception, and exuberant market regimes (R. J. Shiller, 2006) are sometimes followed by crashes (Johansen & Sornette, 2002; Sornette, 2003). These turbulences have increasingly characterised financial markets in the last three decades, with the dot-com bubble that crashed in 2000, the bubble on real-estate and financial securitisation of mortgages that crashed in 2008, the market boom fuelled by quantitative easing that ended in 2018-Q1, the short-lived market exuberance of 2019 ending with the Covid-19 triggered market crash in March 2020 (see e.g. (Sornette & Cauwels, 2015a, 2015b)).

Arguably, policy makers should have a strong interest in stabilizing financial markets by decreasing excess volatility and damping market turbulences such as bubbles and crashes, but without sacrificing economic growth. Biswas et al. (2020) even show that the burst of a bubble can re-

sult in persistent aggregate economic activity below the pre-bubble trend, thus reducing the ex-ante welfare. However, in practice, the stance of most policy makers seems to be summed up by the remarks by Chairman Alan Greenspan at a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming (August 30, 2002): "As events evolved, we recognized that, despite our suspicions, it was very difficult to definitively identify a bubble until after the fact–that is, when its bursting confirmed its existence." Greenspan further confirmed that "instead of trying to contain a putative bubble by drastic actions with largely unpredictable consequences, we chose, as we noted in our mid-1999 congressional testimony, to focus on policies to mitigate the fallout when it occurs and, hopefully, ease the transition to the next expansion" (Greenspan, 2004). Acting to prick a bubble when the bubble is on its way up is generally considered impossible or ill-advised for fear of false positives and the danger that the remedy might be worse than the disease. Apart from some exceptions, the attitude of central banks has thus been in general to act vigorously only after a crash occurred, to provide liquidity as well as cheaper access to credit in the form of lower interest rates.

On the other side of the debate, Cecchetti et al. (2000) and Cecchetti et al. (2002) develop a number of arguments for how asset price misalignments should be used to guide central bank policy. In particular, they show that interest rates should respond to stock price bubbles in order to dampen the overall volatility in economic activity. The theoretical work by B. Bernanke and Gertler (2000); B. S. Bernanke and Gertler (2001) finds that direct asset price targeting might have undesirable side-effects, but changes in the asset price can help to forecast inflationary or deflationary pressure and a flexible inflation-targeting provides macroeconomic and financial stability. In his review of six general arguments favoring monetary targeting of asset bubbles, Roubini (2006) suggested that the standard arguments against policy intervention to thwart bubbles do not hold on close scrutiny. In particular, in the case of endogenous bubbles (defined by their susceptibility to monetary policies), optimal monetary policy calls for attempting to control the bubble. Wadhwani (2008) reviews the justifications for "leaning against the wind", using a number of indicators such as loan-to-value ratios, growth of the value of bank assets and so on. F. Mishkin (2011) argues that policy intervention should monitor credit market conditions and use macroprudential measures to restrain over-exuberance in credit markets as well as financial imbalances. Ikeda (2019) develops a dynamic model with rational bubbles in which bubble-led boom reduces firms' borrowing constraints and keeps inflation from rising. Ramsey-optimal monetary policy is shown to call for tightening to curb the boom. Gali et al. (2020) study the relationship between interest rate policy and bubble dynamics in the laboratory and show in their setting that "leaning against the wind" by increasing interest rates in response to asset price increases reduces the price of the asset bubble on the short term but tends to exacerbate asset price bubbles on the longer term. They also stress that expectations are backward-looking, with adaptive and trend extrapolating elements, rather than rational. Blot et al. (2018) assess the dynamic impact of monetary policy shocks on a new bubble indicator based on a principal component analysis to extract the common pattern of structural, econometric and statistical empirical approaches. Their main result based on their chosen examples is that restrictive monetary policy ("leaning against the wind") cannot help deflating stock or housing price bubbles.

Here, we complement these burgeoning theoretical, empirical and experimental approaches by deploying a realistic agent-based model (ABM) that supports transient endogenous "super-exponential" bubbles followed by crashes, in order to test the consequences of direct policy intervention on the overall financial market and the risk-adjusted performance of the investors. ABMs belong to the broader class of computational economic models and simulate individual operations and interactions of multiple agents. ABMs can incorporate heterogeneous beliefs and have become more widely used in economics in recent years (Kirman, 2012). In principle, ABMs have the great advantage of allowing for the introduction of arbitrary levels of complexity and heterogeneity among a large population of economic agents, thus providing the means to approach a realistic description of real economic and financial systems. ABMs are well suited to represent out-of-equilibrium phenomena such as financial bubbles and are not constrained to stationary conditions.

They can describe transient dynamics, including convergence to or deviations from equilibrium (Sornette, 2014). Consequently, they can be used to test the dynamics of an inherently nonlinear impact of a policy in the presence of complex feedback loops. Kyle (1985) and Black (1986) were among the first who developed ABMs of financial markets that consider the behavior of traders who are influenced by trends and price patterns instead of only the fundamental value of an asset. Later, De Long et al. (1990, 2007) developed an ABM in which financial bubbles emerge from positive feedback caused by noise traders. Extensive reviews of ABMs of financial markets can be found in Samanidou et al. (2007), and more recently in Dieci and He (2018) and C. Hommes and LeBaron (2018).

In the present paper, we build on the ABM consisting of fundamentalists and noise traders developed by Kaizoji et al. (2015) and extended by Westphal and Sornette (2020b). It features realistically looking bubbles and crashes, while reproducing the most important stylized facts of financial markets. The set-up can incorporate the feedbacks of policy decisions on the market and on the other traders. Thus, this market model is used to analyse the impact and consequences of policy intervention on the overall market and to explore the existence of trade-off between economic growth and financial stability objectives.

Previous research (Westphal & Sornette, 2020b) investigated the impact of a "dragon-rider", i.e. a third class of investors in addition to fundamentalists and noise traders, who exploit their ability to diagnose financial bubbles from the endogenous price history to determine optimal entry and exit trading times. The name "dragon-rider" is based on the empirical observation that crashes that follow bubbles are exceptional events, outliers of strong significance (Johansen & Sornette, 2002, 2010), which are named "dragon-kings" to emphasise their special status and specific amplifying mechanisms (Sornette, 2009; Sornette & Ouillon, 2012) (for a pedagogical introduction, see `https://en.wikipedia.org/wiki/Dragon_King_Theory`). Using calibrations of the price dynamics with models of bubbles viewed as transient super-exponential episodes, the dragon-riders obtain a diagnostic of the presence or absence of a bubble and use this information to "ride" the ascending price bubble that they exit when they assess that the burst is close.

Here, we consider the different situation where yet another class of investors are introduced in addition to fundamentalists and noise traders, the "dragon-slayers". As the name implies, the dragon-slayers are taken to represent policy makers whose goal is to prevent the development of bubbles, and thus avert the ensuing crashes. While dragon-riders strive to maximize their risk-adjusted return by exploiting financial bubbles, dragon-slayers aim at reducing or even suppressing bubbles and their subsequent damaging crashes. Their metric of success is not necessarily a risk-adjusted return but how much they calm the markets and prevent the big swings of bubbles and crashes. The analysis followed by our dragon-slayers to decide how to intervene can be decomposed into three components: an estimation of the long-term growth rate of the asset, a diagnostic of a growing bubble, and the anticipation of the crash. The later two are quantified by a crash probability that is estimated via a logistic function whose argument is a measure of the realised excess return over the expected return of the asset. Equipped with their diagnostic, the dragon-slayers invest to increase their allocation to the stock market when they detect a small incipient positive bubble so as to pile up assets that they are then able to sell countercyclically when they diagnose too much exuberance. They follow the opposite strategy for negative bubbles (see e.g. (Sornette & Cauwels, 2015a) for a definition). We investigate the intended and possible unintended consequences of these interventions and quantify the trade-off between financial stability and economic growth.

We find that the policy maker succeeds in preventing bubbles and crashes in our ABM. In simulations without bubbles, the policy maker behaves similarly to the fundamentalists and his impact is negligible, following the principle of "Primum non nocere". In simulations where bubbles form spontaneously as a result of the noise traders's strategies, the policy maker's intervention reduces the average drawdown by a factor of two when his market impact becomes significant. We find that the policy maker intervention improves all analysed metrics of market returns, including volatility, skewness, kurtosis and VaR, making the market less turbulent and more stable. The combination of

fewer bubbles and crashes, lower market risks and the stability of the long-term growth rate make the policy maker intervention to improve the performance of all investors as measured by their risk-adjusted return, increasing the Sharpe ratios from approximately 0.3 to 0.5 for noise traders, from 0.6 to 0.8 for fundamentalists as the market impact of the policy maker increases to the level of the fundamentalists. We also test the sensitivity of these results to variations of the key parameters of the strategy of the policy maker and find very robust outcomes. In particular, the conclusions are unchanged even under very large miscalibrated long-term expected returns of the risky asset.

In the financial literature, the dragon-slayers' interventions are called open market interventions and have become one of the favoured tools of central banks, completing more standard monetary policies via interest rate guidance. Open market operations usually involve buying or selling government bonds in the open markets. Market intervention in the form of stock purchases is a more recent phenomenon, which seems to develop in importance as the severity of economic and market stresses has been mounting. Stock purchases have been documented as a policy measure to end a crash and help the economy to recover, for example in Hong Kong in 1998 and in China in 2015 (Su et al., 2002; Huang et al., 2019). In Hong Kong, the government purchased stocks that are constituants of the Hang Seng Index to restore investors' confidence. The total investment of HK $118 billion (or US $15 billion) in the 33 constituent stocks of the Hang Seng made the government one of the largest shareholders, owning between 2.49% and 12.28% of the outstanding shares of the individual companies. The purchases reversed the trend of declining stock prices, and the higher stock prices persisted even after the intervention period ended (Su et al., 2002). After the crash in the Chinese stock market in mid-June 2015, the government directly and indirectly purchased stocks of more than 1000 firms between July and September. These purchases increased stock demand, reduced default probabilities, and increased liquidity. This increased the value of a subsample of the firms by about RMB 206 billion (Huang et al., 2019). Between 2008 and 2011, the Swiss franc appreciated from 1.6 CHF/EUR to less than 1.1 CHF/EUR. As a response to this massive "over-valuation" and in order to maintain price stability, the Swiss National Bank (SNB) has been intervening in the foreign exchange market, building reserve assets becoming larger than Switzerland's GDP in 2016. As a consequence of its interventions, the SNB became a large public investor globally: in its latest 13F SEC filing for Q2 2020, the Swiss National Bank has disclosed 2,437 total holdings and a US stock portfolio valuation of about 118 billion USD. Following the burst of Japan's asset price bubble in 1990, the Bank of Japan has been taking an ever expanding role in its fight against deflation. After having exhausted the orthodox policy arsenal, the Bank of Japan started buying equity ETFs in 2010. By April 2019, it has become a top ten shareholder of more than 50% of publicly traded companies. As a reaction of the March 2020 market crash, on March 16, Governor Haruhiko Kuroda announced that the central bank would double the pace of its equity purchases to $113 billion per year.

The paper is structured as follows. Section 2 presents the ABM consisting of fundamentalists and noise traders, which is used to test the policy interventions, and describes how the strategy of the dragon-slayer (policy maker) is constructed. Section 3 describes the influence of the dragon-slayer on the price dynamics and analyses the consequences on price peaks and drawdowns. In section 4, the impact of the dragon-slayer's intervention on the risk-adjusted return of the traders is analysed together with the quantification of the fraction of wealth accumulated by the dragon-slayer. The observed improvement on market risk properties resulting from the dragon-slayer's intervention is appraised with respect to parameter miscalibration risks. This is done by varying the parameters controlling the strategy of the dragon-slayer. Section 4 concludes.

3.2 The market model

The market in the ABM consists of two types of assets: a risky asset, which is a dividend paying stock, and a risk-free asset, which pays a constant return in each time-step and represents a risk-free government bond or a bank account. These assets are traded by three types of investors, funda-

mentalists (f), noise traders (n), and a dragon-slayer (d). Fundamentalists are rational risk-averse investors who invest by maximizing their expected utility under a constant relative risk aversion (CRRA) utility function in each time-step. The noise traders invest based on the price momentum and social imitation. The dragon-slayer trades with the objective to prevent bubbles and crashes in the price of the risky asset. The set-up of the model follows the ABM developed by Kaizoji et al. (2015) and modified by Westphal and Sornette (2020b) to introduce dragon-riders as mentioned in the introduction. The present version of the model replaces the dragon-riders of Westphal and Sornette (2020b) by dragon-slayers.

3.2.1 Market set-up

The traders decide in each time-step how to allocate their wealth between the two assets. The risk-free asset has perfect elastic supply and pays a constant return r_f. In contrast, the price P_t of the risky asset is defined endogenously by demand and supply. The asset pays a dividend d_t in each time-step. The dividend process is a discrete stochastic growth process following Westphal and Sornette (2020b). It is defined as

$$d_t = d_{t-1}\left(1 + r_t^d\right) \ , \tag{3.1}$$

where the growth rate r_t^d is a Gaussian process with mean value $r_d > 0$ and variance σ_d^2,

$$r_t^d = r_d + \sigma_d u_t, \tag{3.2}$$

with $u_t \overset{iid}{\sim} \mathcal{N}(0,1)$.

The difference between the return of the risky asset and the risk-free asset is called excess return. It is defined as the sum of the capital return $r_t = \frac{P_t}{P_{t-1}} - 1$, where P_t is the price of the risky asset at time t, and of the return from the dividend d_t of the risky asset minus the risk-free rate r_f:

$$r_{excess,t} = r_t + \frac{d_{t-1} \cdot (1 + r_t^d)}{P_{t-1}} - r_f. \tag{3.3}$$

3.2.2 Fundamentalist strategy

At each time-step, fundamentalists invest a fraction x_t^f of their wealth in the risky asset and the remaining fraction into the risk-free asset, such that they maximize their expected utility with a constant relative risk aversion (CRRA) utility function over one-period. In other words, they are myopic investors who update at each time step their investment decision based on the new information on the risky asset price and its dividend that is obtained at the end of each period. The one-period optimisation is chosen because it quite realistically captures the bounded rationality of real human investors and it allows us to keep the mathematical formulation simple. In that, we follow the tradition of previous investigations, such as Chiarella et al. (2006), C. H. Hommes and Wagener (2009), Kaizoji et al. (2015), and Westphal and Sornette (2020b). Moreover, in the absence of transaction cost (and other limitations on trading), a greedy strategy that only considers one period at a time is optimal, since performance for the current period does not depend on previous holdings (Boyd et al., 2016). Since most of the efforts in developing a good trading algorithm goes into forming good forecasts of the expected return (J. Y. Campbell et al., 1998; Grinold & Kahn, 2000), and given the large noise and difficulties inherent in the one-period return prediction, a multi-period approach to forecast returns would be difficult to justify at the concrete operational level.

Each fundamentalist is equipped with the same information and utility function. Therefore, each fundamentalist decides on the same optimal allocation of their wealth. Thus, the fundamentalists' investment can be considered at the aggregate level as the optimization problem of one representative agent who invests a fraction x_t^f of the cumulative wealth W_t^f in the risky asset.

The CRRA utility function with risk aversion parameter γ is defined as (Ljungqvist & Sargent, 2018):

$$U(W) = \begin{cases} \log(W) & \text{for } \gamma = 1 \\ \frac{W^{1-\gamma}}{1-\gamma} & \text{for } \gamma \neq 1 \end{cases} \tag{3.4}$$

In each time-step, the fundamentalists solve the maximization problem

$$\max_{x_t^f} \mathbf{E}_{t-1}\left[U(W_t^f)\right]. \tag{3.5}$$

The wealth W_t^f evolves as a function of x_t^f. This evolution consists of the wealth $W_{t-1}^f \cdot x_{t-1}^f$ invested in the risky asset, which enjoys the return r_t on the risky asset and the dividend payment d_t per share, plus the wealth $W_{t-1}^f \cdot (1 - x_{t-1}^f)$ invested in the risk-free asset paying an interest at the risk-free interest rate r_f:

$$\begin{aligned} W_t^f &= W_{t-1}^f \left(1 + x_{t-1}^f \cdot \left(r_t + \frac{d_t}{P_{t-1}}\right) + (1 - x_{t-1}^f) \cdot r_f\right) \\ &= W_{t-1}^f \cdot \left(1 + x_{t-1}^f \cdot r_{excess,t} + r_f\right), \end{aligned} \tag{3.6}$$

where $r_{excess,t}$ is defined by expression (3.3). In first order approximation and assuming $d_t \ll P_t$ (Kaizoji et al., 2015), the resulting fraction of wealth invested in the risky asset with CRRA coefficient γ is given by

$$x_{t-1}^f = \frac{1}{\gamma}\frac{E_{t-1}[r_{excess,t}]}{Var_{t-1}[r_{excess,t}]} = \frac{E_{r_t}^f + \frac{d_{t-1}}{P_{t-1}}(1 + r_d) - r_f}{\gamma(\sigma^2 + \frac{d_{t-1}^2 \cdot \sigma_r^2}{P_{t-1}^2})} \approx \frac{E_{r_t}^f + \frac{d_{t-1}}{P_{t-1}}(1 + r_d) - r_f}{\gamma\sigma^2} \tag{3.7}$$

where $E_{r_t}^f$ is the fundamentalists' expectation of the return of the risky asset and σ^2 is its expected variance. Thus, the first order approximation reduces the CRRA optimisation to a mean-variance optimisation. Denoting the number of shares invested in the risky asset by $n_t^f := \frac{x_t^f W_t^f}{P_t}$, the excess demand of the fundamentalists for the risky asset is described by the following equation:

$$\begin{aligned} \Delta D_{t-1\to t} :=& n_t^f P_t - n_{t-1}^f P_t \\ =& x_t^f W_t^f - x_{t-1}^f W_{t-1}^f \frac{P_t}{P_{t-1}} \\ =& W_{t-1}^f \left(x_t^f \left[x_{t-1}^f \left(r_t + \frac{d_t}{P_{t-1}} - r_f\right) + r_f + 1\right] - x_{t-1}^f \frac{P_t}{P_{t-1}}\right) \end{aligned} \tag{3.8}$$

together with (3.7).

3.2.3 Noise trader strategy

The noise traders' investment strategy is based on the analysis of the assets historical price returns and on the opinion of other noise traders. Analogous to Kaizoji et al. (2015) and Westphal and Sornette (2020b), each individual noise trader is either invested in the risky asset or in the risk-free asset and does not diversify their portfolio. In each time-step, the noise traders decide independently and probabilistically to keep their current position or switch their strategy to the other asset. The switching probability is influenced by the opinion of the other noise traders and the price momentum. The price momentum H_t is defined as the exponential moving average of the return of the risky asset

$$H_t = \theta H_{t-1} + (1 - \theta) r_t, \tag{3.9}$$

where $0 \leq \theta \leq 1$ controls the time-span $\sim 1/(1-\theta)$ of the noise traders' memory.

The number of noise traders invested in the risky asset is denoted as N_t^+, the number of noise traders invested in the risk-free asset is N_t^-, and the total number of noise traders is $N^n = N_t^+ + N_t^-$. This total number of noise traders is kept constant during the simulations. The opinion index s_t describes the collective opinion towards the risky asset compared to the risk-free asset in each time-step. It is defined as

$$s_t = \frac{N_t^+ - N_t^-}{N^n} \quad \in [-1, 1]. \tag{3.10}$$

The probability at time-step t for a noise trader who is invested in the risky asset to switch his investment position to the risk-free asset is denoted as p_t^+. Respectively, the switching probability of a trader, who is invested in the risk-free asset, is denoted as p_t^-. These probabilities are given by

$$p_t^{\pm} = \frac{p}{2}\left(1 \mp \kappa_t(s_t + H_t)\right), \tag{3.11}$$

where the constant p control the average holding time of each asset type.

The time-dependent parameter κ_t determines the strength of social imitation and momentum following, which are assumed here to be controlled by the same parameter. Financial markets are characterised by periods of exuberance alternating with periods of pessimistic mood (Sornette, 2003; R. J. Shiller, 2006). We propose to account for the existence of different regimes and for the random switches between them by allowing the noise traders to shift between periods when they have a large tendency to herd and when they have more heterogenous opinions. Thus, their susceptibility to herding is regime dependent. This incorporates the influence of exogenous factors such as economic and geopolitical regimes in the model. We account for these characteristics by allowing the coupling strength κ_t to be time-dependent according to a discretized Ornstein-Uhlenbeck process

$$\kappa_t = \kappa_{t-1} + \eta_\kappa(\mu_\kappa - \kappa_{t-1}) + \sigma_\kappa v_t, \tag{3.12}$$

where η_κ is the strength of mean reversion that controls the persistence time $\sim 1/\eta_\kappa$ of deviations from the mean μ_κ, driven by fluctuations with standard deviation σ_κ with $v_t \overset{iid}{\sim} \mathcal{N}(0,1)$. The expected value of the Ornstein-Uhlenbeck process κ_t starting at an initial value κ_0 is

$$E[\kappa_t] = \kappa_0 \cdot \exp(-\eta_\kappa \cdot t) + \mu_\kappa(1 - \exp(-\eta_\kappa \cdot t)). \tag{3.13}$$

Thus, the estimated time to revert from a value $\kappa_0 > \kappa_c$ to a value $\kappa_c > \mu_\kappa$ is

$$\Delta T = \frac{1}{\eta} \log\left(\frac{\kappa_0 - \mu_\kappa}{\kappa_c - \mu_\kappa}\right). \tag{3.14}$$

Aggregating the independent investment decisions over all noise traders amounts to considering an equivalent representative noise trader who decides on the fraction x_t^n of his wealth invested in the risky asset. The risky fraction is then given by

$$x_t^n = \frac{N_t^+}{N^n} \quad \in [0, 1]. \tag{3.15}$$

The fraction of wealth invested in the risky asset then evolves as

$$x_t^n = \frac{1}{N^n}\left(\sum_{k=1}^{N_{t-1}^+}\left(1 - \xi_k(p_{t-1}^+)\right) + \sum_{l=1}^{N_{t-1}^-}\xi_l(p_{t-1}^-)\right), \tag{3.16}$$

where $\xi(p)$ are Bernoulli random numbers. The drawing of a Bernoulli random number corresponds to the decision of each individual noise trader to switch to the other asset or to stay invested as in the previous time-step.

The corresponding aggregated wealth equation has the same structure as (3.6):

$$W_t^n = W_{t-1}^n \cdot \left(1 + x_{t-1}^n \cdot \left(r_t + \frac{d_t}{P_{t-1}}\right) + (1 - x_{t-1}^n) \cdot r_f\right). \tag{3.17}$$

Combining the previous equations, the resulting aggregated excess demand from the noise traders for the risky asset is described by the following equation:

$$\Delta D_{t-1\to t}^n = W_{t-1}^n \left(x_t^n \left[x_{t-1}^n \left(r_t + \frac{d_t}{P_{t-1}} - r_f\right) + r_f + 1\right] - x_{t-1}^n \frac{P_t}{P_{t-1}}\right) \tag{3.18}$$

together with equations (3.16) and (3.11).

3.2.4 Dragon-slayer strategy

The dragon-slayer's objective is to prevent bubbles and crashes by predicting them and trading the risky asset. He builds an expectation of the future return based on his diagnostic of the bubble, his anticipation of the drawdown that is expected to result from the burst of the bubble and his expectation of the long-term growth rate. Using this expected return, he maximizes his expected utility similarly to the fundamentalists. First, he calculates the excess return momentum y_t above the long-term return $\bar{r}$ in units of daily return as an exponential moving average with memory parameter a:

$$y_t = a \cdot y_{t-1} + (1 - a) \cdot (r_{t-1} - \bar{r}). \tag{3.19}$$

Note that y_t is knowable by the dragon-slayer at time $t-1$ when r_{t-1} is observed. The index t of y_t is used to indicate that the excess return momentum y_t is used by the dragon-slayer to decide on his allocation on the risky asset over the period from $t-1$ to t.

Using y_t, the dragon-slayer estimates the probability that the overpricing will result in a crash according to the logistic function:

$$\lambda_t = 1/(1 + e^{-(|y_t| - l_y)/s}) \tag{3.20}$$

The threshold l_y describes the level of overpricing that the dragon-slayer defines as excessive and s quantifies the confidence of the dragon-slayer in the existence of overpricing. The dragon-slayer invests in the risky asset when he detects a small deviation of the return from the long-term growth rate in order to construct an inventory that he will be able to draw from later to fight a possible future market exuberance. Then, when this deviation between the current growth rate and the long-term growth rate exceeds the dragon-slayer's tolerance level, he starts to sell the risky asset, that he has accumulated earlier, to fight against future price increase.

The dragon-slayer's main objective is to prevent bubbles and crashes. However, as a policy maker, he is also obligated not to squander the wealth (or money creation power) he is entrusted with and to invest it in a way that ensures a reasonable amount of risk-adjusted return. Therefore, the dragon-slayer's strategy is embedded in the same framework as the fundamentalist strategy, maximizing the expected utility with the same risk-aversion parameter γ. The only difference between pure fundamentalists and dragon-slayers lies in the way they form their expectation of the future return. The inclusion of bubble and crash forecasts in the expected return of the risky asset by the dragon-slayer implicitly results in the desired counter-cyclical investment strategy, as we shall see. Moreover, this formulation ensures that the strategy of the dragon-slayer converges towards that of the fundamentalists in the absence of bubbles.

At time $t-1$, the dragon-slayer forms an expectation of the risky asset return according to the following expression

$$E_{r_t}^d = \overbrace{(1 - \lambda_t) \cdot w_y \cdot y_t}^{(I)} - \overbrace{\lambda_t \cdot \mathrm{sign}(y_t) \cdot w_y \cdot l_y}^{(II)} + \overbrace{\bar{r}}^{(III)} \tag{3.21}$$

with w_y being the weight that the dragon-slayer gives to his bubble diagnostic. This expected return $E^d_{r_t}$ is time-dependent and is performed over a time horizon proportional to $1/(1-a)$ as seen from expression (3.19).

Term (I) in (3.21) corresponds to the diagnostic by the dragon-slayer of the degree with which the risky asset can deliver a return above the long term value. This occurs when the excess return momentum y_t becomes positive but not too large so that the crash probability λ_t remains small. This corresponds to $-l_y \leq y_t \leq l_y$.

Term (II) captures the impact on the expected return $E^d_{r_t}$ stemming from the anticipation of the drawdown. This part is proportional to the threshold l_y that is tolerated in the price momentum, which is also the expected amplitude of the drawdown. Furthermore, $\mathrm{sign}(y_t)$ incorporates the direction of the bubble. Term (II) is significant when the estimated probability λ_t that the bubble is going to crash is close to 1, which corresponds to $|y_t| > l_y$.

Term (III) is the long-term growth rate of the risky asset, which is equal to the fundamentalists' expectation of the future return and is equal to the dividend growth rate r_d (Westphal & Sornette, 2020b). This term is not influenced by the bubble diagnostic of the dragon-slayer. When the excess return momentum $y_t \to 0$, $E^d_{r_t} \longrightarrow \bar{r}$, which means that the investment allocation of the dragon-slayer converges towards that of the fundamentalists.

After building his expectation of the return of the risky asset according to (3.21) at time $t-1$, the dragon-slayer chooses to allocate the fraction x^d_{t-1} of his wealth to the risky asset, which is given by expression (3.22). This fraction x^d_{t-1} holds from $t-1$ to t, at which time the dragon-slayer observes r_t and recalculates the new value of the excess return momentum y_{t+1} cascading into a new value $E^d_{r_{t+1}}$ and thus of his allocation x^d_t, and so on. To determine x^d_{t-1}, the dragon-slayer uses the same maximisation process as the fundamentalist, using a CRRA utility with a risk aversion level γ. The difference with eq. (3.7) is the use of $E^d_{r_t}$ rather than $\bar{r} = r_d$ for the one-period expected return of the risky asset. This yields

$$x^d_{t-1} = \frac{E^d_{r_t} + \frac{d_{t-1}}{P_{t-1}}(1+r_d) - r_f}{\gamma\sigma^2} \tag{3.22}$$

The dragon-slayer decides on a level l_y of mispricing that seems unreasonable (or unsustainable) to him. This means that, whenever this level is exceeded, the dragon-slayer thinks the asset is in a bubble and is going to crash soon. For any $0 < y_t < l_y$, the dragon-slayer expects the asset to grow further and his estimation of the expected return of the risky asset for the near future is larger than r_d. However, when $y_t > l_y$, he expects a crash, which leads to his reduced expected return below r_d as a result of the drawdown anticipated to burst the bubble. For a negative bubble corresponding to an underpricing of the risky asset, the above reasoning applies ceteris paribus by changing y_t into $|y_t|$.

Figure 3.1 shows the risky fraction of the dragon-slayer given by (3.22) with (3.21) as a function of the excess return momentum y_t. If the dragon-slayer does not detect any overpricing (or underpricing) ($y_t = 0$), the expected return is equal to that of the fundamentalists. The future expected return increases with y_t as the dragon-slayer expects the bubble to grow further, until y_t becomes too large and he expects a crash.

Analogously to the fundamentalists, the dragon-slayer's excess demand for the risky asset is

$$\Delta D_{t-1\to t} = W^d_{t-1}\left(x^d_t\left[x^d_{t-1}\left(r_t + \frac{d_t}{P_{t-1}} - r_f\right) + r_f + 1\right] - x^f_{t-1}\frac{P_t}{P_{t-1}}\right) \tag{3.23}$$

with eq. (3.22).

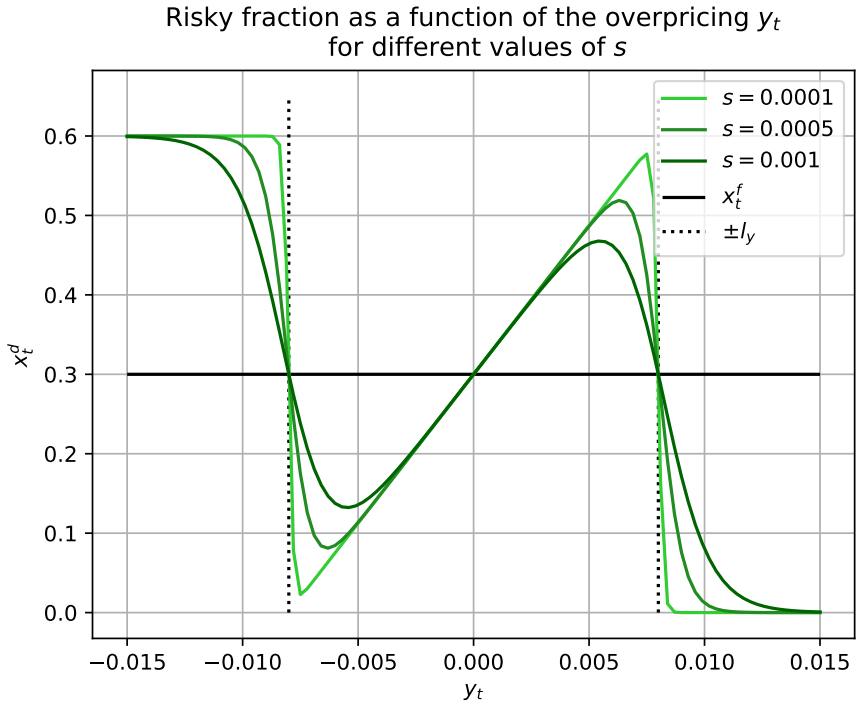

Figure 3.1: The dragon-slayer's risky fraction (3.22) as a function of the excess return momentum y_t. The parameters are $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$. Three values of s are shown, where s defined in expression (3.20) quantifies the confidence of the dragon-slayer in the existence of overpricing. Positive (resp. negative) values of y_t correspond to positive (resp. negative) bubble regimes, namely overpricing (resp. underpricing) of the risky asset.

3.2.5 Market clearing and price equation

As presented for the fundamentalists, noise traders and dragon-slayer above, each trader decides on his excess demand for the next time-step according to eq. (3.8), (3.18), and (3.23) respectively. The price is obtained from the market clearing condition, which balances demand and supply according to Walras' theory of general equilibrium (Walras, 1954):

$$0 = \Delta D^d_{t-1\to t} + \Delta D^n_{t-1\to t} + \Delta D^f_{t-1\to t} \,. \tag{3.24}$$

This yields a quadratic equation of the price at the next time-step, which has a unique positive solution giving the price P_t. The full equations and resulting price equation can be found in appendix 3.6.1.

3.3 Impact of the dragon-slayer on the price dynamics

3.3.1 General conditions of the simulations

The market parameters that are used in all simulations are listed in table 3.2. They are chosen such that each time-step corresponds to approximately one trading day.

We investigate two classes of markets:

- markets with bubbles obtained with time varying social imitation strength κ following an Ornstein-Uhlenbeck (OU) process (3.12) shown in figure 3.2;
- markets without bubbles obtained for constant κ shown in figure 3.3.

These figures plot the time evolution of realisations of the price path P_t, the risky fractions x_t^i invested by the three trader types, the excess return momentum y_t, and the wealth of the three trader types. The price trajectory P_t exhibits bubbles and crashes in the simulation with OU κ. The noise traders' risky fraction and wealth increases during the bubble regimes, but crashes together with the price. The fundamentalists' risky fraction fluctuates less than the other traders' risky fraction and they decrease their exposure to the risky asset during the bubble regime by investing proportionally to the dividend-price ratio. Therefore, their wealth is smaller than the noise traders' wealth during bubbles, but exceeds their wealth in the long-term. In the absence of bubbles, the dragon-slayer's risky fraction is the same as that of the fundamentalists. However, during bubbles, the excess return y_t increases and the dragon-slayer strategy deviates from the fundamentalist strategy. His risky fraction increases until y_t exceeds the threshold l_y, which triggers the dragon-slayer to sell the risky asset. In the simulations with constant κ, which represents markets without bubbles, the dragon-slayer strategy fluctuates around the fundamentalists' strategy.

The Ising-like structure of the noise traders' decision making (Kaizoji et al., 2015; Westphal & Sornette, 2020b) allows for in a phase transition between a disordered and an ordered regime. In the disordered regime, the noise traders' opinions are heterogeneous, in the ordered regime the noise traders polarize, which leads to an increased demand for one of the two assets. This is reflected in the price time-series as a positive or negative bubble. In the simulations with constant κ, we choose its value in the subcritical regime, at 0.98% of the critical value κ_c. In simulations with OU κ, the parameter has the same mean value $0.98\kappa_c$, but fluctuates around it according to a mean reverting OU process. Thus, there are transient regimes in which κ is larger than the critical value. This describes regimes where the noise traders tend to polarize their decisions, as a result of the spontaneous collective organisation of individuals who interact repeatedly and sufficiently strongly.

3.3.2 Dragon-slayers with negligible market impact

In our ex-post analysis, we are interested in characterising how well does the dragon-slayer diagnose bubbles and predict crashes. As a preliminary analysis, we need to identify the price peaks, which can be considered to be the target proxies of the dragon-slayers. We thus define that a price peak occurs at time scale k at time-step t_i if

$$P_{t_i} \geq P_{t_j} \quad \forall t_j \in [t_i - k, t_i + k], \tag{3.25}$$

where P_t is the price at time t and k is the minimum distance between two peaks. Thus, a peak occurs at a given time if the price at this time is larger than the price at the k previous and consecutive times. In this analysis, the minimum distance between two peaks is chosen to be $k = 250$ trading days, which corresponds to approximately one calendar year.

Figure 3.4 shows how the dragon-slayer's bubble diagnostics correlate with the price trajectory, its peaks and subsequent drawdowns. Figure 3.4 shows a simulated price path and its major peaks and the crash probability $\lambda(y_t)$ estimated by the dragon-slayer according to expression (3.20). The black dotted lines and black triangles characterise the occurence times and price heights of the peaks identified ex post for comparison.

In this example, among the 7 peaks diagnosed according to (3.25), 5 are correctly predicted (true positives) by the condition that λ_t exhibits a well-defined peak, while only 3 are correctly predicted if the condition is more stringent, for instance that λ_t should be larger than 0.2. Two peaks, at $t \approx 1000$ and $t \approx 4000$, are not detected by the dragon-slayers (false negatives), because they occur rather close to previous peaks and are relatively smaller than their close predecessors.

It is possible optimise the prediction performance in terms of sensitivity and specificity, for instance, by varying the parameters a, l_y, s, w_y involved in the definition of the dragon-slayer strategy. We refrain from such optimisation in order to focus on the robustness of our conclusions. We examine below how the properties of the price dynamics of the risky asset and the wealth dynamics of the three trader types change upon varying the model and strategy parameters.

3.3.3 Dragon-slayers with significant market impact

We analyse the impact of the dragon-slayer on the price time-series of the risky asset by increasing his initial wealth from 0% to 50% of the total initial wealth of the three trader types. The ABM is simulated with 15 different initial fractions of the total wealth allocated to the dragon-slayer (0%, 1%, 2%, 5%, 10%, 15%, 20%, 25%, 30%, 35%, 40%, 45%, 48%, 49%, 50%). The initial wealth of the noise traders is kept constant, the initial wealth of the fundamentalists decreases by the same amount that the initial wealth of the dragon-slayer increases. This ensures that the impact on the price is due to the increase of the dragon-slayer's wealth and not due to the decrease of the noise traders presence in the market. Each scenario is simulated for the same 1000 random seeds and the same set of parameters. If not denoted otherwise, the dragon-slayer parameters are $a = 0.98$, $l_y = 0.008$, $\bar{r} = 0.00016$ and $s = 0.0005$. The total simulation duration is $T = 25000$.

First, the impact of the dragon-slayer's wealth is shown qualitatively by comparing the price paths with different fractions of the dragon-slayer's wealth in Figure 3.5. Then, Figure 3.6 shows the impact of the dragon-slayer's wealth on the number and amplitudes of the peaks and drawdowns in the price of the risky asset, Table 3.1 quantifies the impact of the dragon-slayer's wealth on the long-term growth of the market, while Figure 3.7 illustrates the impact of the dragon-slayer's wealth on the return distribution of the risky asset.

Figure 3.5 shows sections of a price path and the corresponding fractions invested by the three trader types for OU κ (3.12), the corresponding figure with the same random seed is shown for constant κ in Appendix 3.10. In this example, the price of the risky asset exhibits a large bubble around $t = 1250$ in the simulation where the dragon-slayer impact is negligible. This peak disappears in the simulations performed under the same conditions, except for the initial wealth fraction of the dragon-slayer being 20% or larger. In the cases with 10% and 30%, the main bubble is slightly decreased in amplitude while some secondary peaks appear in the case of 30%. In the 40% and 50% cases, the price trajectory becomes very similar to those obtained with constant κ. In the price path with 20%, all bubbles are eliminated. This shows that the effect of the dragon-slayer impact is not deterministically monotonous as a function of the dragon-slayer wealth fraction and needs to be defined probabilistically as it depends on the specific random realisation of the prices process created by the noise traders. We thus need to perform detailed statistical analysis over many realisations to obtain meaningful conclusions. This analysis is developed below.

In the right panels in figure 3.5, the risky fraction of the dragon-slayer is more volatile for small initial wealth fractions. However, as his initial wealth increases, his strategy becomes more similar to the fundamentalist strategy and the risky fraction fluctuates closely around the fundamentalists risky fraction. By investing proportionally to the dividend-price ratio, the fundamentalist strategy has a stabilizing effect on the market. Despite the absence of bubbles, the noise trader risky fraction fluctuates a lot. This is different from the simulations with constant κ, where the noise trader risky fraction remains between 0.3 and 0.7 most of the time. The fundamentalists' risky fraction is not affected significantly by the change of initial wealth of the dragon-slayer. In the simulations with constant κ, the impact of increasing the initial wealth of the dragon-slayer is much smaller. The price path and risky fractions seem similar to the reference simulation without the dragon-slayer in all six scenarios.

With the ex post identification of peaks given in eq. (3.25), figure 3.6 presents a quantitative description of the dragon-slayer's impact on bubbles and crashes in terms of two metrics: the number of price peaks and the average amplitude of the peak-to-valley drawdown (referred to as the average peak size). The valley following a peak occurs at the time when the price takes its minimum value between two consecutive peaks. The corresponding size of a drawdown d_{t_i} is defined as the difference between the log-price at the time of the peak (t_i) and the log-price at the time of the consecutive valley

$$d_{t_i} = \log(P_{t_{peak}=t_i}) - \log(P_{t_{valley}}). \tag{3.26}$$

Figure 3.6 shows the average number of peaks calculated over 1000 simulations with 1000 different

random number seeds as a function of the initial dragon-slayer wealth fraction. The number of peaks is calculated over 12500 time-steps, which corresponds to approximately 50 years, and for a minimum distance $k = 250$ between price peaks (see definition (3.25)). The figure shows that for, OU κ, the average number of peaks does decrease by 5.7% when increasing the dragon-slayer's wealth fraction from 0% to 50%. In contrast, for constant κ, the average number of peaks remains approximately constant. This result is not necessarily bad news as it is not surprising that peaks subsist in random price trajectories and we should not expect a strong effect for this metric. In contrast, the average peak size is significantly reduced by the dragon-slayer, almost by a factor 2. For OU κ, it decreases from 57% in absence of dragon-slayer to 33% when the dragon-slayer completely replaces all fundamentalists. Thus, we can conclude that the dragon-slayer has a very strong impact in essentially eliminating all bubbles in the dynamics of the risky asset. Figures 3.11 and 3.12 show that the dragon-slayer with a shorter memory length ($a = 0.95$) as defined in expression (3.19) prevents even more bubbles and reduces further the size of the drawdowns more efficiently than the dragon-slayer with slower reactions characterised by $a = 0.99$.

dragon-slayer fraction	**Growth-rate OU κ**	**Growth-rate Constant κ**
0%	$3.87\% \pm 0.45\%$	$3.89\% \pm 0.26\%$
5%	$3.95\% \pm 0.38\%$	$3.89\% \pm 0.26\%$
10%	$3.94\% \pm 0.35\%$	$3.89\% \pm 0.26\%$
20%	$3.93\% \pm 0.34\%$	$3.91\% \pm 0.26\%$
30%	$3.97\% \pm 0.31\%$	$3.91\% \pm 0.26\%$
40%	$3.98\% \pm 0.30\%$	$3.91\% \pm 0.26\%$
50%	$3.98\% \pm 0.29\%$	$3.90\% \pm 0.26\%$

Table 3.1: Annualized growth rate of the price of the risky asset estimated from $t = 5000$ to $t = 25000$ and averaged over 1000 simulations with different random seed for each initial wealth fraction of the dragon-slayer. The dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.004$, $\bar{r} = 0.00016$, and $s = 0.0005$. The theoretical value of the growth rate of the price of the risky asset is 4%.

While the dragon-slayer prevents bubbles and crashes, one could worry that this would come at the cost of impacting the long-term average return of the risky asset. By presenting the average return and its standard deviation calculated over 1000 simulations with different random seeds, Table 3.1 shows that this is not the case. Each value is the annualized growth-rate r_a between $t = 5000$ and $t = 25000$ calculated as

$$r_a = \frac{250}{25000 - 5000} \log\left(\frac{P_{25000}}{P_{5000}}\right). \tag{3.27}$$

The first 5000 time-steps are removed to avoid the influence of transients at the beginning of the simulations. The simulations have been performed with $r_d = 0.00016$ per time step (day), which corresponds to an annualised growth rate of dividends of 4%. Thus, theoretically, the long-term growth rate of the risky asset should also be equal to 4% (Westphal & Sornette, 2020b). For all scenarios with different dragon-slayer fractions, we find that the empirical mean value of the return of the risky asset is less than one standard deviation away from the theoretical value. Thus, one can conclude that the dragon-slayer does not change the long-term growth rate of the risky asset.

Figure 3.7 shows with box plots how the distribution of returns of the risky asset is influenced by the dragon-slayer fraction for OU κ. Figure 3.7 shows the median, quartiles and range of the variance, skew, excess kurtosis and $\text{VaR}_{1\%}$ calculated over 1000 simulations with different random seeds. As observed in the analysis above, the impact on simulations with constant κ is very small, the corresponding figure can be found in Appendix 3.13. For OU κ, the first panel shows that the

median of the variance of the return decreases from $1.47 \cdot 10^{-4}$ to $0.56 \cdot 10^{-4}$ (almost a factor 3) in the presence of the dragon-slayer. Already a fraction of 5% of initial wealth owned by the dragon-slayer decreases the variance by 28%. The skewness of the return is pushed closer to zero as the impact of the dragon-slayer increases. Without the dragon-slayer, the median skewness of the return is -0.081 while, with 50% dragon-slayer, the skewness is -0.037. The median excess kurtosis decreases by 71.1% from 1.50 to 0.43 when the dragon-slayer's wealth fraction increases from 0 to 50%. The median of the absolute value of the 1%-VaR decreases from 0.0294 to 0.0180 for the same change of the dragon-slayer's wealth fraction. The interquartile ranges of all four measures also decrease with the increase of the dragon-slayer's wealth fraction in the market. In particular, the decrease of the variance and of the absolute value of the 1%-VaR demonstrate clearly that the dragon-slayer stabilizes the market. A final noteworthy observation is that the beneficial impact of the dragon-slayer is stronger for small wealth fractions and its marginal effect decreases as its wealth fraction increases. Thus, already a small initial dragon-slayer's wealth fraction can stabilize the market.

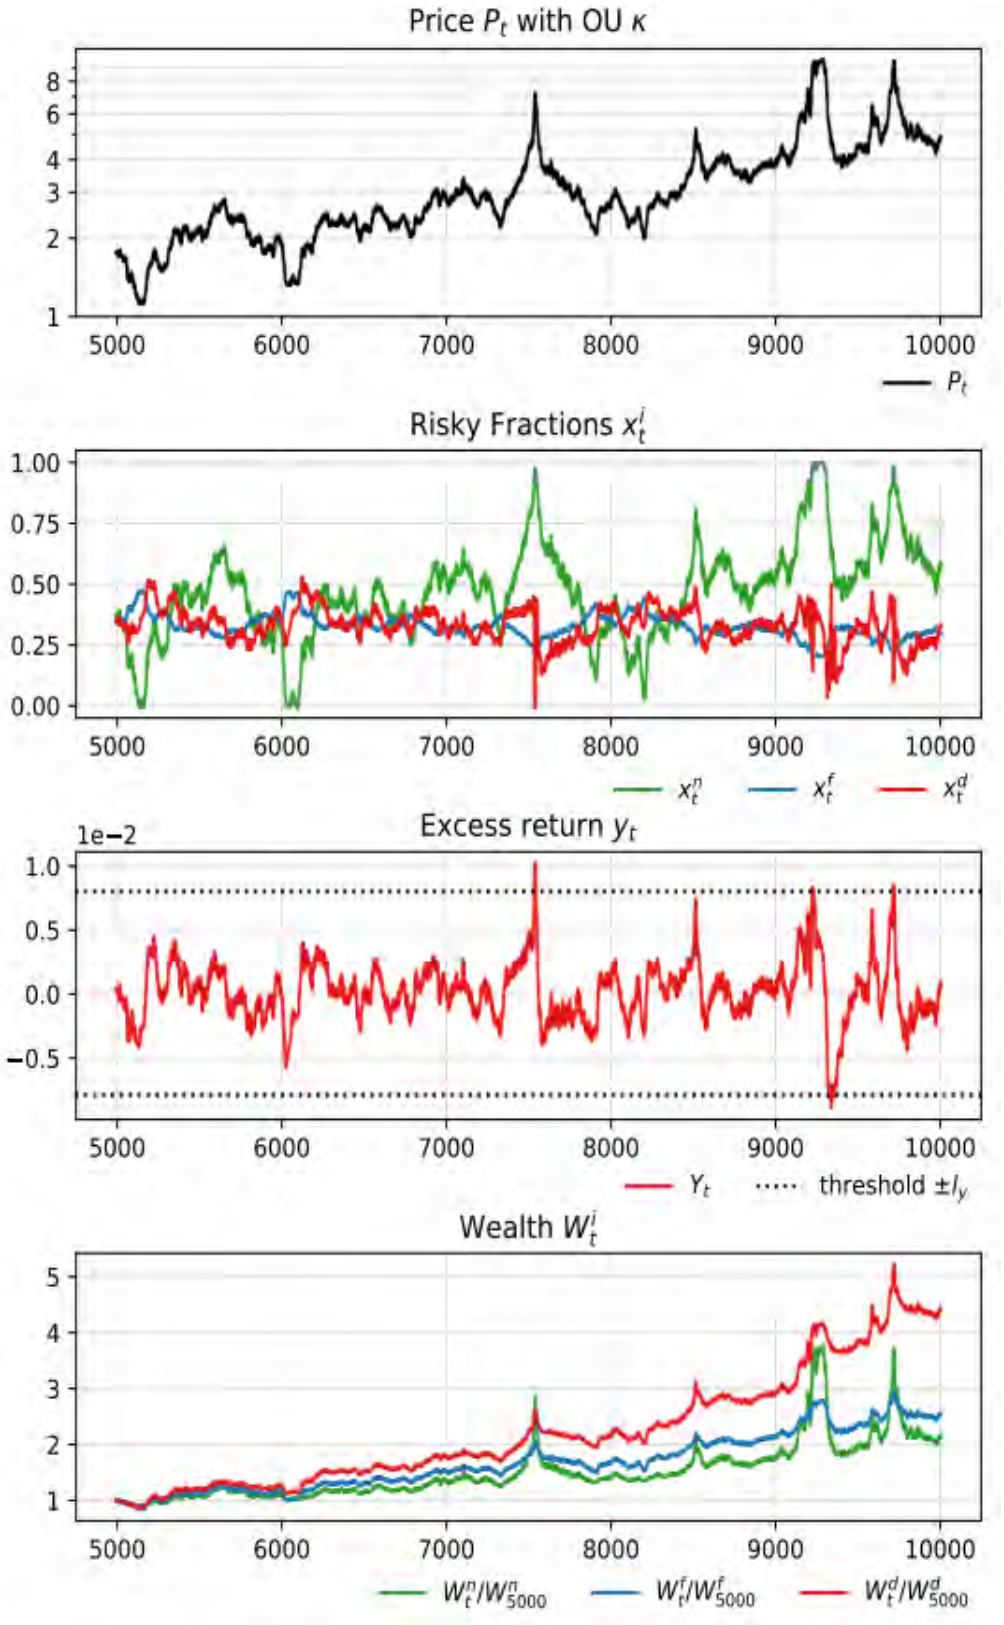

Figure 3.2: Example of a simulated price path P_t with OU κ. The upper panel shows the price P_t in linear-log scale, the second panel shows the risky fraction invested by the three trader types, the third panel shows the excess return momentum y_t, with the dotted lines indicating the threshold $\pm l_y$. The last panel shows the wealth of the three trader types over time. The dragon-slayer has no market impact in the simulation, i.e. these simulations are performed in the case where their wealth is negligible compared to that of the other traders. The parameters used in the strategy of the dragon-slayer are $a = 0.98$, $l_y = 0.008$, $w_y = 0.035$, and $s = 0.0005$. The 5000 time steps that are shown correspond to approximately 20 years given than one time-step corresponds to 1 trading day.

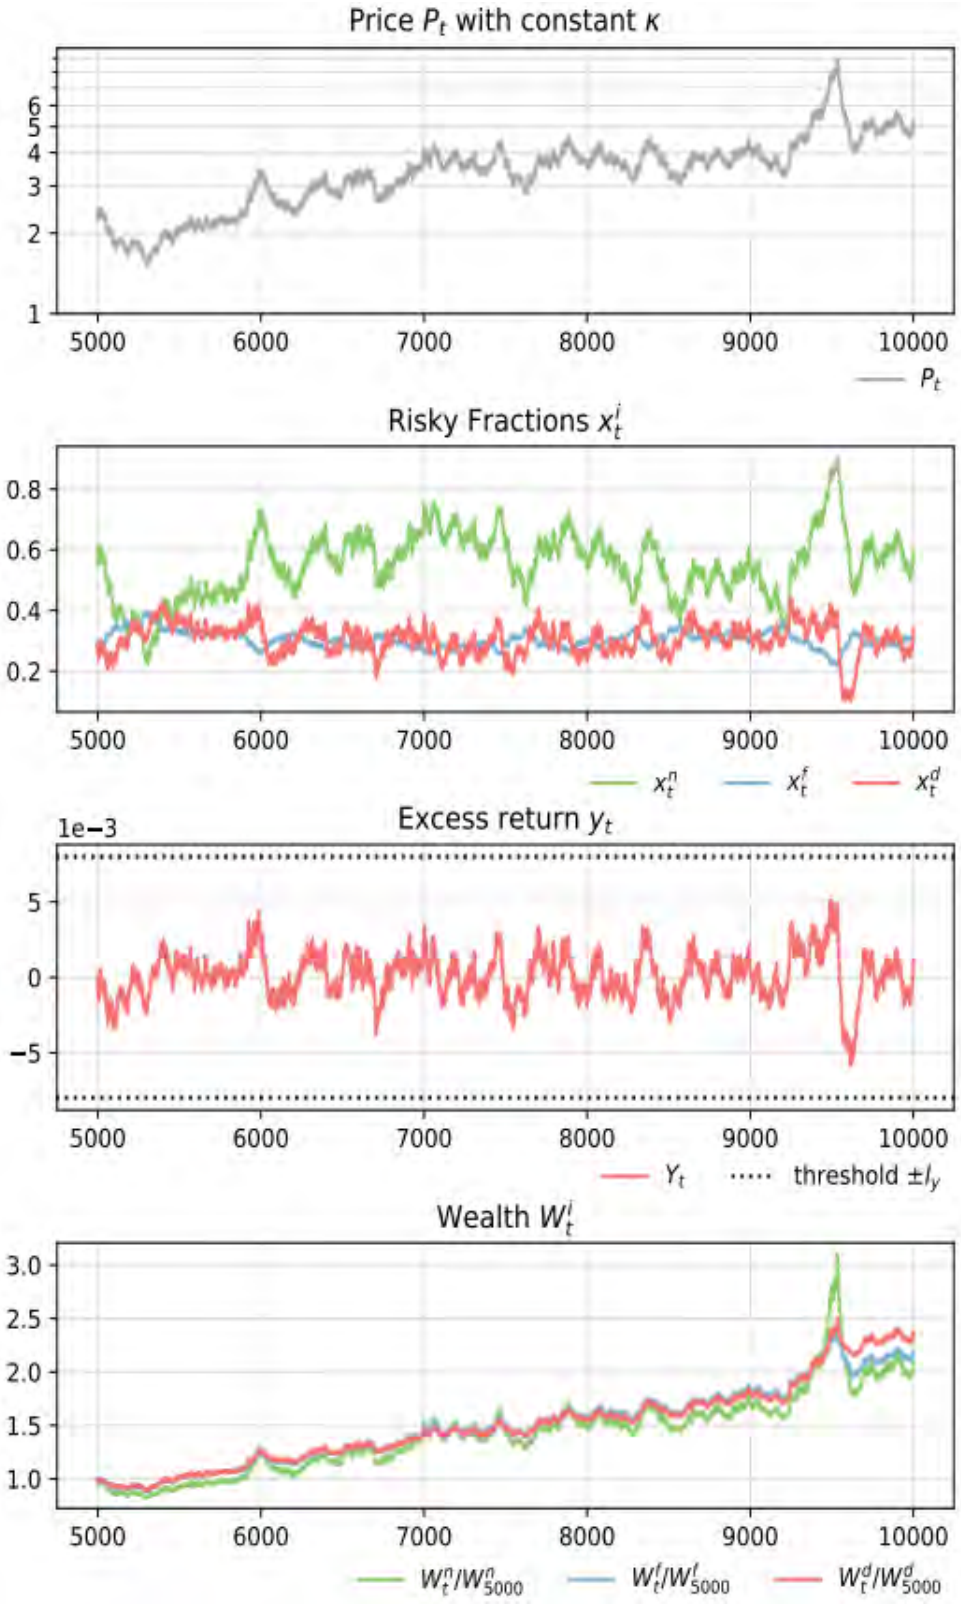

Figure 3.3: Same as figure 3.2, but with constant $\kappa = 0.98\kappa_c$.

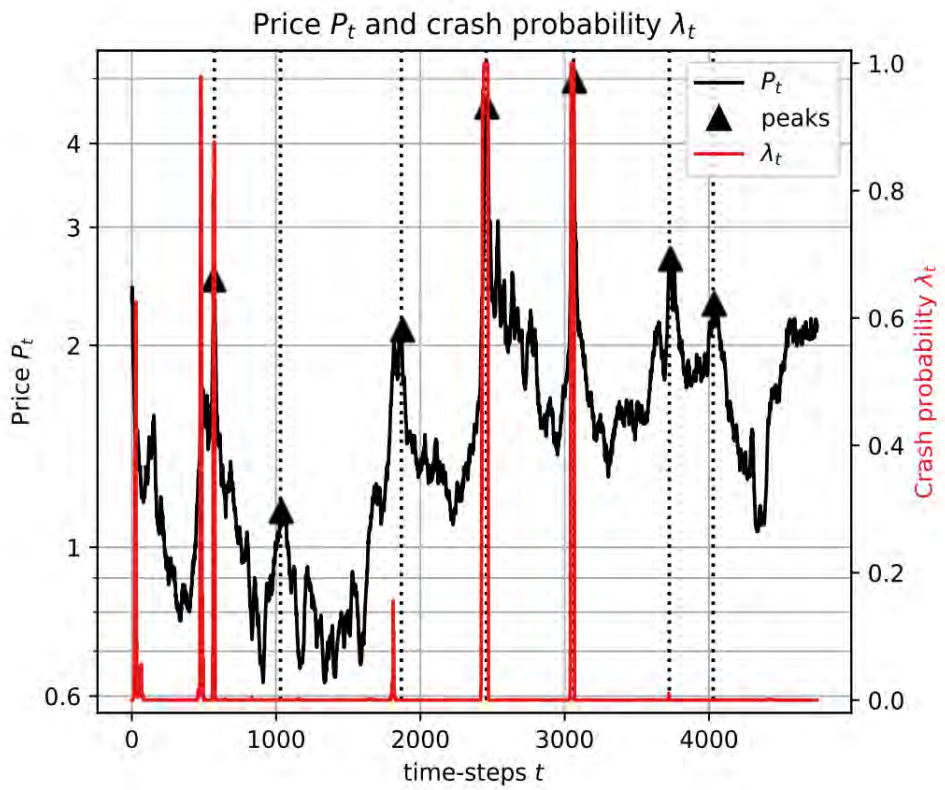

Figure 3.4: Simulated price path (logarithmic scale on the left axis) with the posterior identified peaks indicated as triangles (see main text for their definition) and the crash probability λ_t estimated by the dragon-slayer and given by expression (3.20) on the right axis. The parameters are $a = 0.95$, $l_y = 0.004$, $\bar{r} = 0.00016$, and $s = 0.0005$.

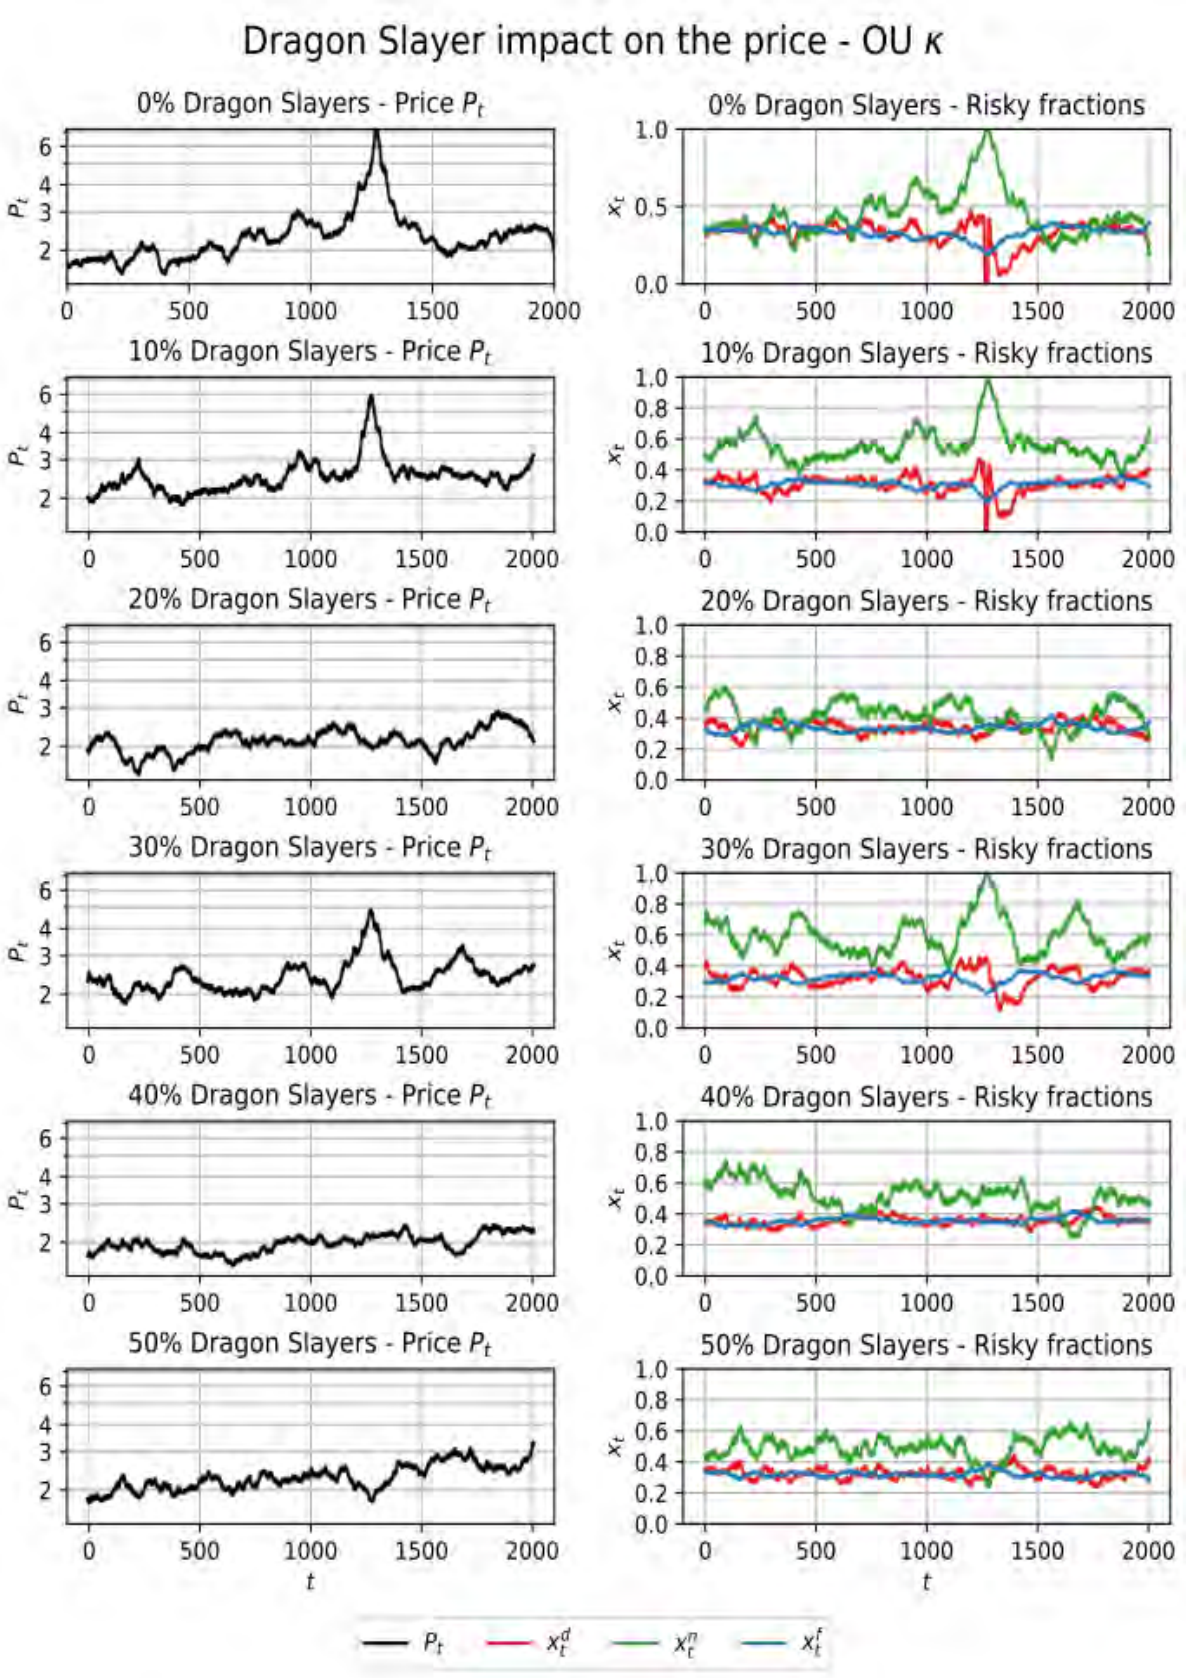

Figure 3.5: Section of a price path (left) and the corresponding fractions invested in the risky asset for the three trader types (right). Each row of panels shows a different dragon-slayer wealth fraction in the market, from 0% (top) to 50% (bottom). The social coupling strength κ follows the OU process (3.12). The dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

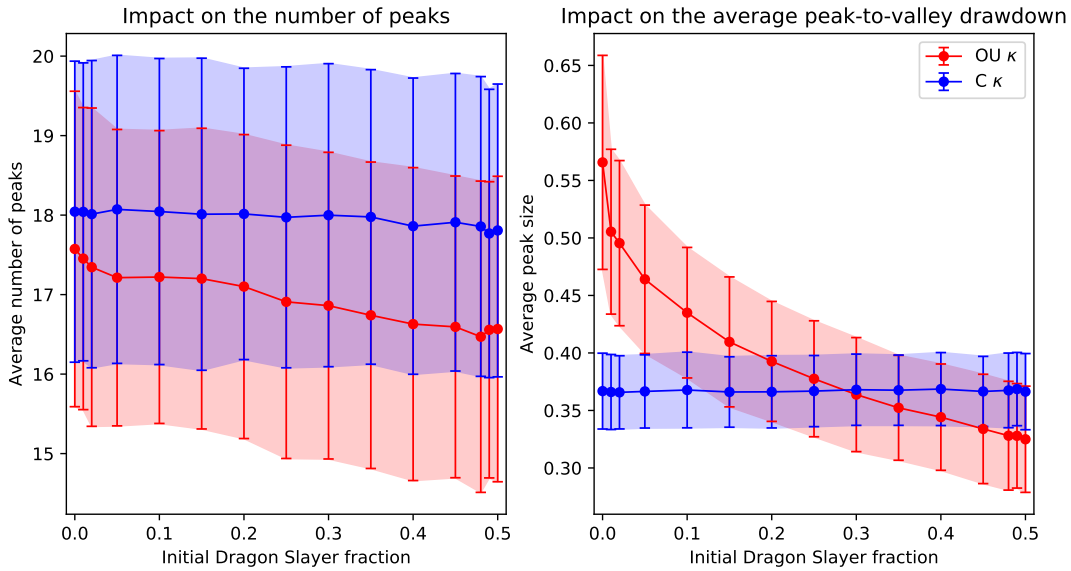

Figure 3.6: Quantification of the impact of the dragon-slayer strategy on the price of the risky asset. The average number of peaks and the average peak-to-valley-drawdown are shown as a function of the initial dragon-slayer wealth fraction in the risky asset. The average quantities are calculated over 1000 simulations with 1000 different random seeds. The error bars represent one standard deviation. Each quantity is calculated over price realisations occurring between $t = 5000$ and $t = 17500$, which corresponds to approximately 50 years. The dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

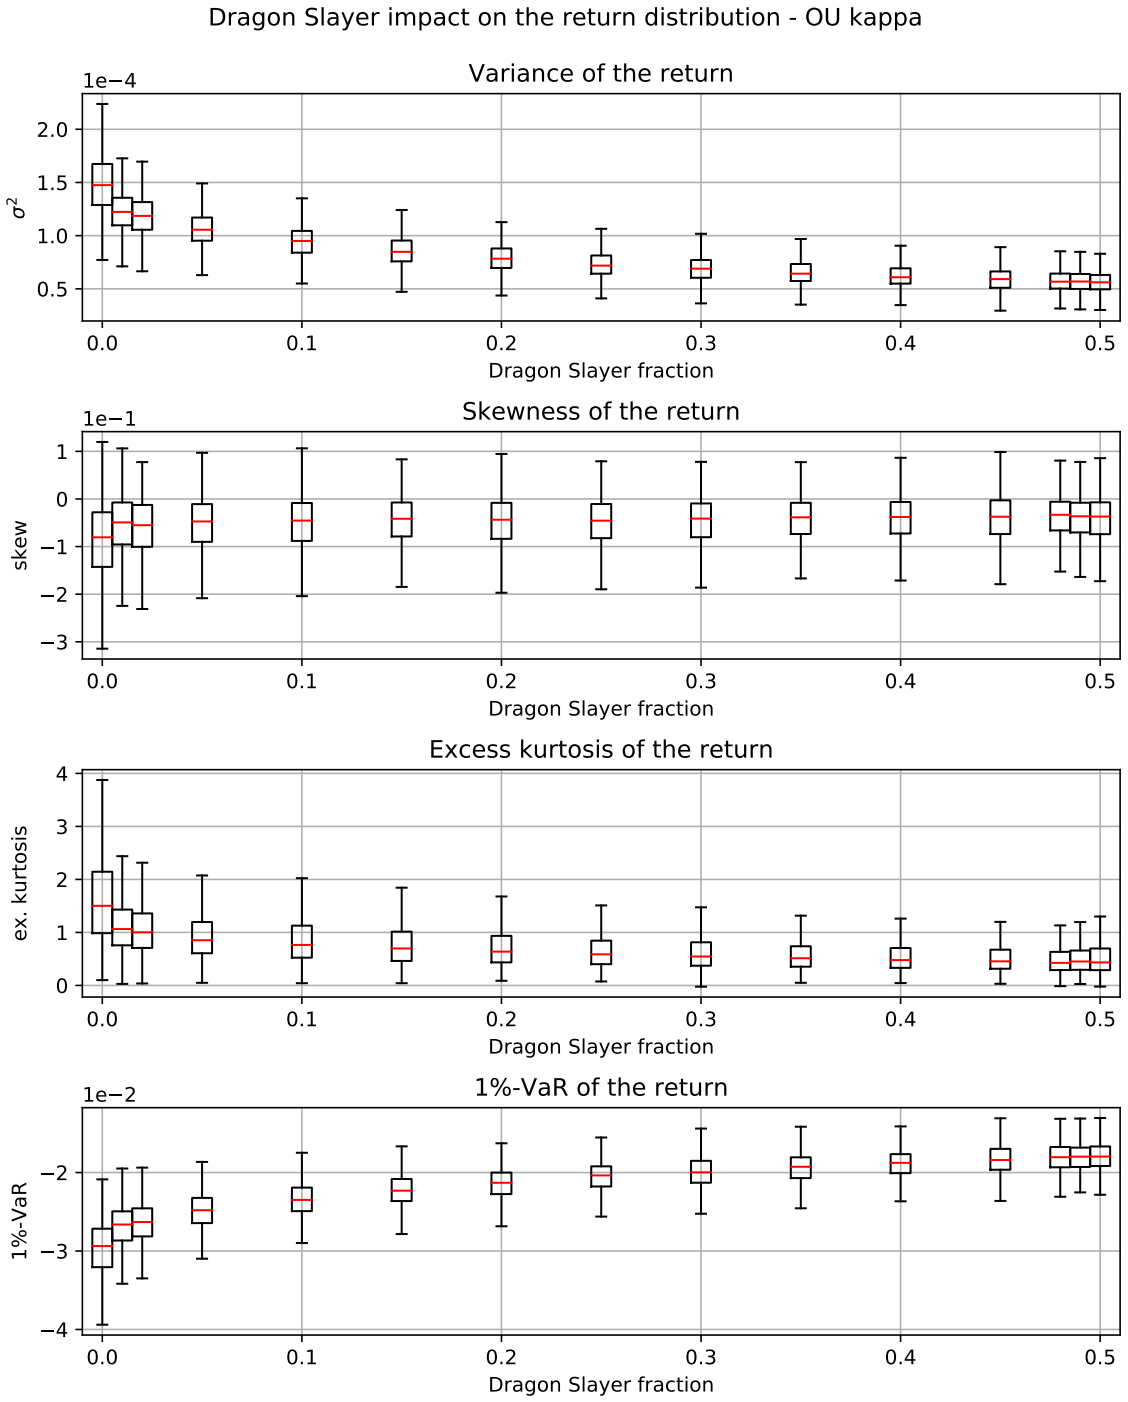

Figure 3.7: Boxplots of the dragon-slayer's impact on the statistics of returns of the risky asset for OU κ, calculated over 1000 simulations with different random seeds for each dragon-slayer fraction. The red lines show the median, the bottom and top of the boxes correspond to the 25% and 75% quartiles and the whiskers indicate the largest and lowest observed value within 1.5 times the interquartile range. The four panels show the variance, skewness, excess kurtosis, and 1%- VaR of the return of the risky asset. For each simulation, the parameters are calculated over the time interval from $t = 5000$ to $t = 17500$. The dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

3.4 Impact of the dragon-slayer on the traders' performance

The previous section showed how the dragon-slayer removes bubbles from the risky asset price and decreases its volatility without influencing the long-term growth rate of the asset. Here, we focus on how his presence affects the performance of the three types of traders present in the market.

Figure 3.8 shows the average Sharpe ratio of the three trader types, calculated over 1000 simulations with different random seeds, as a function of the initial dragon-slayer's wealth fraction in the market. The corresponding figures for $a = 0.95$ and $a = 0.99$ can be found in Appendix 3.14 and 3.15. The Sharpe ratio is calculated over the interval $t \in [5000, 17500]$, which corresponds to 50 years. Figure 3.8 shows that, for OU κ for which bubbles emerge naturally from the strategies of the noise traders, the risk-adjusted return of all three trader types increases with increasing initial wealth fraction of the dragon-slayer. The increases of the Sharpe ratios are economically very significant, from approximately 0.3 to 0.5 for noise traders, from 0.6 to 0.8 for fundamentalists and from 0.8 to close to 1 for the dragon-slayer, as his wealth fraction increases from 0 to 50%. In contrast, for constant κ, the Sharpe ratios remains approximately constant. This is consistent with the observation from figure 3.3, that shows very little impact of the dragon-slayer on the price dynamics of the risky asset for constant κ for which bubbles do not appear. This comes from the fact that, in absence of bubbles, the dragon-slayer's strategy reduces to that of the fundamentalists. This leads to conclude that, in absence of bubbles, the performance of traders remains unchanged when increasing the wealth of the dragon-slayer while it improves significantly when the dragon-slayer removes bubbles that were otherwise present.

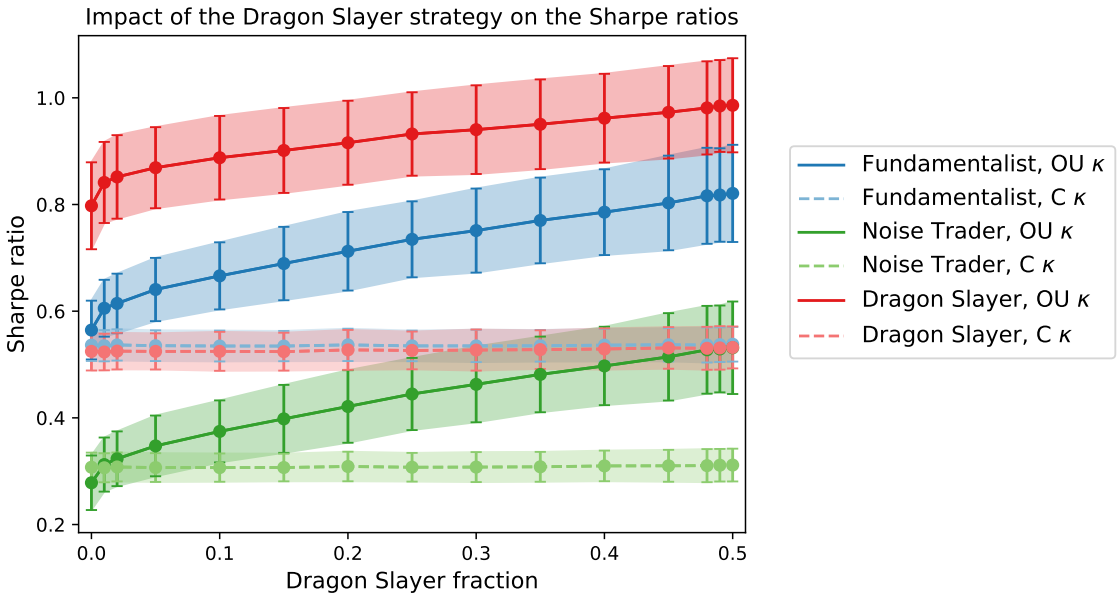

Figure 3.8: Average Sharpe ratios of the three trader types as a function of the dragon-slayer's wealth fraction for OU κ (solid lines) and constant κ (dashed lines) calculated over 1000 realisations with different random seeds and over approximately 12500 time steps. The error bars represent one standard deviation. The dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

While figure 3.8 provides an in-depth analysis of one specific dragon-slayer strategy corresponding to a specific set of parameters, figure 3.9 analyses the sensitivity of the traders' Sharpe ratio to four parameters of the dragon-slayer strategy and to the average growth rate $\bar{r}$ of the risky asset. The three top panels (resp. bottom panel) show the average annualized Sharpe ratios of the traders with OU (solid line) and constant κ (dotted line) as a function of one of the dragon-slayer strategy parameters (resp. $\bar{r}$).

Each scenario is simulated 1000 times with different random seeds, the filled circles indicate the mean value of the Sharpe ratios, each calculated from $t = 5000$ to $t = 17500$. The error bars represent one standard deviation. The dragon-slayer is endowed with 10% of the initial wealth and this is fixed over all simulations when varying the parameters. The Sharpe ratios with 0% dragon-slayer are included as reference values. The first panel shows the sensitivity of the Sharpe ratios to variations in the dragon-slayer's memory parameter a used in the estimation by the dragon-slayer of the excess return momentum y_t (3.19). We scan the values $a \in [0.9, 0.95, 0.98, 0.99, 0.996]$. This means the memory length $1/(1-a)$ is varied between 10 trading days and 250 trading days. For OU κ, the fundamentalists and noise traders performances are better in the presence of a dragon-slayer with small a, corresponding to a short memory. The performance is found to be best for $a = 0.95$, which corresponds to a memory of 20 trading days. For all analysed values of a, the presence of the dragon-slayer is always beneficial. The noise traders also enjoy improved performance when the dragon-slayer is present for all a except for $a = 0.996$ (250 days), where their Sharpe ratio is slightly below the reference value. As expected, the impact of the dragon-slayer is small for constant κ, as the Sharpe ratios are very close to the reference Sharpe ratios for all analyzed memory parameters. For constant κ, for both fundamentalists and noise traders, the Sharpe ratio in the presence of the dragon-slayer is slightly above the reference value for $a \leq 0.95$ and slightly below it for $a > 0.95$. In general, the dragon-slayer with a shorter memory length can react faster to changes in the momentum, performs better and is more beneficial to the market.

The second panel of figure 3.9 shows the Sharpe ratios of the three trader types with 10% dragon-slayer as a function of the overpricing threshold $l_y \in [0.00016, 0.0008, 0.0032, 0.008, 0.016]$. With the long-term daily growth rate $r_d = 0.00016$ (4% annualised), this is equivalent to an excess return threshold between r_d and $100 \cdot r_d$. We find that the risk-adjusted return of the three traders is larger with the larger thresholds l_y. With small l_y's, the dragon-slayer tends to overreact to small deviations of the risky asset price from the long term trend controlled by the average return r_d. Thus, for $l_y = 0.0032$, the dragon-slayer performs even worse than the other traders. The fundamentalists average Sharpe ratio slightly decreases to 0.5342 when a dragon-slayer with $l_y = 0.0032$ is present in the market. The threshold $l_y = 0.0032$ means that the dragon-slayer considers a persistent excess return of 0.32% per day to be unsustainable. While this is 50 times the long-term growth rate, it is significantly smaller than the expected daily volatility which is approximately 1%. Thus, the dragon-slayer jumps between buying and selling the risky asset that are illustrated in Figure 3.1 outside of the dotted lines representing the threshold levels. These rapid portfolio reallocations result in a destabilisation of the market. For thresholds that are larger than 0.5%, the presence of the dragon-slayer improves the performance of the other traders. The conclusion, which should not be a surprise, is that the dragon-slayer should err on the side of discriminating diagnostics of bubbles to avoid over-reacting on too many false positives.

The weight w_y enters in the determination of the expectation $E^d_{r_t}$ of the return by the dragon-slayer in expression (3.21). A dragon-slayer who chooses $w_y = 0$ is identical to the fundamentalists, while the larger w_y is, the more he is concerned with bubbles. Thus, w_y controls the amplitude of the expectation of the dragon-slayer concerning the market return, as shown in figure 3.1. The third panel of figure 3.9 shows that, for OU κ, a medium weight value between $w_y = 0.05$ and 0.1 is optimal for all traders. For all analyzed values of w_y, the traders obtain better Sharpe ratios than without the dragon-slayer. However, for constant κ, large weights ($w_y > 0.05$) slightly decrease the risk-adjusted return of the traders compared to the reference value, because the dragon-slayer uses less the stabilizing fundamentalist strategy of investing proportionally to the dividend-price ratio.

The parameter s enters in the definition of the probability λ_t that the overpricing will result in a crash according to expression (3.20). It controls the reaction to changes in y_t near the threshold l_y. As illustrated in figure 3.1, a large s results in a slower and smoother reaction, while a small s results in an immediate readjustment of the portfolio when the threshold l_y is reached. The fourth panel in figure 3.9 shows that the Sharpe ratios of all traders are monotonously decreasing as a function of s. Thus, smaller s values are beneficial for all trader types, implying that the traders benefit

from a dragon-slayer who reacts determinedly to the detected overpricing. The Sharpe ratios for constant κ are not influenced by a change of s, because y_t fluctuates around 0 and only the reaction to larger deviations is influenced by s. For OU κ, the traders benefit from the presence of the dragon-slayer for all analyzed values of s. The Sharpe ratios are on average significantly larger than in the reference simulations without the dragon-slayer.

The last bottom panel of figure 3.9 shows the average Sharpe ratios of all traders as a function of the expected long-term growth rate $\bar{r}$ of the risky asset. In the default parameter setting, the dragon-slayer uses the true long-term growth rate of the market, which is equal to the dividend growth rate r_d. This is identical to the fundamentalist strategy. However, in real markets, it is difficult to have an accurate estimation of the long-term growth rate of an asset. The figure shows that, even with a wrong estimation of the true growth rate, the traders benefit from the presence of the dragon-slayer. When the dragon-slayer underestimate the growth-rate and use 0.00008 (2% annualised) or 0.00012 (3% annualised) instead of the correct 0.00016 (4% annualised), the traders perform even better than with the true growth rate. When choosing a smaller $\bar{r}$, the overpricing y_t increases by the difference between $\bar{r}$ and the real growth rate, because it is defined as the exponential moving average of the difference between observed return and expected long-term return. However, with an inaccurate estimation of $\bar{r}$, the reaction to bubbles is not symmetric anymore. Thus, choosing a smaller $\bar{r}$ has the same impact as shifting the threshold l_y to the left. In any case, the most important conclusion in practice from the simulations presented in this bottom panel is that the conclusion about the beneficial influence of the dragon-slayer is robust with respect to an error of more than 50% (2% to 6% around the true 4%) in the estimation of the long-term growth rate. This is not surprising given that transient bubbles are characterised by much larger short-term growth rates, which make their detection robust with respect to a miscalibration of the long-term growth rate.

In summary, figure 3.9 shows that the dragon-slayer strategy improves significantly the performance of the other traders over a wide range of strategy parameters. Shorter memory lengths, larger overpricing thresholds, a medium weight on the bubble and crash diagnostic, and a small expected long-term growth rate bring exceptional benefit to the three trader types.

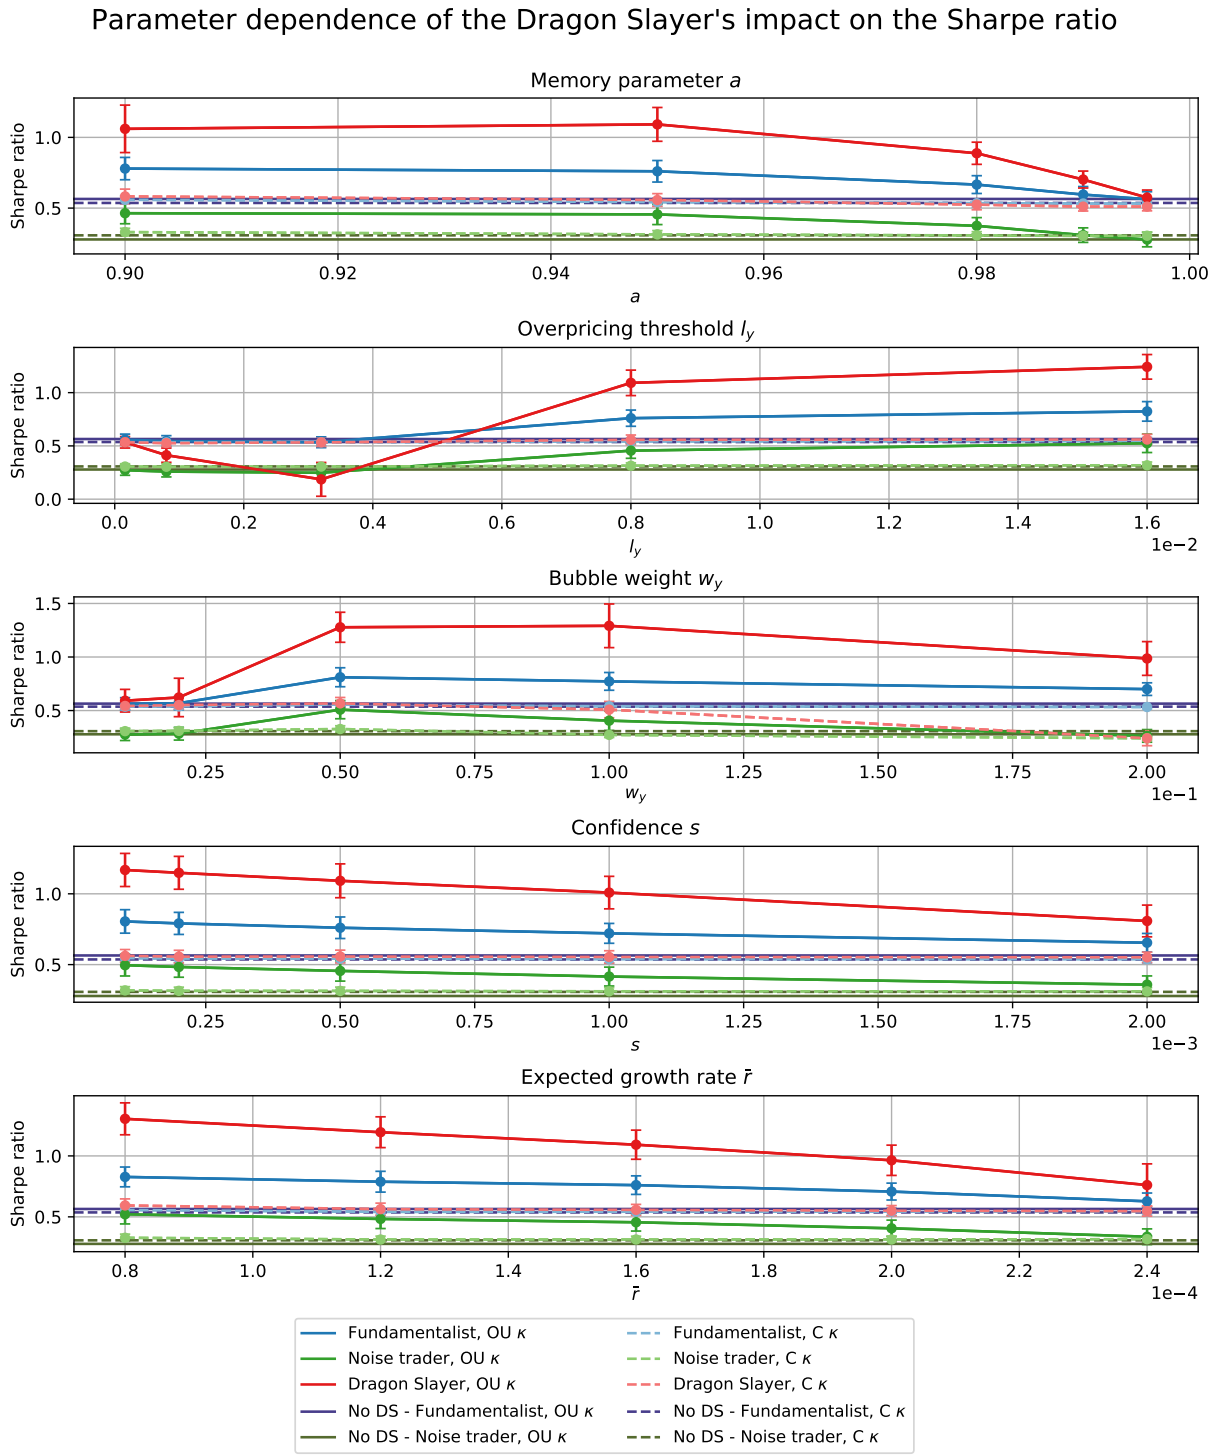

Figure 3.9: Dependence of the Sharpe ratios of the three types of traders as a function of three parameters of the dragon-slayer's strategy and of the average growth rate $\bar{r}$ of the risky asset. The filled circles represent the mean values over 1000 simulations with different random seeds, with the error bars representing one standard deviation. The dragon-slayer is given an initial wealth corresponding to 10% of the total initial wealth over all traders. The solid lines correspond to OU κ, and the dotted lines correspond to constant κ. The reference Sharpe ratios averaged over 1000 simulations in the absence of the dragon-slayer are shown in blue-grey for the fundamentalists and green-grey for the noise traders. In each panel, a single parameter is varied, while the other parameters have the default values from the parameter set $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$ (corresponding to an annualised return of 4%), and $s = 0.0005$. The Sharpe ratios are annualized values calculated over the time interval from $t = 5000$ to $t = 17500$. In the bottom panel, the range of variation of $\bar{r}$ from 0.00008 to 0.00024 (daily) corresponds to a range from 2% to 6% annualised.

3.5 Conclusion

We have presented an extension of a previously studied agent-based model (ABM) originally developed by Kaizoji et al. (2015), which is characterised by the spontaneous formation of bubbles and crashes emerging from the interaction between fundamentalists and noise traders. We have introduced a third type of traders, called dragon-slayer, who represents a policy maker who has the objective of preventing bubbles and crashes by trading between a risky and a risk-free asset. The dragon-slayer diagnoses burgeoning bubbles by forming an expectation of the future return of the risky asset in the form of an exponential moving average of the excess return over the long-term return. When this excess return momentum exceeds a threshold that the dragon-slayer estimates as an unsustainable level, he forms a prediction that a crash may happen with a probability given by a logistic function of the excess return momentum. Equipped with this bubble diagnostic, the dragon-slayer constructs his trading strategy similarly to the fundamentalists but with the advantage of using a real-time dynamical estimation of a transient expected excess return. Specifically, the policy maker invests in the risky asset when he detects a small deviation of the return from the long-term growth rate in order to construct an inventory that he draws upon later to fight future market exuberance. Then, when this deviation between the current growth rate and the long-term growth rate exceeds the policy maker's tolerance level, he starts to sell the risky asset that he has accumulated earlier, in a countercyclical prevent future price increase.

We have found that the dragon-slayer succeeds in preventing bubbles and crashes in the ABM. In simulations without bubbles, the dragon-slayer behaves similarly to the fundamentalists and his impact is negligible. This is a good property in the sense of that any cure should first follow the principle of "Primum non nocere" (first, do no harm). In simulations where bubbles form spontaneously as a result of the noise traders's strategies, the average drawdown is decreased from 57% in absence of the dragon-slayer to 33% when the dragon-slayer is initialized with 50% of the total wealth so that his market impact is very significant. Concomitantly, the average number of peaks is reduced by 5.7%. The stabilising effect of the intervening policy maker is also reflected in the return dynamics of the risky asset. An initial wealth fraction of just 5% for the dragon-slayer reduces the variance of the return by 28%. A larger wealth fraction reduces the variance up to 62%. Furthermore, the skewness of the returns is pushed closer to zero, and the dragon-slayer decreases significantly the excess kurtosis and the absolute value of the 1%-VaR.

While removing bubbles, we find that the dragon-slayer strategy does not affect the long-term growth rate of the risky asset. For all analyzed scenarios, the growth rate is found close to the theoretical value, which is determined by the growth rate of the dividend process underlying the risky asset. This combination of fewer bubbles and crashes, and the stability of the long-term growth rate, leads the dragon-slayer to provide improved performance for the other traders in terms of their risk-adjusted return. Thus, the dragon-slayer increases the wealth of all market participants.

Finally, we have tested the sensitivity of these results to variations of the key parameters of the strategy of the dragon-slayer. We investigated the average Sharpe ratios of fundamentalists and noise traders with a dragon-slayer possessing a wealth amounting to 10% of the total wealth. We found that the dragon-slayer strategy is beneficial to the other traders over a wide range of parameters of his strategy. In general, a faster reacting dragon-slayer with a shorter memory length provides the largest benefit to the other traders. Furthermore, the traders perform better in the presence of the dragon-slayer in the whole analysed range of expected growth rates of the risky asset. Thus, a dragon-slayer who uses a miscalibrated market growth rate that is even larger than 50% off the true growth rate still provides significant increase of the investment performance for all traders.

In sum, our simulations have shown that direct intervention in the stock market to prevent bubbles and drawdowns can be very effective and beneficial for all involved traders.

3.6 Appendix

3.6.1 The Equilibrium Market Price

The market price is defined following Walras' theory of general equilibrium (Walras, 1954). This means that each trader formulates their excess demand for the next time step and the price is calculated as the equilibrium in which supply equals demand. Thus, the market clearing condition states an equilibrium between supply and demand. This is formulated as:

$$0 = \Delta D^d_{t-1\to t} + \Delta D^n_{t-1\to t} + \Delta D^f_{t-1\to t} \tag{3.28}$$

Next, the explicit formulations of the demand from (3.23), (3.8), and (3.18) presented in the main text are inserted into the market clearing condition (3.28):

$$\begin{aligned}
0 =& \Delta D^d_{t-1\to t} + \Delta D^n_{t-1\to t} + \Delta D^f_{t-1\to t} \\
=& W^d_t \left(x^d_t \left[x^d_{t-1} \left(r_t + \frac{d_t}{P_{t-1}} - r_f \right) + r_f + 1 \right] - x^d_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^n_{t-1} \left(x^n_t \left[x^n_{t-1} \left(r_t + \frac{d_t}{P_{t-1}} - r_f \right) + r_f + 1 \right] - x^n_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^f_{t-1} \left(x_t \left[x^f_{t-1} \left(r_t + \frac{d_t}{P_{t-1}} - r_f \right) + r_f + 1 \right] - x^f_{t-1} \frac{P_t}{P_{t-1}} \right) \\
=& W^d_{t-1} \left(\frac{E^d_{r_t} + \frac{d_t}{P_t}(1 + r_d) - r_f}{\gamma \sigma^2} \left[x^f_{t-1} \left(r_t + \frac{d_t}{P_{t-1}} - r_f \right) + r_f + 1 \right] - x^f_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^f_{t-1} \left(\frac{E^f_{r_t} + \frac{d_t}{P_t}(1 + r_d) - r_f}{\gamma \sigma^2} \left[x^f_{t-1} \left(r_t + \frac{d_t}{P_{t-1}} - r_f \right) + r_f + 1 \right] - x^f_{t-1} \frac{P_t}{P_{t-1}} \right) \\
&+ W^n_{t-1} \left(x^n_t \left[x^n_{t-1} \left(r_t + \frac{d_t}{P_{t-1}} - r_f \right) + r_f + 1 \right] - x^n_{t-1} \frac{P_t}{P_{t-1}} \right)
\end{aligned}$$

Finally, the equation is multiplied by P_t, the return $r_t := \frac{P_t}{P_{t-1}} - 1$ is inserted and formula is organised in powers of P_t to obtain

$$
\begin{aligned}
0 =& W^d_{t-1}\left(\frac{E_{r^d_t}P_t + d_t(1+r_d) - r_f P_t}{\gamma\sigma^2}\left[x^r_{t-1}\left(\frac{P_t}{P_{t-1}} - 1 + \frac{d_t}{P_{t-1}} - r_f\right) + r_f + 1\right] - x^r_{t-1}\frac{P_t^2}{P_{t-1}}\right) \\
&+ W^f_{t-1}\left(\frac{E_{r^f_t}P_t + d_t(1+r_d) - r_f P_t}{\gamma\sigma^2}\left[x^f_{t-1}\left(\frac{P_t}{P_{t-1}} - 1 + \frac{d_t}{P_{t-1}} - r_f\right) + r_f + 1\right] - x^f_{t-1}\frac{P_t^2}{P_{t-1}}\right) \\
&+ W^n_{t-1}\left(x^n_t P_t\left[x^n_{t-1}\left(\frac{P_t}{P_{t-1}} - 1 + \frac{d_t}{P_{t-1}} - r_f\right) + r_f + 1\right] - x^n_{t-1}\frac{P_t^2}{P_{t-1}}\right) \\
=& \frac{P_t^2}{P_{t-1}}\left(W^d_{t-1}x^d_{t-1}\left(\frac{E_{r^d_t} - r_f}{\gamma\sigma^2} - 1\right) + W^f_{t-1}x^f_{t-1}\left(\frac{E_{r^f_t} - r_f}{\gamma\sigma^2} - 1\right) + W^n_{t-1}x^n_{t-1}\left(x^n_t - 1\right)\right) \\
&+ P_t\left(W^d_{t-1}\left(x^d_{t-1}\frac{d_t(1+r_d)}{\gamma\sigma^2 P_{t-1}} + \frac{E_{r^d_t} - r_f}{\gamma\sigma^2}\cdot\left(x^d_{t-1}\left(\frac{d_t}{P_{t-1}} - r_f - 1\right) + r_f + 1\right)\right)\right. \\
&+ W^f_{t-1}\left(x^f_{t-1}\frac{d_t(1+r_d)}{\gamma\sigma^2 P_{t-1}} + \frac{E_{r^f_t} - r_f}{\gamma\sigma^2}\cdot\left(x^f_{t-1}\left(\frac{d_t}{P_{t-1}} - r_f - 1\right) + r_f + 1\right)\right) \\
&\left.+ W^n_{t-1}x^n_t\left(x^n_{t-1}\left(\frac{d_t}{P_{t-1}} - r_f - 1\right) + r_f + 1\right)\right) \\
&+ W^d_{t-1}\left(\frac{d_t(1+r_d)}{\gamma\sigma^2}\left(x^d_{t-1}\left(\frac{d_t}{P_{t-1}} - r_f - 1\right) + r_f + 1\right)\right) \\
&+ W^f_{t-1}\frac{d_t(1+r_d)}{\gamma\sigma^2}\left(x^f_{t-1}\left(\frac{d_t}{P_{t-1}} - r_f - 1\right) + r_f + 1\right)
\end{aligned}
$$

Next, we define some auxiliary variables to simplify the equation.

$$v^d_1 := \frac{E_{r^d_t} - r_f}{\gamma\sigma^2} \tag{3.29}$$

$$v^f_1 := \frac{E_{r^f_t} - r_f}{\gamma\sigma^2} \tag{3.30}$$

$$v_2 := \frac{d_t}{P_{t-1}} - r_f - 1 \tag{3.31}$$

$$v_3 := \frac{d_t(1+r_d)}{\gamma\sigma^2} \tag{3.32}$$

The resulting equation is given in the following:

$$
\begin{aligned}
0 =& \frac{P_t^2}{P_{t-1}}\left(W^d_{t-1}x^d_{t-1}\left(v^d_1 - 1\right) + W^f_{t-1}x^f_{t-1}\left(v^f_1 - 1\right) + W^n_{t-1}x^n_{t-1}\left(x^n_t - 1\right)\right) \\
&+ P_t\left(W^d_{t-1}\left(x^d_{t-1}\frac{v_3}{P_{t-1}} + v^d_1\cdot\left(x^d_{t-1}v_2 + r_f + 1\right)\right)\right. \\
&\left.+ W^f_{t-1}\left(x^f_{t-1}\frac{v_3}{P_{t-1}} + v^f_1\cdot\left(x^f_{t-1}v_2 + r_f + 1\right)\right) + W^n_{t-1}x^n_t\left(x^n_{t-1}v_2 + r_f + 1\right)\right) \\
&+ W^d_{t-1}v_3\left(x^d_{t-1}v_2 + r_f + 1\right) + W^f_{t-1}v_3\left(x^f_{t-1}v_2 + r_f + 1\right)
\end{aligned} \tag{3.33}
$$

Equation (3.33) is a quadratic equation of the form $a_t P_t^2 + b_t P_t + c_t = 0$ where

$$a_t = \left(W_{t-1}^d x_{t-1}^d \left(v_1^d - 1 \right) + W_{t-1}^f x_{t-1}^f \left(v_1^f - 1 \right) + W_{t-1}^n x_{t-1}^n \left(x_t^n - 1 \right) \right) / P_{t-1} \tag{3.34}$$

$$b_t = W_{t-1}^d \left(x_{t-1}^d \frac{v_3}{P_{t-1}} + v_1^d \cdot \left(x_{t-1}^d v_2 + r_f + 1 \right) \right) \tag{3.35}$$

$$+ W_{t-1}^f \left(x_{t-1}^f \frac{v_3}{P_{t-1}} + v_1^f \cdot \left(x_{t-1}^f v_2 + r_f + 1 \right) \right) + W_{t-1}^n x_t^n \left(x_{t-1}^n v_2 + r_f + 1 \right) \tag{3.36}$$

$$c_t = W_{t-1}^d v_3 \left(x_{t-1}^d v_2 + r_f + 1 \right) + W_{t-1}^f v_3 \left(x_{t-1}^f v_2 + r_f + 1 \right) \tag{3.37}$$

The equation has the following two solutions:

$$P_t^{1,2} = \frac{-b_t \pm \sqrt{b_t^2 - 4 a_t c_t}}{2 a_t} \tag{3.38}$$

The positive solution of this equation is the price in the next time-step.

3.6.2 Parameter values

Parameter name	Explanation	Value
Market		
T	Simulation length	25000
$seed$	Random seed	6 digit number
Assets		
r_f	Risk free interest rate	0.01/250=0.00004
d_0	Initial dividend	0.04/250=0.00016
r_d	Expected growth rate of the dividend	0.04/250=0.00016
σ_d	Expected standard deviation of the dividend growth rate	0.000016
P_0	Initial price of the risky asset	1
σ_r	Expected standard deviation of the risky asset price	$\sqrt{0.10/250} = 0.02$
N^r	Number of risky assets	1
Noise traders		
x_o^n	Initial fraction of the risky asset held by the noise traders	0.3
W_0^n	Initial wealth of the noise traders	10^9
p	Switching probability	0.2
θ	Memory parameter	0.95
H_0	Initial momentum	0.00016
N^n	Number of noise traders	1000
Fundamentalists		
x_0^f	Initial fraction of the risky asset held by the fundamentalists	0.3
W_0^f	Initial wealth of the fundamentalists	0 to 10^9
E_{R_t}	Expected return of the risky asset $= r_d$	0.00016
Social coupling strength		
κ_0	Initial social coupling strength	0.98
μ_κ	Mean of the OU social coupling strength	0.98
η_κ	Mean reversion of the OU social coupling strength	0.11
σ_κ	Standard deviation of the OU social coupling strength	0.01

Table 3.2: Parameters characterising the model financial markets and used in the simulations. Dividends, interest rates and standard deviations are given as daily values (per time-step).

3.6.3 Additional figures

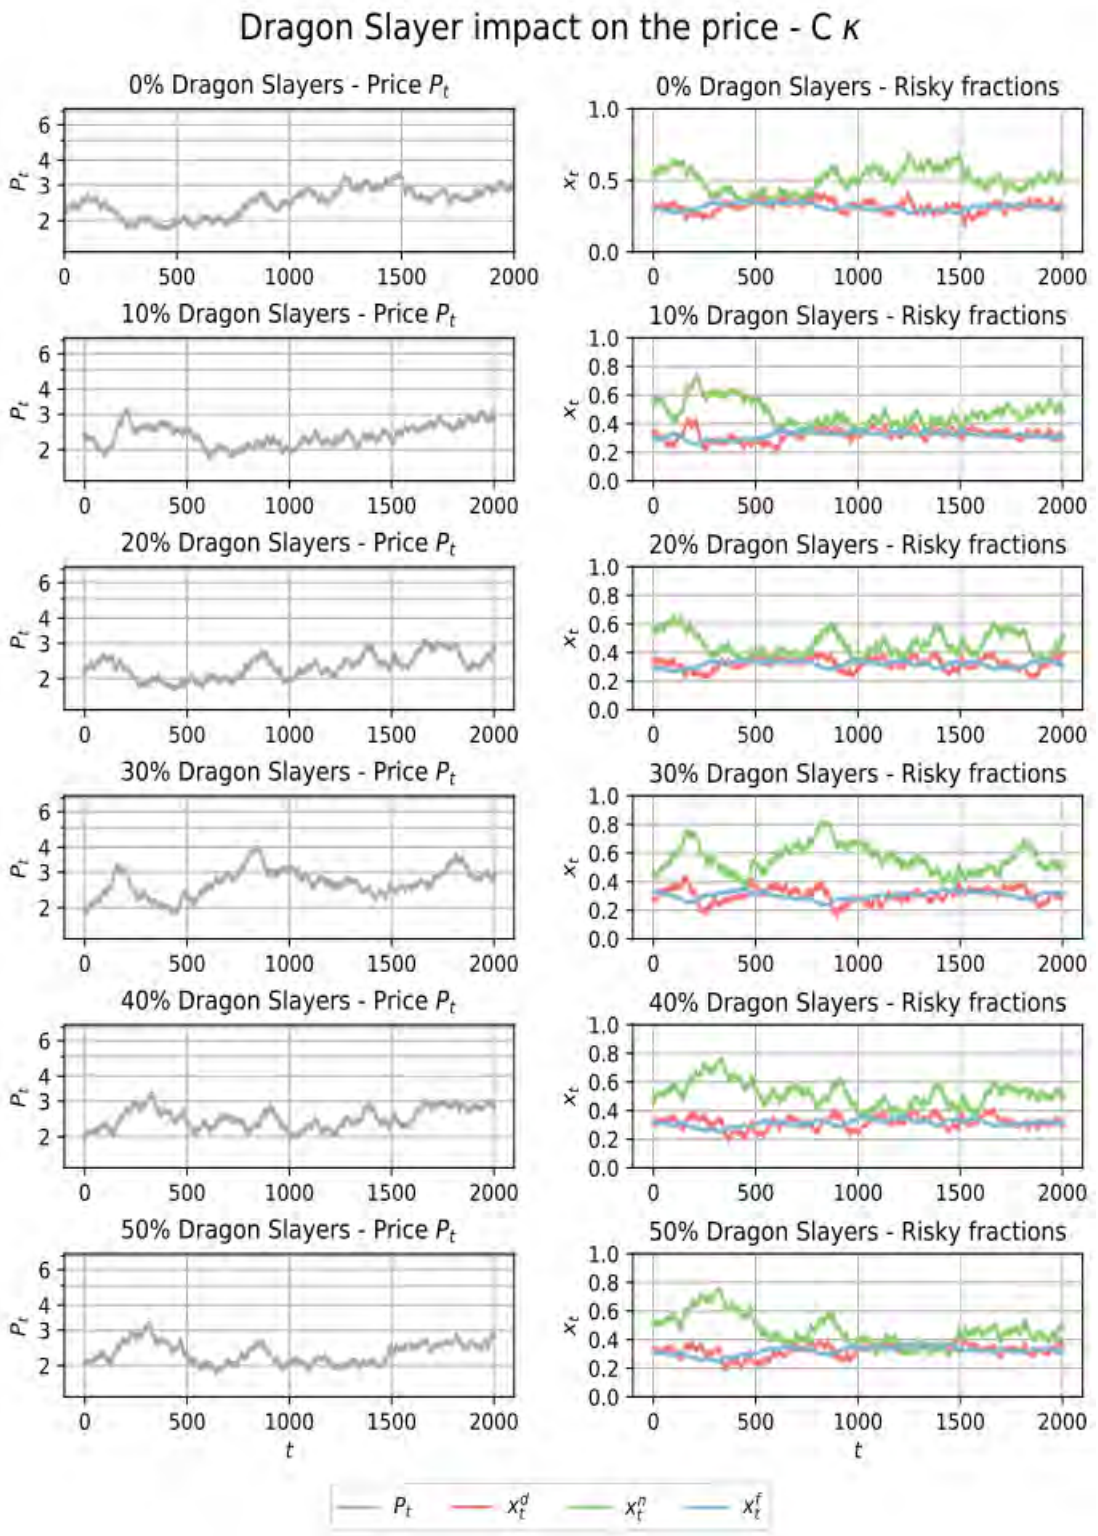

Figure 3.10: Section of a price path (left) and the corresponding fractions invested in the risky asset for the three trader types (right). Each row of panels shows a different dragon-slayer wealth fraction in the market, from 0% (top) to 50% (bottom). The social coupling strength κ is constant. The dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

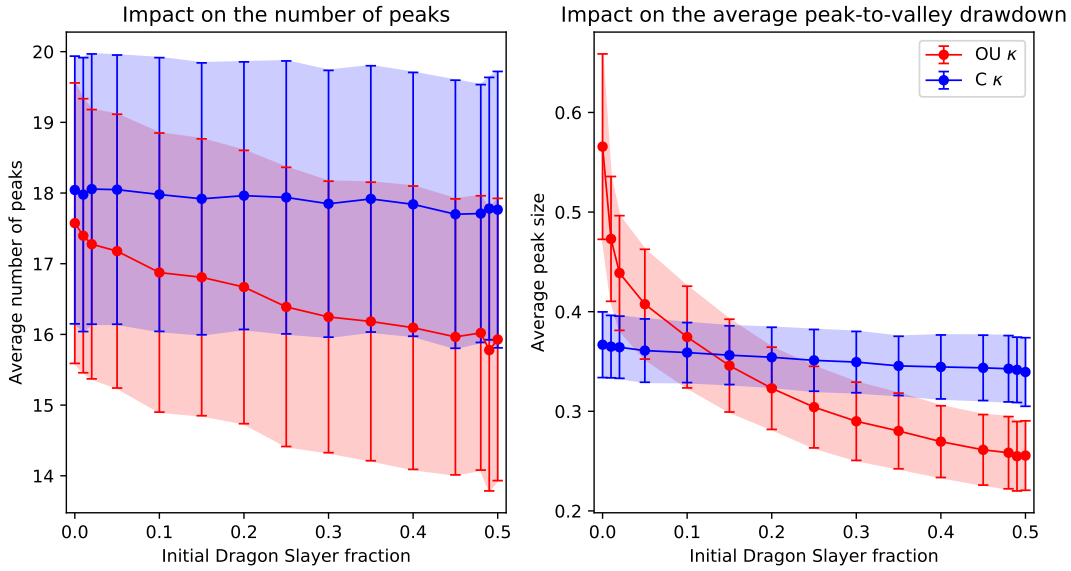

Figure 3.11: The impact of the dragon-slayer strategy on the market. The average number of peaks and the average peak-to-valley-drawdown is shown as a function of the initial dragon-slayer wealth fraction in the market. The averages are calculated over 1000 simulations with different random seeds. The error bars indicate one standard deviation. Each value is calculated between $t = 5000$ and $t = 17500$, which corresponds to approximately 50 years. The dragon-slayer parameters are $a = 0.95$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

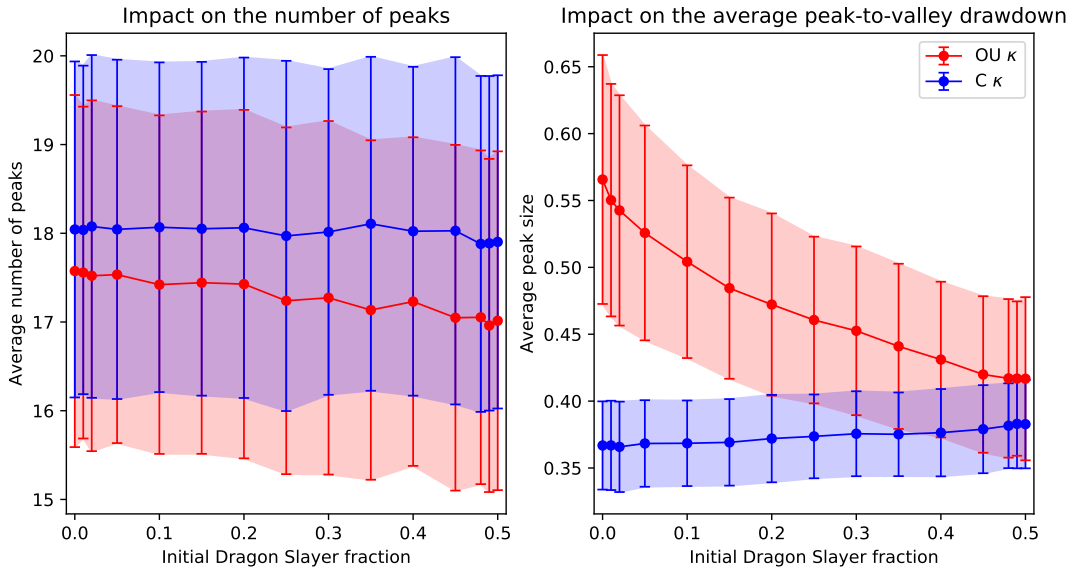

Figure 3.12: The impact of the dragon-slayer strategy on the market. The average number of peaks and the average peak-to-valley-drawdown is shown as a function of the initial dragon-slayer wealth fraction in the market. The averages are calculated over 1000 simulations with different random seeds. The error bars indicate one standard deviation. Each value is calculated between $t = 5000$ and $t = 17500$, which corresponds to approximately 50 years. The dragon-slayer parameters are $a = 0.99$, $w_y = 0.035$, $l_y = 0.004$, $\bar{r} = 0.00016$, and $s = 0.0005$.

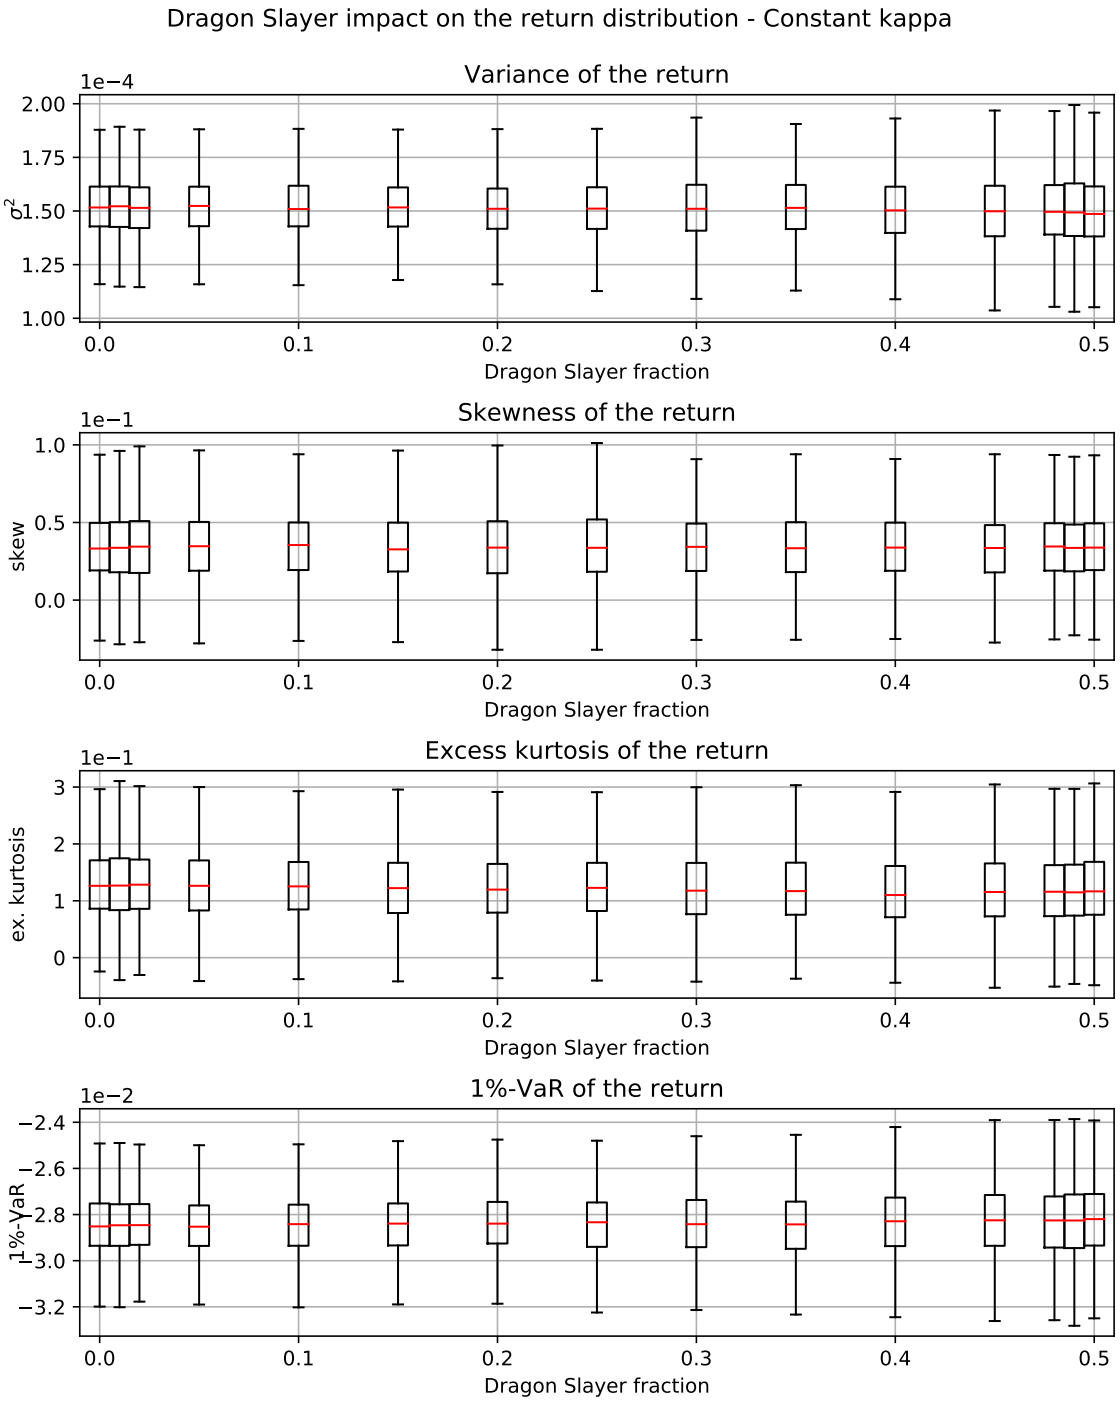

Figure 3.13: Boxplots of the dragon-slayer's impact on the return for constant κ, calculated over 1000 simulations with different random seeds for each dragon-slayer fraction. The red lines show the median, the box contains half of the values and the whiskers indicate the largest and lowest observed value within 1.5 times the interquartile range. The four panels show the variance, the skewness, the excess kurtosis, and the 1%- VaR of the return. For each simulation, the parameters are calculated over the time-interval $t = 5000$ to $= 17500$, the dragon-slayer parameters are $a = 0.98$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

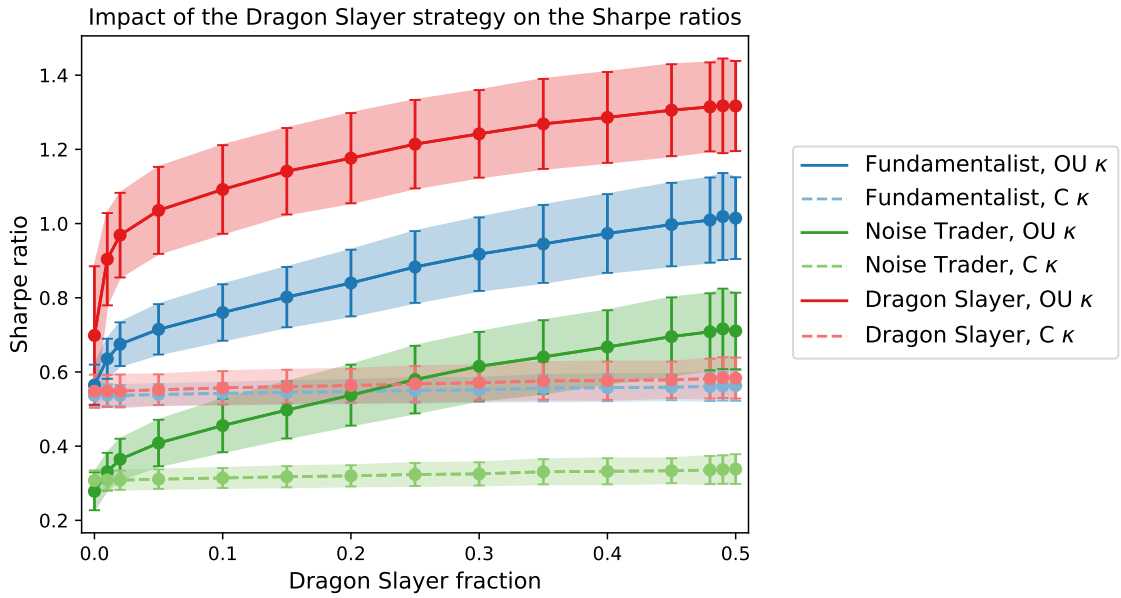

Figure 3.14: Average Sharpe ratios of the three trader types with an increasing fraction of the dragon-slayer's total wealth for OU κ (solid lines) and constant κ (dashed lines) calculated over 1000 realisations with different random seeds and over approximately 12500 time steps. The error bars indicate one standard deviation. The dragon-slayer parameters are $a = 0.95$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

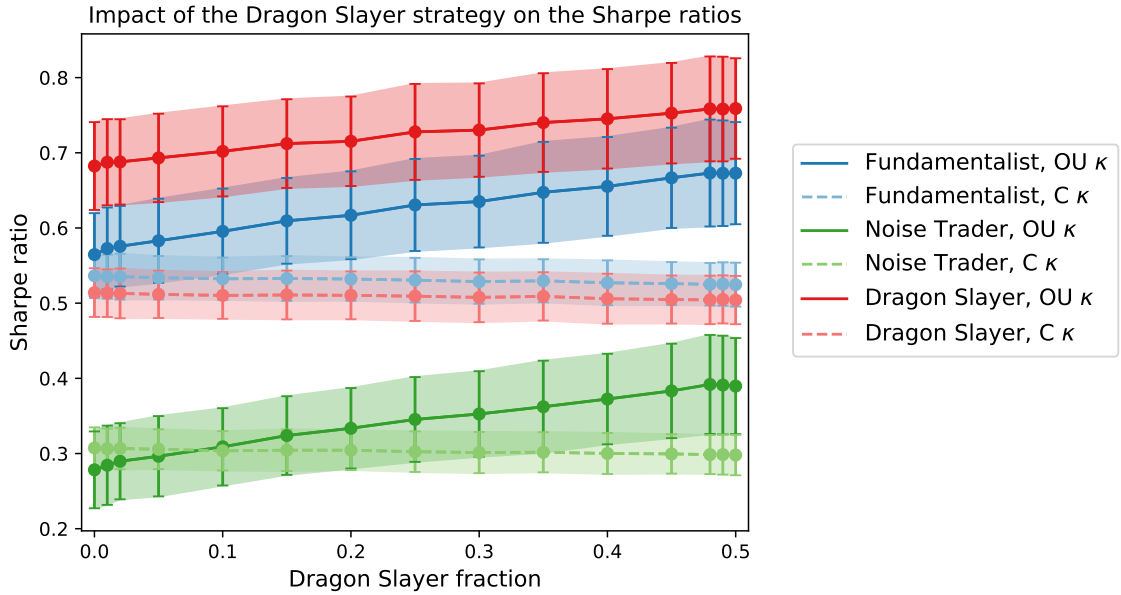

Figure 3.15: Average Sharpe ratios of the three trader types with an increasing fraction of the dragon-slayer's total wealth for OU κ (solid lines) and constant κ (dashed lines) calculated over 1000 realisations with different random seeds and over approximately 12500 time steps. The error bars indicate one standard deviation. The dragon-slayer parameters are $a = 0.99$, $w_y = 0.035$, $l_y = 0.008$, $\bar{r} = 0.00016$, and $s = 0.0005$.

Chapter 4

Agent-based models of multi-asset markets

4.1 Agent-based model generating stylized facts of Fixed Income Markets

This section develops an agent-based model (ABM) of a financial market with multiple assets belonging either to the Fixed Income or to the Equity asset class. The aim is to reproduce main stylized facts of Fixed Income markets with regards to the emerging dynamics of the yield curves. The research is rooted in the market model proposed by Kaizoji et al. (2015) introducing two types of traders: the rational and risk-averse investor referred to as the fundamentalist and the noise traders who invest under the influence of social imitation and price momentum. The investors involved in the present market model diversify their investments between a preferred stock equivalent to a perpetual bond and multiple bonds of selected maturities. Among these assets, a zero-coupon bond provides a constant rate of return, while the remaining coupon-paying bonds' prices are determined at each time step by establishing the equilibrium between the investors' demands and supplies. As a result, the market model provides an evolving yield curve impacted by the investments of the aggregated traders of each type. It is moreover capable to reproduce transient turbulent periods in the prices' time series and a humped term structure of volatility. Ultimately, the comparison of the emerging dynamics arising from different processes governing the evolution of the risk-free rate with those of the historical U.S. treasury market enables to distinguish the capacity of the setup implementing Vasicek's model of interest rates to reproduce the surface of autocorrelation of the individual bonds' yields' volatilities.

Based on the working paper Kopp, A., Westphal, R., & Sornette, D. (2020). Agent-based model generating stylized facts of Fixed Income Markets.

4.1.1 Introduction

This composition aims at unveiling the potential residing in the creation of agent-based models (ABMs) considering the Fixed Income asset class and involving fundamentalist and chartist traders. To the best of the authors' knowledge, it provides the first ABM of a financial market containing multiple assets belonging to the latter asset class. The inclination for this type of computational economic model lies in its capacity to produce transient market phenomena originated by the interactions of the individual agents. As pointed out in Sornette (2014), these models are freed from any stationary condition and can reproduce such periods of market unrest instigated by prices' rallies and their subsequent corrections. Examples of ABMs developed to describe stock markets can be

found in Kyle (1985), Black (1986) and Samanidou et al. (2007), among others.

The literature provides a significant amount of equilibrium models implemented to analyze the dynamics of yield curves and to forecast their evolution through the behaviour of the corresponding forward rates. Such analyses are e.g. found in Nelson and Siegel (1987), Cox et al. (2005), Vasicek (1977), Duffie and Kan (1996), Dai and Singleton (2002), Duffee (2002) and Heath et al. (1992). These models have however proven very limited with regards to their capacity to account for out-of equilibrium conditions corresponding to the ever-recurring extreme transient phenomena observed in reality. Moreover, they omit to consider the relationship between the Equity and Fixed Income asset classes. The market model developed below aims at laying a foundation towards the closure of this gap. The motivation is to propose a new approach to generate stylized facts of Fixed Income markets by a market model including the two previously mentioned asset classes. As exposed in Kaldor (1961), stylized facts represent "broad tendencies" understood as empirical truth observed in the markets of concern. The focus is therefore drawn on the emerging dynamics associated with the yields to maturity and evolving as a function of the traders' investments updated at each time step in conformity with the rules developed in this novel type of ABM.

The present research is based on the market model proposed by Kaizoji et al. (2015) introducing two types of agents investing in two assets. The risk-free asset provides a constant rate of return and the remaining asset's price is subjected to the market clearing process applied at each time step. It also pays dividends determined from a stochastic multiplicative process. The agents are either fundamentalists or noise traders and do not change their strategies in time. This is in contrast to earlier ABMs in which the traders switch between predictor (Arthur et al., 1996) or between chartist and fundamentalist strategies (De Grauwe et al., 1995; Brock & Hommes, 1997). The fundamentalist strategy is as described in Brock and Hommes (1998) and in Chiarella et al. (2009). The chartist strategy is derived from the developments of Lux and Marchesi (1999). In a nutshell, their investments are updated probabilistically at each time step under the influence of selected parameters. These parameters account for the influence of other agents and for the momenta in the prices' time series. On the aggregated level, the model proved to be capable to reproduce fat-tail distributions of returns, slow decaying autocorrelations of absolute returns, fast decaying autocorrelations of signed returns, volatility clustering and faster-than-exponentially growing prices. A contribution proposing a market model characterized by the interactions of fundamentalist and chartist traders through the price-vector of dimension higher than two worth mentioning is provided in Xu et al. (2014). Other models introducing heterogeneous agents investing in a multi-asset market can be found in Borghesi and Bouchaud (2007), Chiarella et al. (2007), Fedyk et al. (2013) and Eckrot et al. (2016).

The following section 4.1.2 introduces the market framework involving the two asset classes of concern. The two aggregated trader types investing in this framework are then presented. The market clearing process is subsequently defined before the unveiling of the parameter selection and of the initialization of the variables achieved in section 4.1.3. The latter section is concluded with the analyses of the time series generated in two characteristic simulations. The first is associated with a quiet market unfolding and the other with a turbulent one. Section 4.1.4 discloses the analyses comparing the dynamics emerging from the market model with those of the U.S. treasury market realized between November 1993 and June 2020. The section starts with the implementation of Vasicek's model of interest rates in the model to account for varying risk-free rates. It continues with the presentation of the term structures of volatility and surfaces of autocorrelations of the yields to maturity associated with each model setup and with the U.S. treasury market. Section 4.2.5 concludes.

4.1.2 Model setup

The market model created contains a stock and multiple bonds of selected maturities. It proposes *endogenous* price dynamics instigated by the traders' investments. One of the bonds represents a zero-coupon bond providing a constant add-on rate throughout the simulations. The agents belong

either to a the fundamentalist or to the noise trader type. The fundamentalist formulates expectations of the future returns as well as of the risk associated with each asset. The fundamentalist's investments are ruled by the need to maximise a constant relative risk aversion (CRRA) expected utility function. The noise traders' investments are subjected to social imitation, adherence to momentum trading and to the influence of a discrete time-varying herding propensity parameter. Their individual investments are updated probabilistically at each time step. Each trader type is represented by a unique agent expressing the aggregate excess demands for each asset at each time step. The market clearing condition imposing the equilibrium between the supplies and demands for the assets is finally at the origin of the *endogenous* evaluation of the asset prices.

The following developments start with the presentation of the market framework considering the Fixed Income asset class. The assets' and wealth dynamics are then exposed. The presentation of both aggregated traders ensues, starting with the fundamentalist and the associated generalized optimization problem. The threefold investment process governing the chartist's investments is subsequently defined before the presentation of the market clearing achieved by the Walrasian auctioneer, as defined by the eponym author in Walras (1954). The latter process provides the system of nonlinear equations ruling the updates of the asset prices at each time step. The selection of the parameter values initializing the market model follows, before the presentation of the time series generated by a configuration of the market model considering six coupon paying bonds and a unique preferred stock.

The Fixed Income market framework

The general notion of the passage of time is tackled through the following assumptions to ensure an ordered unfolding of the market model. The aim is to constrain the model to involve traders following only "aggressive" investment strategies. The selection of such traders is justified by their catalyzing effect on the market through the permanent rebalancing of their portfolio. This is in opposition with passive investors who "freeze" their investments to collect the regular coupon payments and benefit from the yields offered by their securities. This strategy is known as "riding the yield curve".

- Assumption 1: The times to maturity of the bonds evaluated at t are assumed to stay constant in $[t; t + \Delta t]$ for Δt sufficiently small.
- Assumption 2: The investors are supposed to replicate at t the exact same portfolio constructed at $t - 1$ before updating their investments. As soon as Δt is such that the earlier approximation does not hold anymore, each aggregated trader sells the corresponding bonds and buys new ones with the appropriate maturities. This replication has no impact on the bonds' prices as the sell orders of an agent replicating her portfolio are assumed to be compensated by the buy orders of another agent doing the same process. The sole impact on the prices hence results from the excess demands arising from the updated of the agents' investments.
- Assumption 3: The auctions of the coupon paying bonds take place at each time step and are ruled by the Walrasian market making process. They are not reopened, have no impact on the coupon rates and do not distinguish between competitive and non-competitive bidders.
- Assumption 4: The coupon rates are defined by an exogenous stochastic process reflecting the state of the economy. The dependency between the excess demands for the bonds and the coupon rates is ensured by the traders' investment processes.
- Assumption 5: The outstanding of coupon paying bonds remains constant in time.

Table 4.1 below presents the different types of assets involved in the market model and their corresponding maturities. The preferred stock is granted an infinite maturity and is hence considered

as a perpetual bond. The model involves $M = 8$ assets, the first one being the risk-free bond and the last one the preferred stock.

Asset type	Maturity (years)	Identifier
Zero-coupon bond	1	A^1
Coupon-paying bonds	$(k-1) \times 5 \quad \text{for} \quad k \in \{2, ..., M-1\}$	A^k
Preferred stock / perpetual bond	∞	A^M

Table 4.1: Asset types and maturities considered in the market model

The following developments introduce the yields to maturity (YTM) before the remaining parameters involved either in the traders' investment processes or in the analyses of their aggregated portfolio characteristics.

1. **Yields to maturity.** The market model is developed such that the bonds' yields are derived from the associated prices. In fact, the prices are updated at each time step through the market making process, rendering the updates of the corresponding yields necessarily contiguous. The YTMs are defined as internal rates of returns equalizing the present values of the future cash flows (coupons and principal payments) with the bonds' prices. They satisfy Equation (4.88) presented in Appendix 4.1.6. Their approximations are exposed below for each asset A^k with $k \in \{1, ..., M\}$ and where the no-arbitrage prices, times to maturity, principal and coupon payments are respectively represented by PV_k, N_k, FV_k and C_k:

$$\text{YTM}_k \approx \frac{\text{C}_k + \left(\frac{\text{FV}_k - \text{PV}_k}{N_k}\right)}{\frac{\text{FV}_k + \text{PV}_k}{2}}. \tag{4.1}$$

   As the time to maturity of A^M tends to infinity, the associated yield YTM_M is evaluated by:

$$\text{YTM}_M \approx \lim_{N_M \to \infty} \frac{\text{C}_M + \left(\frac{\text{FV}_M - \text{PV}_M}{N_M}\right)}{\frac{\text{FV}_M + \text{PV}_M}{2}} = \frac{\text{C}_M}{\frac{\text{FV}_M + \text{PV}_M}{2}}. \tag{4.2}$$

   The consideration of the preferred stock as a perpetual bond is further justified by considering Equation (4.90) provided in Appendix 4.1.6, which emphasizes the fact that the price of the associated zero-coupon bond vanishes as the time to maturity tends to infinity.

2. **Implied spot rates.** The implied spot rates (ISR) associated with each bond are introduced successively and correspond to the *holding period returns* realized over the entire life of each bond. They are particular forward rates covering the maturities of the bonds and thus provide approximations of the yields to maturity with the advantage to allow to omit the consideration of the rates at which the coupons are reinvested. These rates are evaluated in terms of the corresponding discount factors, as presented in Appendix 4.1.6.

3. **Yield durations and convexities.** The market model considers both modified (ModDur_t^k) and Macaulay (MacDur_t^k) durations associated with each asset. The Macaulay duration is defined as in Marrison (2002). Equations (4.100) to (4.102) of Appendix 4.1.6 provide the approximations of these durations as well as of the corresponding convexities (Conv_t^k).

4. **Portfolio yields, durations, convexities and dispersions.** The aggregated portfolio statistics are defined successively in Appendix 4.1.6, considering each portfolio as a unique entity equivalent to a bigger bond paying bundles of cash flows at each time step. Two approaches are applied to approximate the yields of the traders' portfolio: the market-value-weighted

(MV-weighted) and basis-point-value-weighted (BPV-weighted) approaches. The latter replaces the traders' wealth fractions by the proportion of each asset's BPV to the sum of the BPVs of the assets in which the portfolios are invested. The BPVs are expressed as the change in market value of each asset originated by a variation of its yield by a basis point, i.e. by 0.01%. The market-value-weighted average portfolio modified and Macaulay durations, convexities and dispersions are further defined in the appendix of concern.

This concludes the presentation of the parameters shaping the Fixed Income framework. The assets' returns and updates of the traders' wealth are presented hereafter.

Assets and wealth dynamics

The following developments present the returns associated with each type of assets, allowing to derive the update of the traders' wealth levels. A noteworthy mention is provided to tackle the consideration of the accrued interests in the returns of the coupon-paying bonds of finite maturities.

- **Assets' returns:**
  - **Risk-free asset.** As in Kaizoji et al. (2015), the risk-free asset provides a constant rate of return r_f. A^1 is further defined as a zero-coupon bond having a maturity of one year. The "day-count" convention selected defines a year to correspond to 250 time steps. The issuance price PV^1 and face value FV^1 of A^1 are related by r_f as follows:
  $$\text{FV}^1 = \text{PV}^1(1 + 250 \times r_f). \tag{4.3}$$
  - **Preferred stock.** The preferred stock A^M pays periodical dividends to its holders. Its "one time step" return r_t^M consists of the sum of the price return $\frac{P_t^M}{P_{t-1}^M} - 1$ with the return provided by the dividend payment d_t^M:
  $$r_t^M = \frac{P_t^M + d_t^M}{P_{t-1}^M} - 1 \tag{4.4}$$
  $$= \frac{P_t^M + d_{t-1}^M(1 + r_d + \sigma_d u_t^M)}{P_{t-1}^M}, \tag{4.5}$$
  where r_d is the long-term growth rate, σ_d the standard deviations of the multiplicative process and $u_t^M \sim \mathcal{N}(0,1)$ an i.i.d. random variable.
  - **Coupon-paying bonds of finite maturities.** Assuming the absence of taxation and the simultaneous payments of the coupons and dividends, the returns realized on each bond A^k for $k \in \{2, ..., M-1\}$ are expressed as the sum of the price and coupon returns as follows:
  $$r_t^k = \frac{P_t^k + d_t^k}{P_{t-1}^k} - 1. \tag{4.6}$$

In particular, the coupon payment d_t^k inherently corresponds to the interest accrued between the purchase of the bond A^k at $t-1$ and its sale at t. This arises from the replication process exposed earlier and requiring to buy another bond at t having the same maturity as the bond bought at $t-1$. The accrued interests are related to the annual coupon payments of C_t^k through: $d_t^k = C_t^k \times dt/250$, where the time increment is set to $dt = 1$. The multiplicative process governing the determination of the accrued coupon payments is further expressed as:

$$d_t^k = (1 + \sigma^k \nu_t^k) d_{t-1}^k, \tag{4.7}$$

where σ^k and ν_t^k respectively represent the standard deviation of the process and the associated i.i.d. random variable $\sim \mathcal{N}(0,1)$.

- **Wealth dynamics:**

The wealth dynamics are expressed in terms of the wealth fractions of the corresponding traders and of the assets' returns described above. Let W_t^i represent the wealth level of trader $i \in \{f, c\}$ at time step t. Also, let the number of shares $z_t^{i,k}$ of asset k held by the corresponding trader at t be expressed in terms of the wealth fraction $x_t^{i,k}$: $z_t^{i,k} = x_t^{i,k} W_t^i / P_t^k$. The wealth of the trader allocated in this asset is updated as:

$$\Delta W_{t-1\to t}^{i,k} = x_{t-1}^{i,k} W_{t-1}^i \frac{P_t^k}{P_{t-1}^k} - x_{t-1}^{i,k} W_{t-1}^i. \tag{4.8}$$

Summing over the 7 coupon- and dividend-paying assets, one obtains:

$$\Delta W_{t-1\to t}^i \|_{A^2,...,A^M} = W_{t-1}^i \sum_{k=2}^{8} x_{t-1}^{i,k} r_t^k. \tag{4.9}$$

The following trivial conditions are moreover applied on the agents' wealth fractions due to the prevention of both borrowing and short selling:

$$\begin{cases} 0 \leq x_t^{i,k} \leq 1 \quad \forall t, \\ 0 \leq \sum_{k=2}^{M} x_t^{i,k} \leq 1 \quad \forall t. \end{cases} \tag{4.10}$$

The fraction of wealth $x_t^{i,rf}$ invested in the risk-free asset is thus:

$$x_t^{i,rf} = 1 - \sum_{k=2}^{M} x_t^{i,k}. \tag{4.11}$$

By summing the wealth increments given by Equation (4.9) with the one associated with the risk-free asset, one can express the overall wealth dynamics between two consecutive time steps as:

$$\Delta W_{t-1\to t}^i = W_{t-1}^i \left[\sum_{k=2}^{M} x_{t-1}^{i,k} r_t^k + x_{t-1}^{i,rf} r_f \right], \tag{4.12}$$

which is equivalent to the formulation proposed hereafter, considering the expressions of the total returns exposed earlier:

$$W_t^i = W_{t-1}^i \left[\sum_{k=2}^{M} x_{t-1}^{i,k} \left(\frac{P_t^k + d_t^k}{P_{t-1}^k} \right) + \left(1 - \sum_{k=2}^{M} x_{t-1}^{i,k} \right) (1 + r_f) \right], \tag{4.13}$$

where the coupon and dividend payments associated with each asset A^k for $k \in \{2, ..., 8\}$ are:

$$d_t^k = \begin{cases} (1 + \sigma^k \nu_t^k) d_{t-1}^k & \text{for} \quad k = 2, ..., 7, \\ (1 + r_d + \sigma^k \nu_t) d_{t-1}^k & \text{for} \quad k = 8. \end{cases} \tag{4.14}$$

This concludes the presentation of the wealth updates and leads to the introduction of the fundamentalist's investment strategy.

Fundamentalist trader

The fundamentalist trader is defined in accordance with the descriptions provided in Chiarella et al. (2009), Kaizoji et al. (2015) and Xu et al. (2014). The trader is granted a constant relative risk-aversion (CRRA) utility function required to be maximized w.r.t. the wealth allocations. This objective is represented by a generalized myopic mean-variance optimization problem expressed in terms of the *ex ante* expectation of the utility function depending on the asset returns and wealth allocations introduced previously. Accounting for Equation (4.13), the optimization problem proposed in Kaizoji et al. (2015) is generalized to the multi-asset case hereafter.

- **Generalized optimization problem:**

$$\max_{\mathbf{x}_t^f} \mathrm{E}_t\big[U(W_{t+1}^f)\big] \quad \forall t \in [0,T], \tag{4.15}$$

where $U(W) \in \mathbb{R}$ is the CRRA utility function and $\mathbf{x}_t^f$ the vector containing the wealth fractions associated with each asset at t.

The solution procedure is provided in Appendix 4.1.6 and consists in solving the Hamilton-Jacobi-Bellman partial differential equation (PDE) associated with the generalized optimization problem of concern. The equivalent formulation of the optimization problem as a stochastic optimal control problem is proposed below. This is achieved by considering the expression of Equation (4.13) to define the wealth update in terms of a stochastic differential equation (SDE). To do so, the individual returns r_{t+1}^k of each asset A^k for $k \in \{2, ..., 8\}$ are expressed as in Xu et al. (2014) in terms of the expectations $\mathrm{E}_t\Big[r_{t+1}^k\Big]$ formulated at t of the returns to be realized at $t+1$, of the standard deviations σ_t^k and of the i.i.d. random variables ψ_t^k such that $\psi_t^k \sim \mathcal{N}(0,1)$:

$$r_{t+1}^k = \mathrm{E}_t\Big[r_{t+1}^k\Big] + \sigma_t^k \psi_t^k. \tag{4.16}$$

Considering Equations (4.13) and (4.16), the SDE accounting for the stochasticity of the wealth level of the fundamentalist is hence expressed as:

$$dW_t^f = W_{t-1}^f \left[\left(1 - \sum_{k=1}^{M-1} x_{t-1}^{f,k}\right) r_f + \sum_{k=1}^{M-1} x_{t-1}^{f,k} \mathrm{E}_{t-1}\big[r_t^k\big] \right] + W_{t-1}^f \sum_{k=1}^{M-1} x_{t-1}^{f,k} \sigma_{t-1}^k \psi_{t-1}^k, \tag{4.17}$$

which can be synthesized by the successive formulation proposed in Xu et al. (2014):

$$dW_t^f = \mu_t(W^f)dt + \sigma_t(W^f)\Psi_t, \tag{4.18}$$

where μ_t represents the drift coefficient, σ_t the standard deviation and $\Psi_t \sim \mathcal{N}(0,1)$ an i.i.d. random variable. Considering the unit vector $\mathbf{e} \in \mathbb{R}^{M-1}$ and the transpose operator $(.)^\top$, one can further express $\mu_t(W^f)$ and $\sigma_t(W)$ in both scalar and vector notations:

$$\begin{aligned} \mu_t(W^f) &= W_{t-1}^f \Big[\Big(1 - \textstyle\sum_{k=1}^{M-1} x_{t-1}^{f,k}\Big) r_f + \sum_{k=1}^{M-1} x_{t-1}^{f,k} \mathrm{E}_{t-1}\big[r_t^k\big] \Big] \\ &= W_{t-1}^f \Big[\Big(1 - {\mathbf{x}^f}_{t-1}^\top \mathbf{e}\Big) r_f + {\mathbf{x}^f}_{t-1}^\top \mathrm{E}_{t-1}\big[\mathbf{r}_t\big] \Big], \end{aligned} \tag{4.19}$$

$$\begin{aligned} \sigma_t(W^f) &= \left\{ {W_{t-1}^f}^2 \left[\textstyle\sum_{k=1}^{M-1} {x_{t-1}^{f,k}}^2 {\sigma_{t-1}^k}^2 + \sum_{k=1}^{M-1} \sum_{l=1}^{M-1} x_{t-1}^{f,k} x_{t-1}^{f,l} \sigma_{t-1}^{kl} \right] \right\}^{\frac{1}{2}} \\ &= W_{t-1}^f {\mathbf{x}^f}_{t-1}^\top \sigma_{t-1}, \end{aligned} \tag{4.20}$$

where it is required to distinguish $\sigma_t(W^f)$ from σ_t, the former accounting for the diffusion term of Equation (4.18) and the latter representing the vector containing the standard deviations of the individual assets' returns. Combining Equations (4.18), (4.19) and (4.20), the SDE of concern is given by:

$$dW_t^f = W_{t-1}^f \Big[\Big(1 - {\mathbf{x}^f}_{t-1}^\top \mathbf{e}\Big) r_f + {\mathbf{x}^f}_{t-1}^\top \mathrm{E}_{t-1}\big[\mathbf{r}_t\big] \Big] dt + W_{t-1}^f {\mathbf{x}^f}_{t-1}^\top \sigma_{t-1} \Psi_t, \tag{4.21}$$

where Ψ_t represents a one-dimensional Brownian motion and where the covariance matrix is expressed as $\Sigma = \sigma\sigma^T \in \mathbb{R}^{(M-1)\times(M-1)}$. Letting $\mu \in \mathbb{R}^{M-1}$ account for $\mathrm{E}_{t-1}\big[\mathbf{r}_t\big]$, one can finally write:

$$dW^f = W^f \Big[(1 - \mathbf{x}^\top \mathbf{e}) r_f + \mathbf{x}^\top \mu \Big] dt + W^f \mathbf{x}^\top \sigma \Psi. \tag{4.22}$$

The cost functional subjected to the previous wealth dynamics is subsequently introduced as follows:

$$\mathcal{J}(\mathbf{x}) = \mathrm{E}[U(W^f)] \quad \in \mathbb{R}. \tag{4.23}$$

The problem hence requires to find the optimal control $\mathbf{x} \in \mathbb{R}^{M-1}$ maximizing the cost functional subjected to the wealth defined by the stochastic system of Equation (4.22).

- **Generalized stochastic optimal control problem:**

$$\mathcal{J}(\mathbf{x}^*) = \max_{\mathbf{x}:[0,T]\mapsto\mathbb{R}^{M-1}} \mathcal{J}(\mathbf{x}). \tag{4.24}$$

The solution is given by the optimal state trajectory $\mathbf{x}^* : [0, T] \mapsto \mathbb{R}^{M-1}$. Introducing the risk aversion parameter $\gamma \neq 1$ ($\gamma > 0$ always holds), the cost-to-go function is given bys:

$$\mathcal{J}(\mathbf{x}, t) = \mathrm{E}\left[\frac{W^f(t)^{1-\gamma}}{1-\gamma}\right]. \tag{4.25}$$

Consequently, the stochastic optimal control problem is expressed as follows, considering the differential form of the standard Brownian motion $\Theta(t) \in \mathbb{R}$:

$$\max_{\mathbf{x}:[0,T]\mapsto\mathbb{R}^{M-1}} \mathrm{E}\left[\frac{W^f(t)^{1-\gamma}}{1-\gamma}\right] \tag{4.26}$$

subjected to:

$$\begin{aligned} &dW^f = \mu(W^f)dt + \sigma(W^f)d\Theta \\ &W^f(0) = W_0^f. \end{aligned} \tag{4.27}$$

According to Appendix 4.1.6, the solution of the generalized stochastic optimal control problem is:

$$x_k{}^* = \frac{1}{\gamma}\sum_{l=1}^{M-1} \Sigma_{kl}^{-1}(\mu_l - r_f) \quad \text{for} \quad k \in \{2, ..., M\}. \tag{4.28}$$

The remaining unknowns consist of the covariance matrix and of the vector containing the expectations of the future returns. The former matrix $\mathbf{\Sigma} = \sigma\sigma^\top$ is defined as diagonal and the individual variances remain constant. One obtains the following notations introducing the Kronecker delta δ_{kl}:

$$\begin{cases} \Sigma_{kl} = \delta_{kl}{\sigma_{kl}}^2, \\ \Sigma_{kl}^{-1} = \frac{\delta_{kl}}{{\sigma_{kl}}^2}. \end{cases} \tag{4.29}$$

The expectations of the future returns are further considering a constant and a variable term for each of the yields to maturity and implied spot rates. This enables to account for the trader's own opinion about the returns provided by each bond (constant term) while still considering the reality of the market (variable term). The constant terms consist of the initial values of the parameters. The variable terms considered are those evaluated at $t-1$ as a consequence of the replication process explained in section 4.1.2. Therefore, for each bond A^k where $k \in \{2, ..., M-1\}$, one expresses:

$$\begin{aligned} \mu_t^k = \omega_1 \frac{1}{2}\Big\{\Big[(1+\mathrm{YTM}_0^k)^{\frac{1}{250}} - 1\Big] + \Big[(1+\mathrm{ISR}_0^k)^{\frac{1}{250}} - 1\Big]\Big\} + \omega_2\Big[(1+\mathrm{YTM}_{t-1}^k)^{\frac{1}{250}} - 1\Big] \\ + \omega_3\Big[(1+\mathrm{ISR}_{t-1}^k)^{\frac{1}{250}} - 1\Big] + \frac{d_t^k}{P_t^k}, \end{aligned} \tag{4.30}$$

where ω_1, ω_2 and ω_3 are the weights respectively associated with the initial values of the yields to maturity and implied spot rates as well as with their evaluations at $t-1$. As explained in Appendix 4.1.6, the implied spot rates cannot be computed for the preferred stock. Moreover, due to the particular identity of the asset, the varying term associated with the yield to maturity is neglected. For $k = M$, one hence obtains:

$$\mu_t^M = (1 + \text{YTM}_0^M)^{\frac{1}{250}} - 1 + \frac{d_t^M(1+r_d)}{P_t^M}. \tag{4.31}$$

The wealth fractions of the fundamentalist trader are ultimately expressed as:

$$x_t^{f,k} = \frac{1}{\gamma\sigma_k^2}(\mu_t^k - r_f) \quad \text{for} \quad k \in \{2,..,M\}, \tag{4.32}$$

with the boundary condition given by $\sum_{i=1}^{M} x_t^{f,i} = 1$.

The subsequent developments introduce the investment process of the chartist trader taking the other side of the trades.

Noise traders

As in Kaizoji et al. (2015), the individual chartist traders do not diversify their allocations but instead select an individual asset in which they invest their entire wealth. Their investments are updated probabilistically at each time step through an Ising-like set-up under the influence of the momentum in the assets' prices' time series and the opinion of others. At the aggregated level, the representative noise trader allocates her wealth proportionally to the number of individual traders invested in the corresponding assets. This assumption is assumed not to affect the long-term dynamics of the aggregated chartist's wealth.

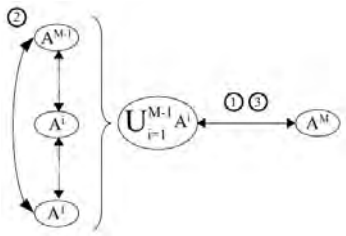

Figure 4.1: Representation of the market by the chartists: differentiation of the two pools of assets and corresponding steps of the investment process. The arrows guide the transit of the agents' wealth.

As shown in Figure 4.1, the noise trader identifies two mutually-exclusive pools of assets. The first one contains all the bonds including the risk-free one and is identified as $\mathcal{P}$. The second one is solely constituted of the preferred stock A^M. Built on this market representation, three distinct steps are defined below to compose the investment process of the aggregated chartist. Figure 4.1 further illustrates which particular pool is concerned by each of these steps. The first one tackles the rebalancing of the aggregated chartist's portfolio between the two asset classes considered. The second one concerns the update of the investments of the individual agents inside the bonds of $\mathcal{P}$. The third one finally reconciles the two previous updates simultaneously in order to derive the wealth fractions of the aggregated trader.

Step 1: Aggregated chartist's portfolio rebalancing between $\mathcal{P}$ and A^M:

The first step of the investment process of the aggregated chartist deals with the update of the allocations in either of the two groups of assets defined previously. It derives an estimate of the amount of individual traders transiting between the two asset classes. Both social imitation and adherence to a momentum-following strategy influence the outcome of the present step.

The number of individual chartists invested in A^M is N^M and the number of remaining agents is $N^{1\circlearrowleft M-1}$. The total fraction of wealth invested in A^M at t is hence:

$$x_t^c|_{A^M} = \frac{N_t^M}{N_t^M + N_t^{1\circlearrowleft M-1}}. \tag{4.33}$$

Following the definition of Kaizoji et al. (2015), the associated opinion index s_t^c is defined as:

$$s_t^{\text{step1}} = \frac{N_t^M - N_t^{1\circlearrowleft M-1}}{N_t^M + N_t^{1\circlearrowleft M-1}} = 2\, x_t^c|_{A^M} - 1 \quad \in [-1, 1]. \tag{4.34}$$

The parameter accounting for momentum is moreover defined as:

$$H_t^{\text{step1}} = H_t^M - H^{1\circlearrowleft M-1}, \tag{4.35}$$

introducing the momentum associated with the preferred stock:

$$H_t^M = \theta H_{t-1}^M + (1-\theta)\left(\frac{P_t^M}{P_{t-1}^M} - 1\right) \tag{4.36}$$

and the one related to $\mathcal{P}$ as:

$$H_t^{1\circlearrowleft M-1} = \theta H_{t-1}^{1\circlearrowleft M-1} + (1-\theta)\left[(1+\text{YTM}_t^{M-1})^{\frac{1}{250}} - (1+\text{YTM}_t^{1})^{\frac{1}{250}}\right]. \tag{4.37}$$

Equation (4.35) introduces the simple "rule of thumb" followed by the individual chartists to compare the price returns of the isolated asset with the "term-spreads" provided by $\mathcal{P}$. As shown in Equation (4.37), this spread is defined as the difference between the yields to maturity of the bonds situated at the extremities of the yield curve, fitted to the "one-time-step" returns of the preferred stock. This spread moreover represents a shorthand evaluation of the steepness of the yield curve giving insights about potential future returns obtainable in excess of those given by the risk-free bond. Considering the fact that $\text{YTM}_t^1 = \text{cst.}$, the higher the yields of the longer-term bond, the more the aggregated chartist is influenced to allocate her wealth in the Fixed Income asset class.

As exposed in Kaizoji et al. (2015), the herding propensity κ_t is introduced in order to represent the noise traders' susceptibility to herding and their propensity to adhere to a momentum following strategy. This parameter is either defined to be constant or to follow a stochastic mean-reverting process and incorporates the alternating regimes of pessimistic mood and exuberance observed in financial markets, as detailed in R. J. Shiller (2006) and in Sornette (2017).

By definition, the transition probabilities p_t^+ and p_t^- are associated with the following actions: an individual noise trader holding the preferred stock A^M at t shifts her investment to $\mathcal{P}$ with a probability of p_t^+ and another noise trader invested in $\mathcal{P}$ at the same moment transits to A^M with a probability of p_t^-. These probabilities are consequently expressed as:

$$p_t^{\pm}(s_t^{\text{step1}}, H_t^{\text{step1}}) = \frac{1}{2}p^{\pm}\big[1 \mp \kappa_t\big(s_t^{\text{step1}} + H_t^{\text{step1}}\big)\big], \tag{4.38}$$

where $p^+ = \text{cst.}$ and $p^- = \text{cst.}$ control the average holding time associated with each group of assets.setting $p^- > p^+$ is further found to imply the fact that an up- or downward increment of either of H_t^{step1} or s_t^{step1} is not engendering the same reaction of the agents. This is translated by a favour given to A^M to the detriment of $\mathcal{P}$.

Introducing the Bernoulli random variables $\xi_k(p)$ taking the value of 1 with a probability of p and 0 otherwise, one can express the subsequent master equations:

$$\begin{cases} N_t^M = \sum_{k=1}^{N_{t-1}^M}\left[1 - \xi_k(p_{t-1}^+)\right] + \sum_{k=1}^{N_{t-1}^{1\circlearrowleft M-1}} \xi_k(p_{t-1}^-), \\ N_t^{1\circlearrowleft M-1} = \sum_{k=1}^{N_{t-1}^{1\circlearrowleft M-1}}\left[1 - \xi_k(p_{t-1}^-)\right] + \sum_{k=1}^{N_{t-1}^M} \xi_k(p_{t-1}^+). \end{cases} \tag{4.39}$$

The net amount of agents transiting between both pools of assets at t is finally estimated by the absolute value of the variable $\Delta N_t^{+\to-}$ defined as:

$$\left|\Delta N_t^{+\to-}\right| = \left| \sum_{k=1}^{N_{t-1}^M} \xi_k(p_{t-1}^+) - \sum_{k=1}^{N_{t-1}^{1\circlearrowleft M-1}} \xi_k(p_{t-1}^-) \right|. \tag{4.40}$$

Step 2: Aggregated chartist's investments within $\mathcal{P}$:

The second step of the allocation process tackles the allocations of the chartists inside $\mathcal{P}$. Their updates are achieved independently for each pair of neighbouring assets (A^i, A^{i+1}) and (A^1, A^{M-1}). Again, both social imitation and adherence to a momentum-following strategy influence the outcome of these updates. The opinion index associated with the pair (A^i, A^{i+1}) is defined below, fitting the variable to the momenta introduced subsequently and introducing N_t^i as the number of chartists invested in A^i at t:

$$s_t^{\text{step2}}\Big|_{i,i+1} = \frac{1}{2}\frac{N_t^{i+1} - N_t^i}{N_t^{i+1} + N_t^i}. \tag{4.41}$$

The momentum parameter associated with the same pair of assets is moreover expressed as:

$$H_t^{\text{step2}}\Big|_{i,i+1} = H_t^{i+1} - H_t^i, \tag{4.42}$$

introducing the momentum parameters associated with each asset. These are further expressed in terms of the sub-parameters accounting for the momenta observed in the yields and in the prices:

$$H_t^i = \alpha H_t^{i,\text{price}} + (1-\alpha) H_t^{i,\text{yield}}, \tag{4.43}$$

where $\alpha \in [0,1]$. In particular:

$$H_t^{i,\text{price}} = \begin{cases} \theta H_{t-1}^{i,\text{price}} + (1-\theta)\left(\frac{P_t^i}{P_{\text{auction}}^i} - 1\right) & \text{for} \quad i = 2, ..., M-1, \\ r_f & \text{for} \quad i = 1. \end{cases} \tag{4.44}$$

$$H_t^{i,\text{yield}} = \theta H_{t-1}^{i,\text{yield}} + (1-\theta)\Big[(1+\text{YTM}_{t-1}^i)^{\frac{1}{250}} - (1+\text{YTM}_t^i)^{\frac{1}{250}}\Big], \tag{4.45}$$

where the yields are fitted to the "one time step" periods over which the price returns are determined. Equation (4.45) is moreover introducing the consideration of decreasing yields by the chartists. This enforces the assessment of positive realized price returns and contrasts with the previous step in which the agents valued higher yields implying higher potential price returns when looking at the overall market. Once invested inside the Fixed Income asset class, these agents naturally favor decreasing yields corresponding to positive realized price returns.

In accordance with the formulation of Equation (4.38), the transition probabilities associated with each pair of consecutive assets are defined as:

$$p_t^{i,i+1\pm} = \frac{p^\pm}{2}\Big[1 \mp \kappa_t\Big(s_t^{\text{step2}}\Big|_{i,i+1} + H_t^{\text{step2}}\Big|_{i,i+1}\Big)\Big], \tag{4.46}$$

where $p_t^{i,i+1^+}$ quantifies the probability to invest in A^i when starting from A^{i+1}. Accordingly, $p_t^{i,i+1^-}$ accounts for the probability to invest in A^{i+1} when starting from A^i.

The amounts of agents transiting between two neighbouring assets in either direction are hence:

$$\begin{cases} \Delta N_t^{i+1\to i} = \sum_{k=1}^{N_{t-1}^{i+1}} \xi_k\big(p_{t-1}^{i,i+1^+}\big), \\ \Delta N_t^{i\to i+1} = \sum_{k=1}^{N_{t-1}^{i}} \xi_k\big(p_{t-1}^{i,i+1^-}\big). \end{cases} \tag{4.47}$$

Ultimately, the net amount of agents leaving or entering each bond i is expressed as:

$$\Delta N_t^{\text{step2}}\Big|_i = \Delta N_t^{i-1\to i} - \Delta N_t^{i\to i-1} + \Delta N_t^{i+1\to i} - \Delta N_t^{i\to i+1}. \tag{4.48}$$

Step 3: Execution of the update of the aggregated chartist's investments:

The last step executes the update of the investments on the aggregated level. The following opinion index is first introduced for each of the $M-1$ bonds of fixed maturities:

$$s_t^{\text{step3}}\Big|_i = \frac{(M-1)N_t^i - \sum_{k=1,k\neq i}^{M-1} N_t^k}{(M-1)N_t^i + \sum_{k=1,k\neq i}^{M-1} N_t^k}. \tag{4.49}$$

The performance indicators accounting for the popularity of each bond and for the trend observed in their returns are defined as follows, considering the parameters H_t^i introduced in Equation (4.43):

$$\text{Perf}_t^i = s_t^{\text{step3}}\Big|_i + H_t^i. \tag{4.50}$$

Building on these parameters, the categorical probabilities are further defined below:

$$p_t^{\text{step3}}\Big|_i = \begin{cases} \frac{\text{Perf}_t^i - \text{Perf}_t^{\text{worst}}}{\sum_{k=1}^{M-1}(\text{Perf}_t^k - \text{Perf}_t^{\text{worst}})} & \text{if} \quad \Delta N_t^{+\to-} > 0, \\ \frac{\text{Perf}_t^{\text{best}} - \text{Perf}_t^i}{\sum_{k=1}^{M-1}(\text{Perf}_t^{\text{best}} - \text{Perf}_t^k)} & \text{if} \quad \Delta N_t^{+\to-} < 0, \end{cases} \tag{4.51}$$

where $\text{Perf}_t^{\text{worst}}$ and $\text{Perf}_t^{\text{best}}$ are respectively $\min_{i\in\{1,\dots,M-1\}} \text{Perf}_t^i$ and $\max_{i\in\{1,\dots,M-1\}} \text{Perf}_t^i$.

The weights ω_t^i associated with each bond are finally defined as:

$$\omega_t^i = \frac{N_t^i \; p_t^{\text{step3}}\Big|_i}{\sum_{k=1}^{M-1} N_t^k \; p_t^{\text{step3}}\Big|_k}. \tag{4.52}$$

As a matter of fact, ω_t^i corresponds to the share of the expected value of the Bernoulli trial associated with A^i and endowed with a probability of success of $p_t^{\text{step3}}\Big|_i$ over the sum of each of the $M-1$ expectations. These weights are introduced to adjust the dispersion of the agents inside $\mathcal{P}$ as can be inferred from the expression of Equation (4.53) presented below.

The number of agents transiting either to or from each bond of $\mathcal{P}$ is subsequently introduced as $\Delta N_t^{\text{step3}}\Big|_i$. By considering the result of the update achieved in the first step of the allocation process, one can hence write:

$$\Delta N_t^{\text{step3}}\Big|_i = \lfloor \omega_t^i \Delta N_t^{+\to-} \rfloor, \tag{4.53}$$

where $\lfloor . \rfloor$ represents the floor operator dropping the decimal part. Being granted this variable, one can execute the update of the individual investments considering the outcome of the two previous steps:

$$N_t^i = N_{t-1}^i + \Delta N_t^{\text{step2}}\Big|_i + \Delta N_t^{\text{step3}}\Big|_i \quad \geq 0 \quad \forall t. \tag{4.54}$$

Accordingly, the updated number of agents invested in the preferred stock is given by:

$$N_t^M = N_{t-1}^M - \sum_{k=1}^{M-1} \Delta N_t^{\text{step3}}\Big|_i. \tag{4.55}$$

As a result, the updated share of wealth of the aggregated chartist trader invested in each asset A^i for $i \in \{1, ..., M\}$ is ultimately expressed as:

$$x_t^c|_i = \frac{N_t^k}{\sum_{k=1}^{M} N_t^k}. \tag{4.56}$$

This concludes the presentation of the aggregated noise trader and leads to the formulation of the market making process.

Market clearing process

The present section tackles the market clearing process executed at each time step by the Walrasian auctioneer. First, the expressions of the excess demands formulated by both aggregated traders are developed. The system of nonlinear equations resulting from the satisfaction of the Walrasian equilibrium condition is provided subsequently.

First and foremost, the excess demand of the fundamentalist trader $\Delta D^{f,i}_{t-1\to t}$ associated with each asset A^i for $i \in \{2, ...M\}$ is expressed at each t as:

$$\begin{aligned}\Delta D^{f,i}_{t-1\to t} =& W^f_{t-1}\left(A^i_t + \frac{B^i_t}{P^i_t}\right)\left[\left(1 - \sum_{k=2}^{M} x^{f,k}_{t-1}\right)(1 + r_f)\right. \\ &\left. + \sum_{k=2}^{M} x^{f,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - x^{f,i}_{t-1} W^f_{t-1}\frac{P^i_t}{P^i_{t-1}},\end{aligned} \tag{4.57}$$

where A^i_t and B^i_t are obtained from Equation (4.32) presenting the expression of the wealth fractions associated with each asset:

$$A^i_t = \frac{\Omega^i_t - r_f}{\gamma \Sigma_{ii}}, \tag{4.58}$$

$$B^i_t = \begin{cases} \frac{d^i_t}{\gamma \Sigma_{ii}} & \text{for} \quad i = 2, ..., M-1, \\ \frac{d^M_t (1+r_d)}{\gamma \Sigma_{MM}} & \text{for} \quad i = M. \end{cases} \tag{4.59}$$

As a reminder and considering Equation (4.30), Ω^i_t for $i \in \{2, ...M-1\}$ is further expressed as:

$$\begin{aligned}\Omega^i_t =& \omega_1 \frac{1}{2}\left\{\left[(1 + \text{YTM}^i_0)^{\frac{1}{250}} - 1\right] + \left[(1 + \text{ISR}^i_0)^{\frac{1}{250}} - 1\right]\right\} \\ &+ \omega_2\left[(1 + \text{YTM}^i_{t-1})^{\frac{1}{250}} - 1\right] \\ &+ \omega_3\left[(1 + \text{ISR}^i_{t-1})^{\frac{1}{250}} - 1\right].\end{aligned} \tag{4.60}$$

And for $i = M$, Ω^M_t is:

$$\Omega^M_t = (1 + \text{YTM}^M_0)^{\frac{1}{250}} - 1. \tag{4.61}$$

On the other hand, the excess demand of the aggregated chartist trader associated with the same asset is expressed as:

$$\Delta D^{c,i}_{t-1\to t} = W^c_{t-1} x^{c,i}_t \left[\left(1 - \sum_{k=2}^{M} x^{c,k}_{t-1}\right)(1 + r_f) + \sum_{k=2}^{M} x^{c,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - x^{c,i}_{t-1} W^c_{t-1}\frac{P^i_t}{P^i_{t-1}}. \tag{4.62}$$

As has been exposed in the second step of the allocation process of the chartist trader presented in 4.1.2, the price returns of the individualized bonds are evaluated from the auction prices. These being further set to $P^i_{\text{auction}} = 1$, the price return of each bond realized at t is hence $P^i_t / P^i_{\text{auction}} - 1 = P^i_t - 1$. A^M being further exempted from any auction, its price returns are directly evaluated by P^M_t / P^M_{t-1}. As a consequence, the following rule is applied in Equation (4.62):

$$P^i_{t-1} = \begin{cases} 1 & \text{for} \quad i \in \{2, ..., M-2\}, \\ P^i_{t-1} & \text{for} \quad i = M. \end{cases} \tag{4.63}$$

For all $i \in \{2, ..., M\}$, the Walrasian equilibrium condition is finally expressed as:

$$\Delta D^{f,i}_{t-1\to t} + \Delta D^{c,i}_{t-1\to t} = 0. \tag{4.64}$$

The system of nonlinear equations involving the unknown prices remaining to be evaluated at each time step is finally obtained by inserting the expressions of Equations (4.57) and (4.62) in the latter condition. The resulting system is presented below and the developments leading to this formulation are provided in Appendix 4.1.6.

$$P_t^{i\,2}(\alpha_{ii}-\beta_i)+P_t^i(\zeta_i+\chi_{ii})+\sum_{k=2,k\neq i}^{M}P_t^kP_t^i\alpha_{ik}+\sum_{k=2,k\neq i}^{M}P_t^k\chi_{ik}+\lambda_i=0, \tag{4.65}$$

where $\alpha \in \mathbb{R}^{(M-1)\times(M-1)}$, $\beta \in \mathbb{R}^{M-1}$, $\zeta \in \mathbb{R}^{M-1}$, $\chi \in \mathbb{R}^{(M-1)\times(M-1)}$ and $\lambda \in \mathbb{R}^{M-1}$ are defined as:

$$\begin{cases} \alpha_{ij} = \frac{x_{t-1}^{f,j}W_{t-1}^fA_t^i+x_{t-1}^{c,j}W_{t-1}^cx_{t-1}^{c,i}}{P_{t-1}^j}, \\ \beta_i = \frac{x_{t-1}^{f,i}W_{t-1}^f+x_{t-1}^{c,i}W_{t-1}^c}{P_{t-1}^i}, \\ \zeta_i = W_{t-1}^fA_t^i\left(1-\sum_{k=2}^M x_{t-1}^{f,k}\right)(1+r_f)+W_{t-1}^cx_t^{c,i}\left(1-\sum_{k=2}^M x_{t-1}^{c,k}\right)(1+r_f) \\ \qquad\qquad +\sum_{k=2}^M\left[\frac{d_t^k}{P_{t-1}^k}\left(x_{t-1}^{f,k}W_{t-1}^fA_t^i+x_{t-1}^{c,k}W_{t-1}^cx_t^{c,i}\right)\right], \\ \chi_{ij} = \frac{x_{t-1}^{f,j}B_t^i}{P_{t-1}^j}, \\ \lambda_i = W_{t-1}^fB_t^i\left[\left(1-\sum_{k=2}^M x_{t-1}^{f,k}\right)(1+r_f)+\sum_{k=2}^M x_{t-1}^{f,k}\frac{d_t^k}{P_{t-1}^k}\right]. \end{cases} \tag{4.66}$$

Ultimately, this system of $M-1$ equations can be synthesized hereafter by the functions $f^i(t) \in \mathbb{R}$:

$$f^i(P_t^2,..,P_t^i,...,P_t^{M-1})=0 \quad \text{for} \quad i\in\{2,M-1\}. \tag{4.67}$$

An initial guess of the solution is provided by the previously realized prices $(P_{t-1}^2,...,P_{t-1}^i,...P_{t-1}^{M-1})$ and is further refined by the following process introducing $\mathbf{P}$, the vector notation of the prices of the coupon and dividend paying assets:

$$\mathbf{P}\rightarrow\mathbf{P'}=\mathbf{P}-\mathcal{J}^{-1}\mathbf{f}(\mathbf{P}), \tag{4.68}$$

where $\mathcal{J}$ is the Jacobian matrix defined by $\mathcal{J}_{ij}=\partial f^i/\partial P^j$. This concludes the presentation of the market clearing process.

4.1.3 Model dynamics

This section starts by presenting the initialization of the parameters introduced above and is concluded with the analyses of the emerging dynamics observed in typical simulations of the market model.

The focus is first shed on the herding propensity parameter κ_t mentioned in 4.1.2. As explained in Kaizoji et al. (2015), κ_t represents the strength of the influence that the parameters quantifying the opinion of other chartists and the trends in asset prices exercise on the update of the chartists' investments. The consideration of this parameter is justified by the need to account for the non-stationarity of the environment in which the agents evolve. In the context of this research, this environment embodies the economic environment as well as other non-economic aspects. Two distinct processes are hence ruling the evolution of κ_t: a constant and a stochastic mean-reverting process, respectively representing a stationary and a non-stationary market environment. The discrete linear stochastic mean-reverting process accounting for the varying environment is further defined as:

$$\kappa_t=\kappa_{t-1}+\eta_k(\mu_k-\kappa_{t-1})+\sigma_k\nu_t, \tag{4.69}$$

where the initial condition is given by κ_0 and where ν_t is an i.i.d. discrete-time white noise process with a mean of 0 and a standard deviation of 1. The mean-reversion rate is further represented by η_k, the mean value by μ_k and the diffusion associated with the Wiener process ν_t by σ_k.

Statistics	**Instantaneous evaluation**	**Long-run evaluation**
$\mathrm{E}[\kappa_t]$	$\mu_k - (\mu_k - \kappa_0)e^{-\eta_k t}$	$\lim_{t\to\infty} \mathrm{E}[\kappa_t] = \mu_k$
$\mathrm{Var}[\kappa_t]$	$\frac{\sigma_k^2}{2\eta_k}\left(1 - e^{-2\eta_k t}\right)$	$\lim_{t\to\infty} \mathrm{Var}[\kappa_t] = \frac{\sigma_k^2}{2\eta_k}$

Table 4.2: Expectation and variance of the herding propensity κ_t.

The instantaneous and long-run evaluations of the expectation and variance of the herding propensity are successively given in Table (4.2) above. The process of concern is hence found to be stationary and normally distributed in the long-run, such that the following expression holds:

$$\kappa_t \sim \mathcal{N}\left(\mu_k, \frac{\sigma_k^2}{2\eta_k}\right). \tag{4.70}$$

The statistical estimation of the parameters of Equation 4.201 is finally achieved as in Kaizoji et al. (2015) thrpugh the following expressions:

$$\eta_k = \frac{1}{\Delta T} \log\left(\frac{0.2}{1-\mu_k}\right), \tag{4.71}$$

$$\sigma_k = 0.1\sqrt{2\eta_k}, \tag{4.72}$$

where ΔT is the time frame during which κ_t is set to revert when being in the super-critical regime identified by the bounds located two standard deviations away form the mean.

The initialization of the parameters introduced in the market model is subsequently exposed hereafter. The analyses of the time series generated in typical simulations ensue.

Parameter selection

The wealth fractions of the aggregated traders are first initialized below, starting with the fundamentalist trader. Given the initial wealth fraction $x_0^{f,M}$, the constant variance σ_M^2 and the initial expectation of the future return $\mu_0^M = \mathrm{E}_{r,t}^M + d_0(1+r_d)/P_0$ associated with A^M, the constant risk aversion parameter γ is defined *a priori* by the following expression:

$$\gamma = \frac{\mu_0^M - r_f}{x_0^{f,M}\sigma_M^2}. \tag{4.73}$$

The expressions of the initial wealth fractions associated with the coupon paying bonds A^i for $i \in \{2, ..., M-1\}$ are obtained accordingly:

$$x_0^{f,i} = \frac{\mu_0^i - r_f}{\gamma\sigma_i^2}, \tag{4.74}$$

by ensuring that the following condition is satisfied:

$$0 \leq \sum_{k=2}^{M} x_t^{f,i} \leq 1 \quad \text{for} \quad t = 0. \tag{4.75}$$

The expectations of the future returns μ_0^i associated with these assets are defined below and the constant values of the variances σ_i^2 are provided in Table (4.3).

On the other hand, the initial wealth fractions of the aggregated chartist are obtained from the value of $x_0^{c,M}$ through the following expression:

$$x_0^{c,i}|_{i\neq 1,M} = \frac{1 - x_0^{c,M}}{M-1}. \tag{4.76}$$

Considering the floor operator $\lfloor . \rfloor$ introduced earlier and the total amount of chartists N, the amount of individual agents initially invested in the coupon and dividend paying assets is hence given by:

$$N_0^i = \begin{cases} 1 - \sum_{k=2}^{M} N_0^k & \text{for} \quad i = 1, \\ \left\lfloor N \frac{1 - x_0^{c,M}}{M-1} \right\rfloor & \text{for} \quad i \in \{2, ..., M-1\}, \\ \left\lfloor N x_0^{c,i} \right\rfloor & \text{for} \quad i = M. \end{cases} \tag{4.77}$$

This further enables to initialize the opinion indices involved in the allocation process governing the aggregated chartist's investments:

$$s_0^{\text{step1}} = 1 - 2x_0^{c,M}, \tag{4.78}$$

$$s_0^{\text{step2}}|_{i,i+1} = \frac{1}{2} \frac{N_0^{i+1} - N_0^i}{N_0^{i+1} + N_0^i}, \tag{4.79}$$

$$s_0^{\text{step3}}|_i = \frac{(M-1)N_0^i - \sum_{k=1,k\neq i}^{M-1} N_0^k}{(M-1)N_0^i + \sum_{k=1,k\neq i}^{M-1} N_0^k}. \tag{4.80}$$

Table (4.3) above unveils the initialization of the constant parameters as well as of the variables involved in the market model. The values included in parentheses respectively correspond to each of the assets in terms of increasing order of maturities.

The annualized yields to maturity are successively initialized as follows:

$$\text{YTM}_0^i = \begin{cases} \frac{2\times 250\times r_f}{2+250\times r_f} & \text{for} \quad i = 1 \\ \frac{d_0^i \times 250 + \frac{\text{FV}^i - P_0^i}{N}}{\frac{P_0^i + \text{FV}^i}{2}} & \text{for} \quad i \in \{2, ..., M-1\}, \\ \frac{d_0^M \times 250}{P_0^M} & \text{for} \quad i = M. \end{cases} \tag{4.81}$$

One can hence observe that the initial yield curve is flat as each yield initially equals 2.5%, with the exception of the yield of the perpetual bond initialized as $\text{YTM}_0^M = 3\%$. This precise setup is justified by the need to analyze the impacts of the agents' investments on a neutral configuration of the yield curve. This further explains the selection of the face values of the corresponding assets.

Building on these developments, the evaluations of the initial artificial changes of the market values, modified and Macaulay durations, convexities, discount factors and implied spot rates of each coupon-paying bond are straightforward. The according values are presented in Table (4.4) of Appendix 4.1.6.

Finally, given the previous developments, the expectations of the future returns are initialized as:

$$\mu_0^i = \begin{cases} \frac{1}{2}\left\{ \left[(1 + \text{YTM}_0^i)^{\frac{1}{250}} - 1 \right] + \left[(1 + \text{ISR}_0^i)^{\frac{1}{250}} - 1 \right] \right\} + \frac{d_0^i}{P_0^i} & \text{for} \quad i \in \{2, ..., M-1\}, \\ \left[(1 + \text{YTM}_0^i)^{\frac{1}{250}} - 1 \right] + \frac{d_0^M}{P_0^M} & \text{for} \quad i = M. \end{cases} \tag{4.82}$$

This concludes the presentation of the initialization of the constant parameters and variables defined previously and leads to the subsequent analyses of the phenomena emerging from numerical simulations of the market model.

Entity	Parameter	Parameter value	Equation
Assets	Number of assets	$M = 8$	
	Risk-free rate	$r_f = 8 \times 10^{-5}$	(4.3)
	Face values	$\text{FV} = (1.06667, 1.14286, 1.23077,$ $1.33333, 1.45455, 1.60000, \text{n.a.})$	(4.1) (4.2)
	Dividend growth rate	$r_d = 1.2 \times 10^{-4}$	(4.14)
	Standard deviation of the dividend processes	$\sigma_d = 1.6 \times 10^{-5}$	(4.14)
	Initial prices	$\text{P}_0 = (1, 1, 1, 1, 1, 1, 1)$	(4.4) (4.6)
	Initial coupons or dividend	$\text{d}_0 = (5, 5, 5, 5, 5, 5, 12) \times 10^{-5}$	(4.14)
	Maturities	$N = (1, 5, 10, 15, 20, 25, 30, \infty)$	(4.1)
Fundamentalist trader	Initial fraction of wealth invested in A^M	$x_0^{f,M} = 0.3$	(4.73)
	Initial wealth	$W_0^f = 10^6$	(4.13)
	Standard deviations of the assets' returns	$\sigma^2 \approx (3, 3, 3, 3, 3, 3, 4) \times 10^{-5}$	(4.32)
	Weights involved in the expected returns	$(\omega_1, \omega_2, \omega_3) = (0.9, 0.05, 0.05)$	(4.30)
Noise traders	Initial fraction of wealth invested in A^M	$x_0^{c,M} = 0.3$	(4.76)
	Initial wealth	$W_0^c = 10^6$	(4.13)
	Number of chartists	$N_0^{1 \circlearrowleft M-1} + N_0^M = 1500$	(4.77)
	Initial momentum of A^M	$H_0^M = r_d$	(4.36)
	Initial term spread	$H_0^{1 \circlearrowleft M-1} = (1 + \text{YTM}_0^{M-1})^{\frac{1}{250}}$ $-(1 + \text{YTM}_0^1)^{\frac{1}{250}}$	(4.37)
	Memory parameter	$\theta = 0.95$	(4.36) (4.37) (4.44) (4.45)
	Initial price and yield momenta of the bonds of $\mathcal{P}$	$(H_0^{i,\text{price}}, H_0^{i,\text{yield}}) = (0.0, 0.0)$	(4.43)
	Weight associated with the latter momenta	$\alpha = 0.5$	(4.43)
	Constants involved in the transition probabilities	$(p^+, p^-) = (0.19995, 0.20026)$	(4.46) (4.38)
Herding propensity	Initial value	$\kappa_0 = \mu_k$	(4.201)
	Long-run average	$\mu_k = 0.98$	(4.201)
	Mean-reversion rate	$\eta_k = 0.05$	(4.201)
	Diffusion of the associated Wiener process	$\sigma_k \approx 3 \times 10^{-2}$	(4.201)

Table 4.3: Set of constant parameters and initialization of the variables involved in the market model.

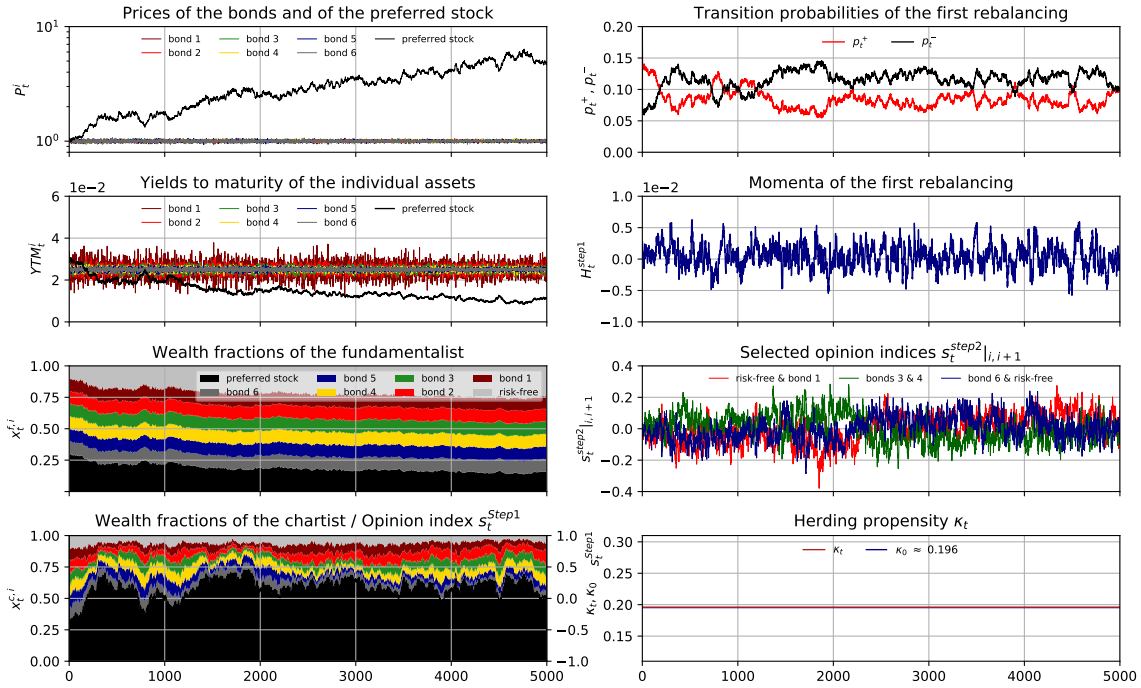

Figure 4.2: Time dependence of the variables associated with the market model. The simulation length is set to $T_{sim} = 5000$ time steps and the herding propensity remains constant throughout the simulation.

Time series analyses

The following section specifies the dynamics emerging in typical simulations. Figures 4.2 and 4.3 show representative time-series of the *endogenous* prices P_t^i, yields to maturity YTM_t^i, wealth fractions $x_t^{f,i}$ and $x_t^{c,i}$, transition probabilities $p_t^{\pm}(s_t^{\mathrm{step1}}, H_t^{\mathrm{step1}})$, momentum H_t^{step1}, and three representative opinion indices $s_t^{\mathrm{step2}}|_{i,i+1}$ created with a constant and a mean-reverting herding propensity. The herding propensity κ_t corresponds to the inverse temperature of the underlying Ising-structure of the noise traders' decision process. In the simulations involving a constant κ and where $\kappa < \kappa_c$, the noise traders are in the disordered regime. The stochastic mean reverting herding propensity has the same mean value $\kappa_\mu = 0.98 \times \kappa_c$ as the constant one, but fluctuates transiently above the critical value. This results in a polarization of the noise traders' opinions and can be observed through bubbles in the corresponding prices' time-series. These two different types of simulations are analyzed hereafter.

The time-series of the asset prices P_t^i show the divergence of the preferred stock's prices from those of the bonds. This divergence is governed by the growth rate of r_d. The subcritical herding behaviour caused by the constant nature of the herding propensity only results in subtle deviations from the average trajectory given by the previous rate. On the other hand, the simulation involving a varying herding propensity is found to generate periods of extreme transient phenomena during which the preferred stock's prices undergo turbulent patterns. Such periods are e.g. observed for $t \in [1300, 1700]$. Analyses of the super-exponential growth characterizing these bubbles are provided in Kaizoji et al. (2015), among others. Besides their prices, the bonds are characterized by their yields to maturity YTM_t^i. In both simulations, the yields fluctuate around their initial values with an amplitude related to the maturities of the bonds. These fluctuations reach higher values for bonds of low maturities. This constitutes an emerging phenomenon referred to as the humped "term structure of volatility". Moreover, according to the well known price-yield relationship, the yields of the perpetual bond reach their lowest values during the peak of the price bubbles, e.g. for $t \in [1400, 1700]$ in Figure 4.3. The reverting of the prices of A^M to levels directed by the growth rate of r_d as shown $t \in [3000, 3100]$ symmetrically drives the yields back to their initial value of 3%.

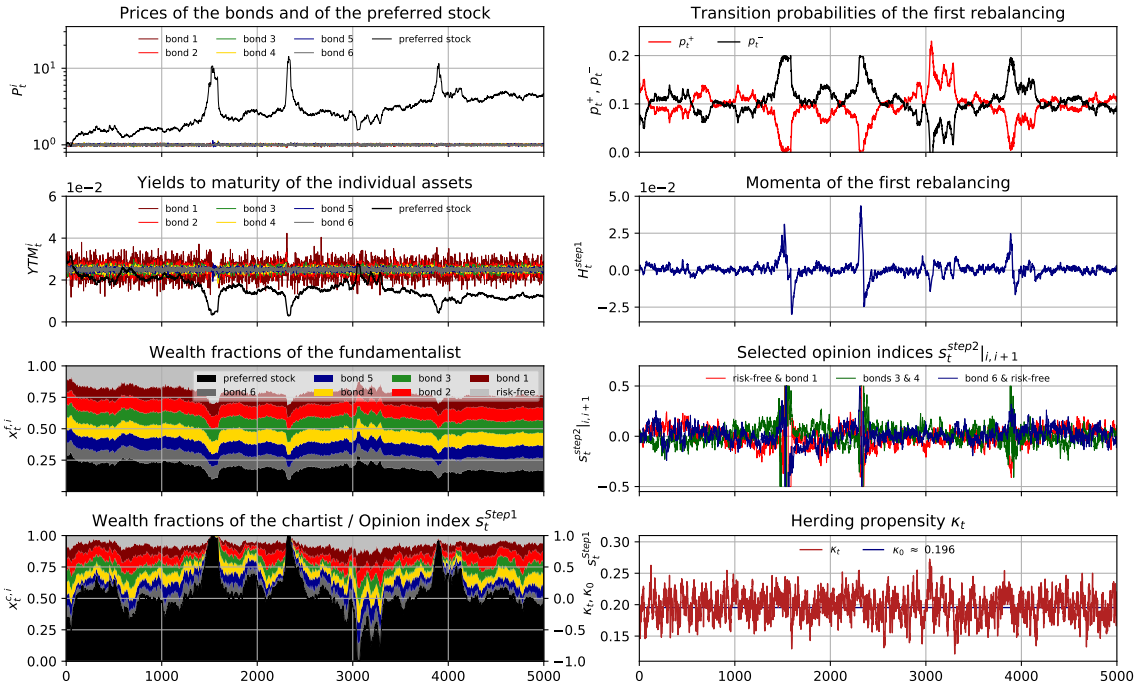

Figure 4.3: Time dependence of the variables associated with the market model. The simulation length is set to $T_{sim} = 5000$ time steps and the herding propensity follows a stochastic mean-reverting process.

The noise traders switch their investments between the different assets with a probability expressed in terms of the momentum of the prices of the preferred stock and of the opinion of others. The dynamics of the momentum H_t^M of the preferred stock, of the probabilities $p_t^{\pm}(s_t^{\text{step1}}, H_t^{\text{step1}})$ governing the rebalancing between the stock and the Fixed Income portfolio as well as of three opinion indices $s_t^{\text{step2}}|_{i,i+1}$ are detailed in the following. Overall, the momentum parameter H_t^{step1} evaluated in terms of the two sub-parameters H_t^M and H_t^{step1} replicates the main trends of the preferred stock's prices. This is due to the high amplitudes of the exponentially-weighted moving average of the returns of A^M in comparison with the stable spreads proposed by the yield curve during the simulations. The latter amplitudes are nevertheless found to be of approximately an order of magnitude lower than those associated with the opinion index s_t^{step1}. This explains the strong in influence given to the latter variable in the outcome of the Bernoulli trials. The transition probabilities fluctuate symmetrically around $(p^+ + p^-)/2$ and are bounded in $[0, \approx 0.15]$ for the simulation governed by a constant herding propensity. The smooth fluctuations observed in this simulation further explain the lack of abrupt changes in the preferred stock's price dynamics, as can be expected from the outcome of the Bernoulli trials involving these probabilities and generating no significant rebalancing of the aggregated chartist's portfolio between the two pools of assets. In contrast, the transition probabilities observed in the simulation involving the stochastic herding propensity increases sharply during the emergence or burst of the bubbles. The noise traders' opinion formation is illustrated by the following pairs of assets selected to avoid any redundancy in the analyses and to assess the integration of A^1 in $\mathcal{P}$: (A^1, A^2), (A^7, A^1) and (A^3, A^4). All three indices are found to fluctuate around a long-run average of zero with the same amplitudes, indicating an equilibrated dispersion of the agents in $\mathcal{P}$. The risk-free asset is hence not introducing any bias in the allocation process. This statement is emphasized by the fact that the trends in the fluctuations of the indices $s_t^{\text{step2}}|_{1,2}$ and $s_t^{\text{step2}}|_{7,1}$ replicate each other and are approximately symmetrically mirrored by those of $s_t^{\text{step2}}|_{3,4}$. The emphasis is furthermore directed towards the ascertainment of the abrupt saturation of the opinion indices observed during periods of herding towards the preferred stock proposed by the simulation involving a varying herding propensity. Considering the fact that a bad performing bond is penalized twice: the first time by the loss of agents transiting to the neighbouring bonds

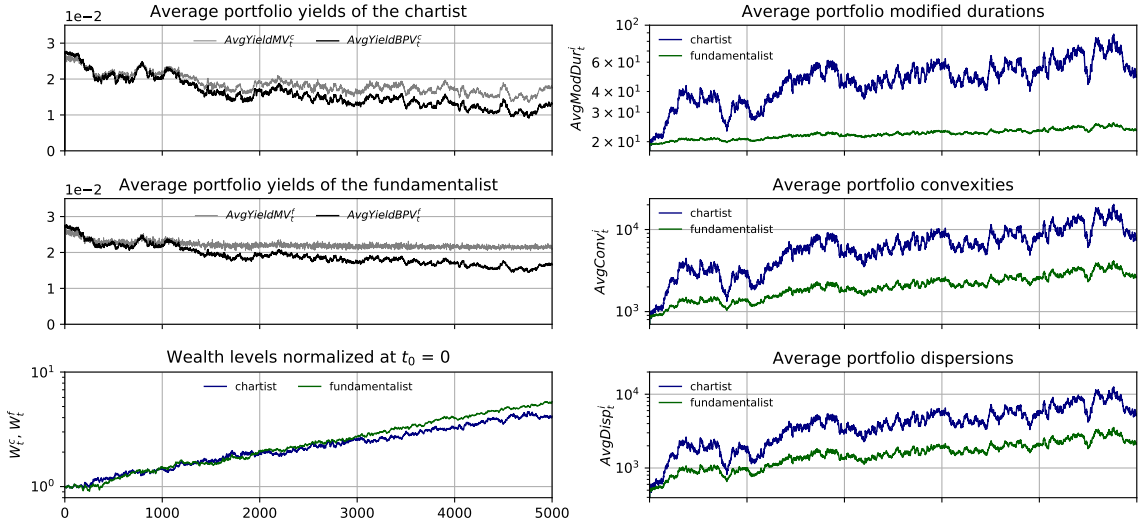

Figure 4.4: Time dependence of the variables associated with the market model. The simulation length is set to $T_{sim} = 5000$ time steps and the herding propensity remains constant throughout the simulation.

and the second time by the outflow of agents towards A^M, it appears that a very small amount of agents remain invested in the bonds during the culmination of the preferred stock's price bubbles. As a result, the opinion indices become very sensitive to the transition of agents as dictated by the outcome of the Bernoulli trials, explaining their abrupt variations at the corresponding time steps.

The transition probabilities define the aggregated noise trader's investment decisions as they determinate the update of the number of traders invested each asset. The resulting time-series of the wealth fraction x_t^c invested in each asset are analysed subsequently. A particular mention tackles the transient dynamics of $x_t^{c,M}$ occurring in the simulation involving the varying κ_t. For $t \in [2200, 2400]$, s_t^{step1} is found to be entirely polarized at the culmination of the corresponding price bubble. The lock-in effect occurring at this moment is straightforwardly accompanied by the following observations: $x_t^{c,M} = 1$ and $p_t^+(s_t^{\text{step1}}, H_t^{\text{step1}}) = 0$. In this case, the reaching of a threshold value by the herding propensity satisfies the statement proposed in Sornette (1994) and granting the responsibility of the "sweeping of an instability" of the latter parameter to instigate a self-reinforcing loop leading to a gradual feedback between H_t^M and s_t^{step1} provoking an outflow of agents from $\mathcal{P}$ to A^M. At the apogee of this phenomenon, no agent is spurred to transit to the Fixed Income portfolio. This situation is however unsustainable and the flux of agents between the two asset classes is reversed as soon as $p_t^+ > 0$, leading to the inevitable burst of the bubble. The market model is moreover found to be subjected to an initial phase of auto-regulation occurring at the beginning of each simulation. The quintessence of this observation lies in the increase of $x^{c,M}$ from 30% to a stable level of $\approx 50\%$ on the long-run. As a consequence of the resulting increase of the preferred stock's price level, $x^{f,M}$ is further reduced at the benefit of $x^{f,1}$. This is due to the counter-cyclical reaction of the fundamentalist trader to the evolution of the market. More generally, this "price taker" characteristic is found to be revealed in the time series of the corresponding wealth fractions mirroring the fluctuations of the prices of A^M.

Figures 4.4 and 4.5 provide a comparison of the fundamentalist's and aggregated chartist's portfolio statistics for simulations respectively involving a constant and a mean-reverting herding propensity. They show the evolution of the average portfolio yields AvgYieldMV_t^i and AvgYieldBPV_t^i, aggregate wealth levels W_t^i as well as approximated portfolio modified durations ApproxModDur_t^i, convexities ApproxConv_t^i and dispersions AvgDisp_t^i.

First, the average market-value-weighted and basis-point-value-weighted portfolio yields to maturity (AvgYieldMV_t^i and AvgYieldBPV_t^i) are analysed. The portfolio yields computed from the two different approaches provide different assessments of the attractivity of each investment strategy. The MV-weighted yields tend to replicate the patterns of the yields of the preferred stock less

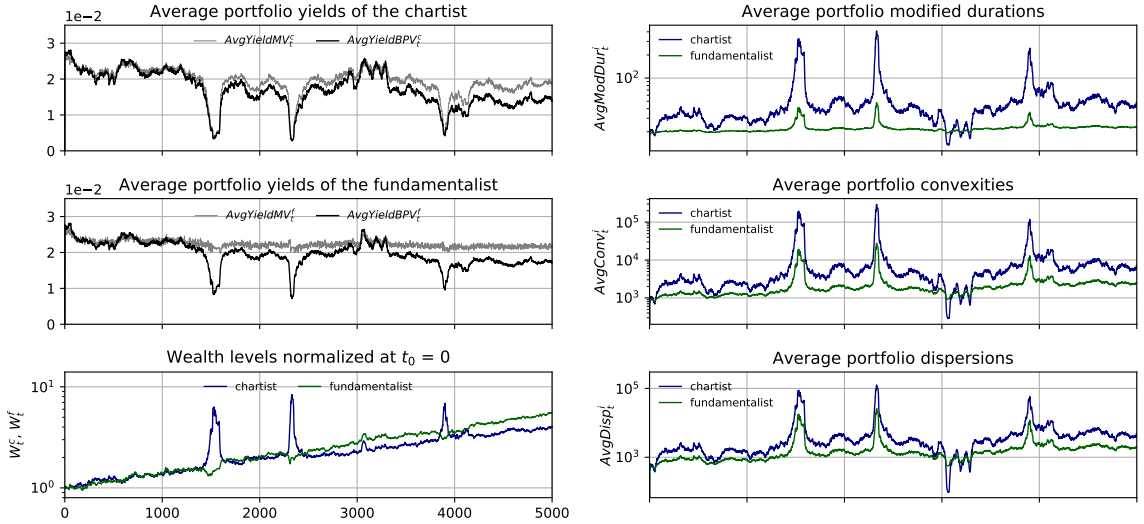

Figure 4.5: Time dependence of the variables associated with the market model. The simulation length is set to $T_{sim} = 5000$ time steps and the herding propensity follows a stochastic mean-reverting process throughout the simulation.

than is achieved by the BPV-weighted yields. This is striking for the fundamentalist trader who constrains her endowments to A^M during the simulations. The BPV-weighted yields contrast with the other yields in accordance to the expression of Equation (4.114). In particular, as the perpetual bond is granted a much higher duration than the other bonds, its prices are much more sensitive to variations of its yields, explaining why A^M takes a significantly greater share in the sum of the absolute price variations in the determination of the yields of concern.

The average portfolio modified durations AvgModDur_t^i tend to replicate the dynamics of the modified durations associated with the traders' predominant investment. The fundamentalist is hence found to have an almost stable portfolio modified duration over time with only few exceptions arising during the transient market unfolding. As a result, the latter variable is hence fluctuating between 20 and 30 years during periods of quietude. The portfolio modified duration of the fundamentalist is also more sensitive to increases of the perpetual bond's duration than to the according decrease of $x_t^{f,M}$. This is illustrated by the humps in the time series of the corresponding variable generated in the simulation involving the varying herding propensity. The present market model moreover provides a striking observation during periods of extreme transient phenomena and unveiling the limitation of the well known Fixed Income pricing formula expressed in Equation (4.103). In fact, the relationship implying, among others, that simultaneous increases of the duration and decreases of the yields lead to an increase in the corresponding asset's market value is found not to hold in turbulent periods for the portfolio of the fundamentalist trader. As one can observe during any of the price bubbles associated with A^M, the wealth of the fundamentalist decreases despite the fact that the portfolio modified duration increases and that the portfolio yield decreases. The convexity is besides found to increase, adding no interference to the previous statement. This sheds lights into the boundedness of the earlier mentioned pricing formula and explains why one might loose commonly-used landmarks when managing a fundamentalist portfolio during periods of market unrest. Consequently, surfing partly on the price bubble might help the fundamentalist to contain the losses caused by the massive exodus of the chartists from the other assets to the perpetual bond subjected to the bubbles of concern. Following the tendencies observed for the portfolio modified durations, the average portfolio convexities AvgConv_t^i and dispersions AvgDisp_t^i propose patterns strikingly sensitive to the dynamics of the perpetual bond's prices. One can again observe the fact that the rebalancing executed by the fundamentalist as a reaction to the price deviations of A^M is not sufficient to counteract the overall tendencies of the parameters of concern to follow the dynamics of the latter asset.

The comparison of the normalized wealth levels W_t^f of the fundamentalist and W_t^c of the ag-

gregated chartist reveals the fact that the fundamentalist's strategy outperforms the chartist's one irrespectively of the process governing the evolution of the herding propensity. The notion of opportunity costs is however crucial to evaluate the reversal of the latter trends on the short-term during turbulent periods. As a matter of fact, such costs appear to be incurred by the fundamentalist during the build up of the bubbles of the prices of A^M. Nevertheless, the subsequent bursts instigated by the flee of individual chartists towards $\mathcal{P}$ bringing the variable $x_t^{c,M}$ to anterior levels happen to cancel all the gains of the aggregated chartist, while the fundamentalist kept benefiting from the constant price return of the risk-free asset. As a result, the decreases of W^f observed during the build up of the price bubbles are adjusted by the still invariable price returns granted by A^1 left unimpaired by the variations of $x_t^{f,1}$ and $x_t^{c,1}$. This explains the slight increase of the difference of the wealth levels of the aggregated traders after each burst of a bubble.

4.1.4 Analyses

Application of the Vasicek model of interest rates

The previous developments introduced a constant risk-free rate provided by the zero-coupon bond of one year maturity. As a matter of fact, the rates provided by such risk-free assets are not constant in reality. The solution of the Vasicek model of interest rates developed by the eponym author in Vasicek (1977) is applied in the market model to provide a varying *exogenous* arbitrage-free risk-free rate. As a result, the comparison of the emerging dynamics associated with this new setup with those of the one considering $r_f = \text{cst}$. shall demonstrate the capacity of the new model to reproduce strikingly similar autocorrelations of the volatilities of the yields to maturity in comparison with those observed in the U.S. treasury market. The statistical estimation of the parameters involved in the expression of the varying risk-free rate is first achieved below, before the presentation of the emerging dynamics arising in typical simulations. The analyses of the autocorrelations are presented in 4.1.4.

The "one time step" rates are determined from following closed-form expression:

$$r_{f,t} = r_{f,t-1}e^{-\lambda} + \theta(1 - e^{-\lambda}) + \nu_t\sqrt{\frac{\sigma_r^2(1 - e^{-2\lambda})}{2\lambda}}, \tag{4.83}$$

where ν_t is an i.i.d. random variable such such that $\nu_t \sim \mathcal{N}(0, 1)$ and where λ, θ and σ_r are estimated in the same way as achieved for the herding propensity. The long-run average of the risk-free rate is equal to the constant value of the rate applied earlier, namely such that: $\theta = 8 \times 10^{-5}$. The mean-reversion strength is subsequently estimated according to:

$$\lambda = \frac{1}{\Delta T}\log_{10}\Big(\frac{r_{max} - \theta}{r_{critic} - \theta}\Big), \tag{4.84}$$

where $r_{max} = 1 \times 10^{-4}$ and $r_{critic} = 9.5 \times 10^{-5}$. The value attributed to θ is known and $\Delta T = 60$ time steps. One can hence evaluate the mean-reversion strength as $\lambda \approx 2 \times 10^{-3}$. Finally, σ_r is obtained as a function of the mean-reversion rate as follows:

$$\sigma_r = 2\lambda \mathrm{Var}_{r_{f,t}}, \tag{4.85}$$

where the variance of the risk-free rate is such that $\mathrm{Var}_{\mathrm{r_t}} = 2 \times 10^{-5}$, twice as low as the variance of the perpetual bond's returns. The initial value of the risk-free rate is finally set to $r_{f,0} = \theta$.

Figure 4.6 below presents the time series generated by a typical simulation involving a varying herding propensity as well as the stochastic process governing the evolution of $r_{f,t}$. The following developments tackle the inputs provided by the mean-reverting nature of $r_{f,t}$ on the emerging dynamics of the market model. The evolution of the risk-free rate is found to be reflected on $x^{f,1}$, while the preferred stock's prices show a clear symmetrical mirroring of its dynamics. The evolution of the wealth levels further provides an example of an indirect impact of the varying nature of the risk-free rate. The reason for the high sensitivity of the fundamentalist's wealth fractions to the

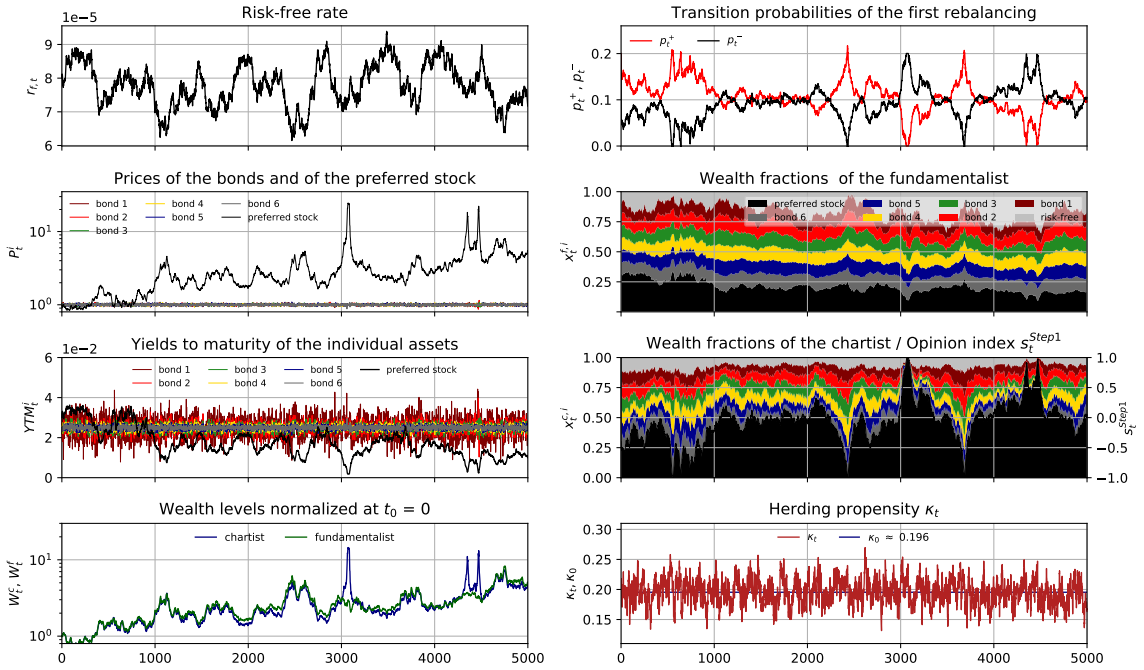

Figure 4.6: Time dependence of the variables associated with the market model. $T_{sim} = 5000$ time steps and the herding propensity follows a stochastic mean-reverting process. The evolution of the risk-free rate $r_{f,t}$ is governed by the Vasicek model of interest rates.

variations of $r_{f,t}$ is straightforward when considering Equations (4.31) and (4.32). In this regard, the fundamentalist's investments are hence found to have a major impact on the assets' prices, being mainly driven by the evolution of the risk-free rate and having a significant impact on the wealth levels of both aggregated traders. There is however no form of self-reinforcing feedback arising between the fundamentalist's investments and the corresponding impacts on the assets' prices, the latter agent having no incentive to follow series of positive returns and being predominantly reacting to the stochastic variations of the risk-free rate. This further explains why one cannot find clear trends in the evolution of the difference of both agent's normalized wealth levels, especially after the periods of extreme transient phenomena, where the earlier mention about the opportunity costs is not holding anymore as one is not guaranteed to obtain a stable return in the money market instrument during the culmination of the preferred stock's price bubbles.

Dynamics of the Fixed Income parameters

The following developments focus on the emerging dynamics of the yields to maturity, implied spot rates and discount factors. For the sake of clarity, the l.h.s. (respectively r.h.s.) of each illustration is associated with the simulation involving the constant herding propensity (respectively varying).

Time dependencies of the Fixed Income parameters

Figures 4.7, 4.8 and 4.9 present the surfaces generated from the time series associated with the yields to maturity, implied spot rates and discount factors obtained from the two typical simulations presented in 4.2.3. The yields of the perpetual bond are represented separately in Figure 4.7. Their distinction from the other yields is achieved for the sake of clarity as the corresponding asset is omitted in the computation of the yield curves. An ensemble of $T_{sim} = 5000$ yield curves is generated in each simulation, i.e. one yield curve per time step. The impacts of the aggregated agents' investments are straightforward considering the initial flat and neutral configuration of the yield curve. The increasing volatile character of the yields to maturity in terms of decreasing maturities is striking in both simulations and further reveals the aptitude of the present market model to produce a humped term structure of volatility. This paradigm is subsequently tackled in 4.1.4.

The implied spot rates defined according to Equation (4.98) provide accurate approximations

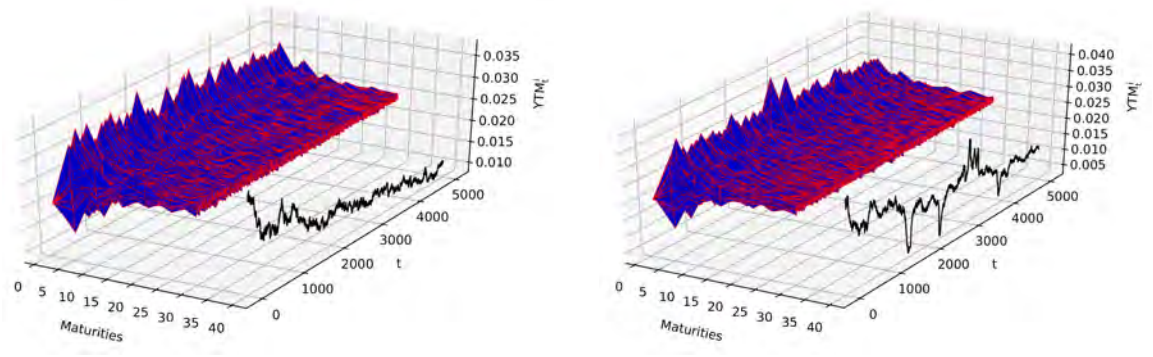

Figure 4.7: Surface generated from the time dependencies of the yields to maturity of the bonds of finite maturity. The yields of the perpetual bond are identified separately. The l.h.s. (respectively r.h.s.) is associated with the simulation involving a constant herding propensity (respectively mean-reverting).

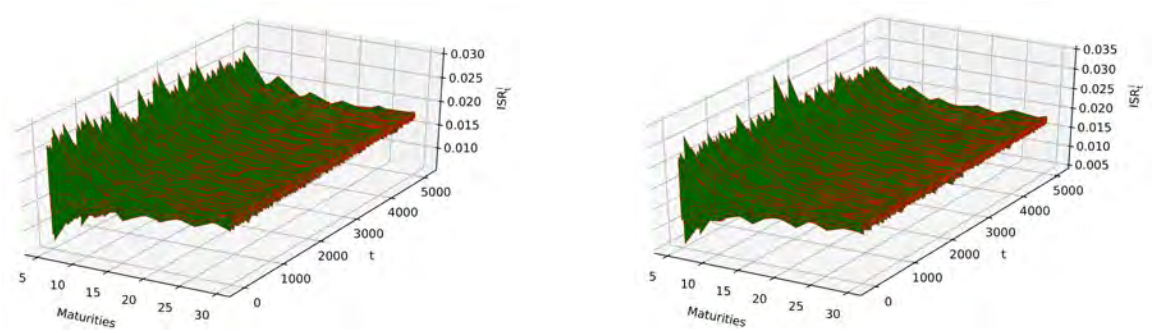

Figure 4.8: Surface of the implied spot rates of the bonds of finite maturity. The l.h.s. (respectively r.h.s.) is obtained from aconstant herding propensity (respectively mean-reverting).

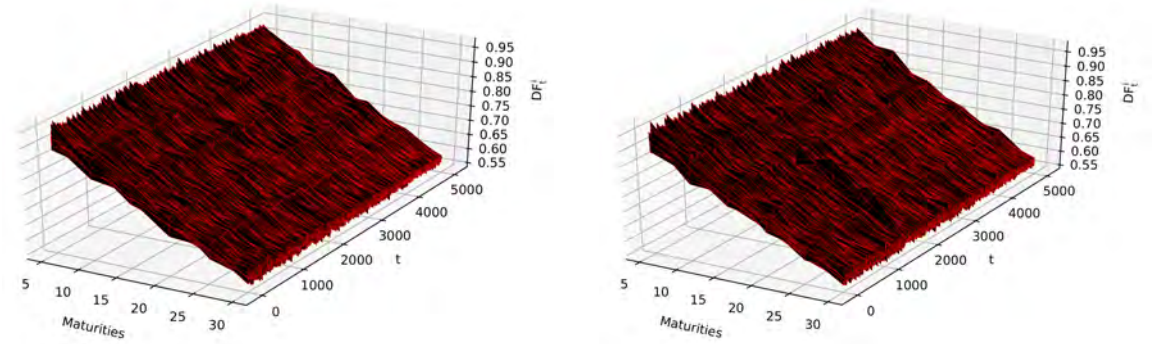

Figure 4.9: Surface of the discount factors of the bonds of finite maturity. The l.h.s. (respectively r.h.s.) is obtained from a constant herding propensity (respectively mean-reverting).

of the yields to maturity. As a result, the surfaces presented in Figure 4.8 are very similar to those of Figure 4.7. The spikes appearing at $t \in [2300, 2600]$ in the former very accurately replicate those emerging at the same period in the latter illustration. Figure 4.9 finally shows the surfaces generated by the linear interpolations of the discount factors associated with each of the bonds of fixed maturities. As one can observe, the disruptions caused by the aggregated chartist's investments on the assets' prices have non-negligible effects on these factors for bonds of high maturities.

This concludes the presentation of the evolution of the yields to maturity, implied spot rates and discount factors.

Volatilities of the yields to maturity

Figure 4.10 below presents the surfaces of obtained from the volatilities of the yields to maturity. The volatilities are computed over moving windows of 125 time steps in both simulations. They are evaluated at each time step from the following definition:

$$\sigma_{\mathrm{YTM}^i,t} = \sqrt{\left\langle (\mathrm{YTM}^i_k - \langle \mathrm{YTM}^i_k \rangle)^2 \right\rangle}\Bigg|_{k=t-\tau}^{t}, \tag{4.86}$$

where the sample average operator is represented by $\langle . \rangle$ and the selected moving window by $\tau = 125$.

Figure 4.11 shows the realized volatilities of the yields of U.S. Treasury bonds evaluated between the 18^{nth} of November 1993 and the 6^{th} of June 2020, considering moving windows of 250 business days. The selection of this moving window is justified by the need to amplify the trends of the fluctuations of the realized volatilities. Notwithstanding the fact that the volatilities of the yields of the U.S. Treasury bonds are subjected to the *exogenous* policies of the Federal Reserve, the comparison between those realized *endogenously* in the present model with those of the U.S. Treasury market enables to assess the capacity of the present market model to account for the main trends observed in reality.

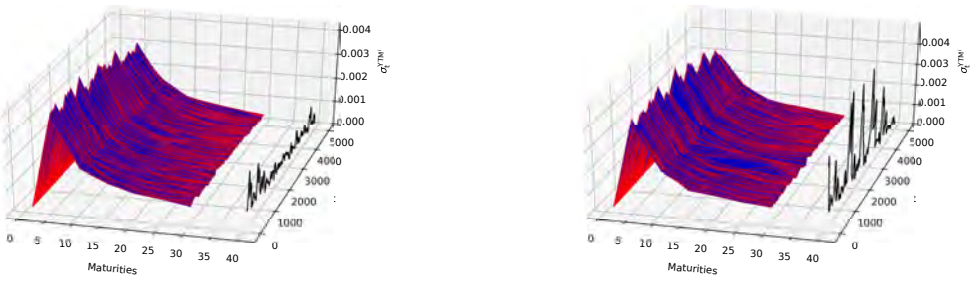

Figure 4.10: Surface of the volatilities of the yields to maturity. The l.h.s. (respectively r.h.s.) is obtained from a constant κ_t (respectively mean-reverting).

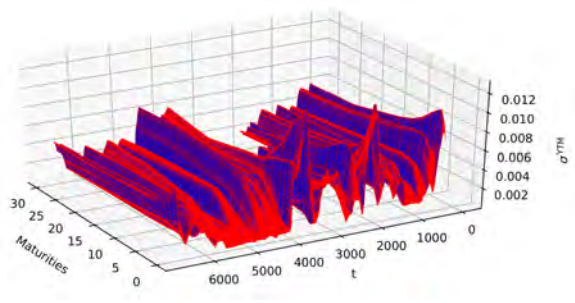

Figure 4.11: Surface of the volatilities of the yields to maturity of the corresponding U.S. treasury bonds realized between the 18^{nth} of November 1993 and the 6^{th} of June 2020.

The model produces a humped term structure of volatility in both simulations. The term structure of volatility is further decreasing when considering the bonds of maturities of higher than one year. Such phenomena are found episodically in the U.S. Treasury market, e.g. between 2004 and 2008, corresponding to $t \in [3200, 3950]$ in Figure 4.11. On the other hand, the *exogenous* imposition of low discount rates by the Federal Reserve are commonly known to be reflected on the yields of the U.S. Treasury bonds of low maturities, explaining the boundedness of their volatilities

during the according time frames ($t \in [4000, 5500]$). The increases of these bounds hence result in increasing volatilities of the yields of bonds of low maturities. As revealed in Cieslak and Povala (2016) and Bertocchi et al. (2005), the volatilities of the U.S. Treasury yield curve are found to be following such a hump-like shape on average, comforting the assertion of the capacity of the present market model to account for the main trends observed for the volatilities of the yields of U.S. Treasury bonds.

Autocorrelations of the volatilities of the yields to maturity

Figures 4.12 and 4.13 present the surfaces generated from the coefficients of autocorrelation of the volatilities of the yields to maturity computed for each bond, excluding the risk-free one and including the perpetual one. The former and the latter respectively consider the simulations involving a constant and a varying risk-free rate. The surfaces are generated by computing the average of the coefficients of autocorrelations from a sample of 300 simulations. The first 500 steps of each simulation are furthermore omitted in the computations due to the auto-regulation phenomenon commented in 4.2.3. The above-mentioned coefficients are determined for each simulation from Equation (4.87) below, considering the volatilities $\sigma_{\mathrm{YTM}^i,t}$ determined earlier and the lags $l \in [0, 500]$:

$$\begin{aligned} \mathrm{ACF}_l\left(\sigma_{\mathrm{YTM}^i}\right) &= \frac{\mathrm{Cov}\left[\sigma_{\mathrm{YTM}^i,t}\sigma_{\mathrm{YTM}^i,t-l}\right]}{\sqrt{\mathrm{Var}\left[\sigma_{\mathrm{YTM^i,t}}\right]\mathrm{Var}\left[\sigma_{\mathrm{YTM^i,t-l}}\right]}}, \\ &= \frac{\left\langle \sigma_{\mathrm{YTM}^i,t}\sigma_{\mathrm{YTM}^i,t-l}\right\rangle - \left\langle \sigma^2_{\mathrm{YTM}^i,t}\right\rangle}{\left\langle \sigma^2_{\mathrm{YTM}^i,t}\right\rangle - \left\langle \sigma_{\mathrm{YTM}^i,t}\right\rangle^2}. \end{aligned} \tag{4.87}$$

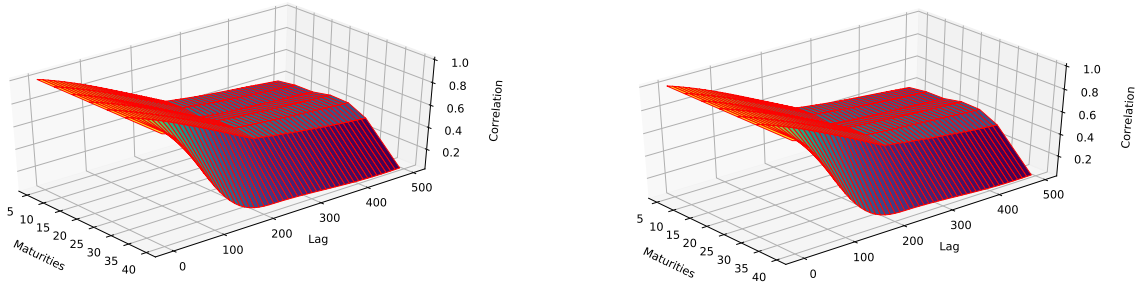

Figure 4.12: Surface of autocorrelation associated with the simulation involving a constant risk-free rate and a constant κ_t on the l.h.s. (respectively mean-reverting κ_t on the r.h.s.).

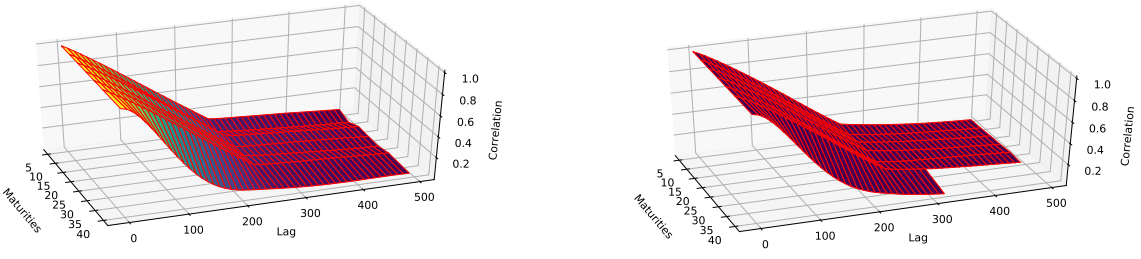

Figure 4.13: Surface of autocorrelation associated with the simulation involving a mean-reverting risk-free rate and a constant κ_t on the l.h.s. (respectively mean-reverting κ_t on the r.h.s.).

The comparison of Figures 4.12 and 4.13 reveals the fact that the surfaces associated with the market model involving a varying risk-free rate decay much faster than the others in terms of increasing lags. This emerging phenomenon is hence attributed to the stochastic character of the

evolution of the risk-free rate. There is moreover no difference on the aggregated level for different selections of the processes governing the evolution of the herding propensity.

The comparison of the surfaces presented in Figures 4.12 and 4.13 with the surface of autocorrelation associated with the U.S. treasury market provided in Figure 4.14 and computed from the earlier-mentioned volatilities further enables to distinguish the capacity of the simulations involving the varying risk-free rates to replicate the same emerging phenomena as they reproduce similar autocorrelation coefficients at the corresponding lags. More precisely, the latter simulations are capable to replicate the steeper gradients of the autocorrelation surfaces until the 250^{th} lag as well as the milder ones for higher lags directing the surfaces towards null autocorrelations at the 500^{th} lag. This concludes the presentation of the emerging dynamics of the market model consisting of the first ABM of a Fixed Income market, to the best of the authors' knowledge.

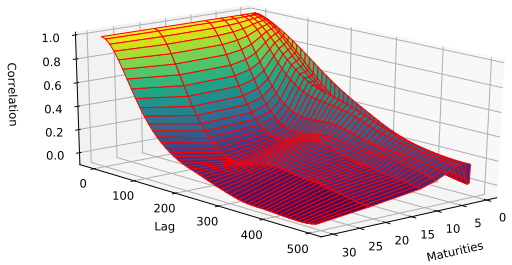

Figure 4.14: Surface of autocorrelation computed on the volatilities of the realized U.S. treasury yields between the 18^{nth} of November 1993 and the 6^{th} of June 2020.

4.1.5 Conclusion

The market model proposed in this research extended the model proposed in Kaizoji et al. (2015) by introducing multiple assets, among which several bonds of fixed maturities. The inclusion of the various Fixed Income parameters allowed to frame the investment processes of both aggregated traders to the peculiarities of the market in which they evolve. The market setup initialized by a flat yield curve demonstrated its capacity to reproduce several stylized facts of the U.S. Treasury market. The simulations involving a varying herding propensity showed that the well known bond-pricing formula involving the durations, convexities and shifts in yields to maturity does not hold on the fundamentalist's portfolio level during periods of market unrest. Subsequently, the implementation of Vasicek's model of interest rates to simulate the evolution of the risk-free rate was found to allow the linearly dependent fundamentalist's wealth fractions to replicate its dynamics and have a significant impact on the asset prices. The model endowed with such a varying rate ultimately prevented to distinguish a better performing investment strategy on the long run among the two applied in the present research. The analyses of the volatilities of the yields to maturity further enabled to assess the capacity of the market model to produce a humped term structure of volatility as has is observed on average in the U.S. Treasury market. Finally, the market model introducing the mean-reverting process governing the dynamics of the risk-free rate also proved superior to the one involving $r_f = \text{cst.}$ with regards to its capacity to reproduce the main characteristics of the surface of autocorrelation of the volatilities of the yields to maturity of the U.S. Treasury bonds for selected the time-frame.

4.1.6 Appendix

Parameters of the Fixed Income market model

- **Security level:**

The following developments introduce the parameters mentioned in 4.1.2. First and foremost, the famous relationship between an asset's no-arbitrage price PV, face value FV, coupon payments C, time to maturity N and yield to maturity YTM is expressed as:

$$\text{PV} = \frac{\text{C}}{\text{YTM}}\left[1 - \frac{1}{(1+\text{YTM})^N}\right] + \frac{\text{FV}}{(1+\text{YTM})^N}. \tag{4.88}$$

Representing the discount factor associated with the payments occurring at maturity by $\delta(N)$, the no-arbitrage price can be further expressed a as:

$$\text{PV} = \frac{\text{C}}{\text{YTM}}(1-\delta(N)) + \delta(N)\text{FV}, \tag{4.89}$$

exacerbating the expression of the bonds' dirty prices in terms of those of a perpetuity and of a zero coupon bond as follows:

$$\text{PV} = \text{PV}_{\text{perpetuity}}(1-\delta(N)) + \delta(N)\text{PV}_{\text{zero coupon bond}}. \tag{4.90}$$

Yields to maturity. Considering the fact that each asset A^i for $i \in \{2, ..., M\}$ pays a coupon or a dividend of d_t at each time step and consisting of the accrued interest associated with $dt = 1$, the corresponding yields to maturity are evaluated as:

$$\text{YTM}_t^i \approx \frac{250 \times d_t^i + \left(\frac{\text{FV}^i - \text{PV}_t^i}{N^i}\right)}{\frac{\text{FV}^i + \text{PV}_t^i}{2}}, \tag{4.91}$$

where the day-count convention defines a year as consisting of 250 trading days. The accrued interests AI_t^i of each asset A^i are further determined at each time step t by:

$$\text{AI}_t^i = \text{FV}^i \times \frac{\text{C}^i}{\text{FV}^i} \times \frac{dt}{T_{\text{year}}}, \tag{4.92}$$

where $T_{\text{year}} = 250$. Accordingly, the running yield of the preferred stock A^M is given by:

$$\text{YTM}_t^M \approx \frac{2 \times 250 \times d_t^M}{\text{FV}^M + \text{PV}_t^M}. \tag{4.93}$$

Discount factors. The discount factors provide a practical bootstrapping technique upon which the spot rates are determined. The methodology applied to determine the discount factors is detailed hereafter:

1. Start by determining the discount factor of the risk-free bond:

$$\text{DF}_t^1 = \frac{\text{PV}_t^1}{\text{FV}^1}. \tag{4.94}$$

Setting $\text{PV}^1 = 1$ and $\text{FV}^1 = 1 + 250 \times r_f$, one obtains:

$$\text{DF}_t^1 = \frac{1}{1 + 250 \times r_f} \quad = \text{cst.} \quad \forall t. \tag{4.95}$$

1. Determinate the remaining DF_t^i in an increasing order of maturities:

$$\text{DF}_t^i = \frac{\text{PV}_t^i - 250 \times d_t^i \sum_{k=1}^{i-1} \text{DF}_t^k}{250 \times d_t^i + \text{FV}^i}. \tag{4.96}$$

As one can observe from the expression of Equation (4.96), discount factors can only be assigned to bonds of fixed maturities. These factors are further related to the implied spot rates ISR_t^i by:

$$\text{DF}_t^i = \frac{1}{\left(1 + \frac{\text{ISR}_t^i}{\text{PER}}\right)^{N \times \text{PER}}}, \tag{4.97}$$

where the periodicity accounting for semiannual coupon payments is set to PER $= 2$.

Implied spot rates. The implied spot rates are accordingly isolated by the subsequent expression:

$$\text{ISR}_t^i = \left[\left(\frac{1}{\text{DF}_t^i}\right)^{\frac{1}{N^i \times \text{PER}}} - 1\right] \times \text{PER}. \tag{4.98}$$

Variations of market values. Assuming that the prices of each bond are continuous and twice differentiable w.r.t. to the yields to maturity, one can recall the famous expression:

$$\Delta\text{MV}_t^i \approx \left(\frac{\partial \text{MV}_t^i}{\partial \text{YTM}_t^i} d\text{YTM}_t^i\right) + \left(\frac{1}{2}\frac{\partial^2 \text{MV}_t^i}{\partial \text{YTM}_t^{i\,2}} (d\text{YTM}_t^i)^2\right), \tag{4.99}$$

from which the following parameters are introduced, considering only flat variations of the yield curve.

Durations and convexity.

The Macaulay duration is first expressed as:

$$\text{MacDur}_t^i = -\frac{\partial \text{MV}_t^i}{\partial \text{YTM}_t^i}\frac{1 + \text{YTM}_t^i}{\text{MV}_t^i}, \tag{4.100}$$

followed by the definition of the modified duration:

$$\text{ModDur}_t^i = -\frac{\partial \text{MV}_t^i}{\partial \text{YTM}_t^i}\frac{1}{\text{MV}_t^i} = \frac{\text{MacDur}_t^i}{1 + \text{YTM}_t^i}. \tag{4.101}$$

The convexity is finally formulated as:

$$\text{Conv}_t^i = \frac{\partial^2 \text{MV}_t^i}{\partial \text{YTM}_t^{i\,2}}\frac{1}{\text{MV}_t^i}, \tag{4.102}$$

allowing to express Equation (4.99) in terms of the above-introduced parameters:

$$\frac{d\text{MV}_t^i}{\text{MV}_t^i} \approx -(\text{ModDur}_t^i \times d\text{YTM}_t^i) + \frac{1}{2}(\text{Conv}_t^i \times (d\text{YTM}_t^i)^2). \tag{4.103}$$

These parameters are further evaluated from to the approximations provided in D. J. Smith (2014) and are first applied to individual securities before being associated with the aggregated traders' portfolios.

The modified duration is approximated by:

$$\text{ApproxModDur}_t^i = \frac{\Delta\text{MV}_t^i|_{\text{YTM}_t^{i\downarrow}} - \Delta\text{MV}_t^i|_{\text{YTM}_t^{i\uparrow}}}{2 \times \Delta\text{YTM}_t^i \times \text{MV}_t^i}, \tag{4.104}$$

where $\Delta\text{MV}_t^i|_{\text{YTM}_t^{i\downarrow(\uparrow)}}$ represents the variation of the bond's market value resulting from an artificial decrease (increase) of the yield by $-(+)0.001$.

The convexity is furthermore approximated by:

$$\text{ApproxConv}_t^i = \frac{\Delta\text{MV}_t^i|_{\text{YTM}_t^{i\downarrow}} + \Delta\text{MV}_t^i|_{\text{YTM}_t^{i\uparrow}} - 2 \times \text{MV}_t^i}{\Delta\text{YTM}_t^{i\,2} \times \text{MV}_t^i}. \tag{4.105}$$

In particular, considering the absence of arbitrage, the assets' market values and their corresponding variations arising from artificial shifts of the yields are determined according to:

$$\text{MV}_t^i = \frac{250 \times d_t^i}{\text{YTM}_t^i}\left(1 - \frac{1}{(1 + \text{YTM}_t^i)^N}\right) + \frac{\text{FV}^i}{(1 + \text{YTM}_t^i)^N}, \tag{4.106}$$

$$\mathrm{MV}_t^i|_{\mathrm{YTM}_t^{i\downarrow}} = \frac{250 \times d_t^i}{\mathrm{YTM}_t^i - 0.001}\left(1 - \frac{1}{(1+\mathrm{YTM}_t^i - 0.001)^N}\right) + \frac{\mathrm{FV}^i}{(1+\mathrm{YTM}_t^i - 0.001)^N}, \quad (4.107)$$

$$\mathrm{MV}_t^i|_{\mathrm{YTM}_t^{i\uparrow}} = \frac{250 \times d_t^i}{\mathrm{YTM}_t^i + 0.001}\left(1 - \frac{1}{(1+\mathrm{YTM}_t^i + 0.001)^N}\right) + \frac{\mathrm{FV}^i}{(1+\mathrm{YTM}_t^i + 0.001)^N}. \quad (4.108)$$

The Macaulay duration is finally approximated by:

$$\mathrm{ApproxMacDur}_t^i = \mathrm{ApproxModDur}_t^i(1 + \mathrm{YTM}_t^i). \quad (4.109)$$

- **Portfolio level:**

Durations and convexity. The market-value-weighted (MV-weighted) average modified portfolio durations are successively expressed as:

$$\begin{aligned} \mathrm{AvgModDur}_t^i &= \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{ModDur}_t^k \\ &\approx \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{ApproxModDur}_t^k, \end{aligned} \quad (4.110)$$

where $i \in \{f, c\}$ depicts the trader of concern. Similarly, the market-value-weighted average Macaulay portfolio durations are introduced as:

$$\begin{aligned} \mathrm{AvgMacDur}_t^i &= \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{MacDur}_t^k \\ &\approx \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{ApproxMacDur}_t^k. \end{aligned} \quad (4.111)$$

The average portfolios convexities are further determined as:

$$\begin{aligned} \mathrm{AvgConv}_t^i &= \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{Conv}_t^k \\ &\approx \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{ApproxConv}_t^k. \end{aligned} \quad (4.112)$$

Yields to maturity. The market-value-weighted average portfolio yields are evaluated by:

$$\mathrm{AvgYieldMV}_t^i = \sum_{k=1}^{M} x_t^{i,k} \times \mathrm{YTM}_t^k. \quad (4.113)$$

The basis point values (BPVs) of each asset A^k are determined by each trader i as:

$$\begin{aligned} \mathrm{BPV}_t^{i,k} &= \mathrm{ModDur}_t^k \times x_t^{i,k} \times W_t^i \times 0.0001 \\ &\approx \mathrm{ApproxModDur}_t^k \times x_t^{i,k} \times W_t^i \times 0.0001. \end{aligned} \quad (4.114)$$

Hence, the average basis point value-weighted (BPV-weighted) yields of each trader are evaluated as:

$$\mathrm{AvgYieldBPV}_t^i = \frac{\sum_{k=1}^{M} \mathrm{BPV}_t^{i,k} \times \mathrm{YTM}_t^k}{\sum_{k=1}^{M} \mathrm{BPV}_t^{i,k}}. \quad (4.115)$$

Dispersion. Ultimately, the average portfolios dispersions are introduced as:

$$\mathrm{AvgDisp}_t^i = \mathrm{AvgConv}_t^i(1 + \mathrm{AvgYieldBPV}_t^i)^2 - \mathrm{AvgMacDur}_t^{i\,2} - \mathrm{AvgMacDur}_t^i. \quad (4.116)$$

Solution of the generalized mean-variance optimization problem

The stochastic optimal control problem assigned to the fundamentalist trader is tackled in the following lines. The optimal feedback solution is derived after the formulation of the well known Hamilton-Jacobi-Bellman (HJB) equation associated with the problem of concern. The solution is furthermore given by the optimal control trajectory $\widehat{\mathbf{x}} : [0,T] \mapsto \mathbb{R}^{M-1}$ at the origin of the optimal state trajectory $\widehat{W^f} : [0,T] \mapsto \mathbb{R}$. The former is expressed as:

$$\widehat{\mathbf{x}}(t) = \tilde{\mathbf{x}}(\widehat{W^f}(t), \nabla_{W^f}\mathcal{J}(\widehat{W^f}(t),t),t) \quad \text{for} \quad t \in [0,T]. \tag{4.117}$$

In a nutshell, the control $\mathbf{x} = \mathbf{x}(W^f, \mathcal{J}_{W^f}, \mathcal{J}_{W^fW^f}, t)$ maximizing the r.h.s. of the HJB Equation (4.118) involving the cost-to-go function $\mathcal{J}(W^f,t)$ is first sought for. The resulting expression of the control law is then inserted back into the HJB equation before the resulting PDE expressed in terms of the cost-to-go function is solved. The expression of $\mathcal{J}(W^f,t)$ is finally replaced in the expression of the control parameter $\mathbf{x}$, unveiling the formulation of the optimal control law.

The HJB equation associated with the present framework is expressed as:

$$\mathcal{J}_t + \max_{\mathbf{x}\in\mathbb{R}^{M-1}} \left\{ \mathcal{J}_{W^f} W^f \Big[r_f + \mathbf{x}^\top(\mu - \mathbf{e}r_f) \Big] + \frac{1}{2}\mathcal{J}_{W^fW^f} {W^f}^2 \mathbf{x}^\top \sigma\sigma^\top \mathbf{x} \right\} = 0. \tag{4.118}$$

Recalling the fact that the covariance matrix associated with the SDE of the wealth update is expressed as $\mathbf{\Sigma} = \sigma\sigma^\top$, Equation (4.118) can be reduced to:

$$\frac{d}{d\mathbf{x}} \left\{ \mathcal{J}_{W^f} W^f \Big[r_f + \mathbf{x}^\top(\mu - \mathbf{e}r_f) \Big] + \frac{1}{2}\mathcal{J}_{W^fW^f} {W^f}^2 \mathbf{x}^\top \mathbf{\Sigma}\mathbf{x} \right\} = 0, \tag{4.119}$$

which further yields the following expression:

$$(\mu - \mathbf{e}r_f) W^f \mathcal{J}_{W^f} + \mathcal{J}_{W^fW^f} {W^f}^2 \mathbf{\Sigma}\mathbf{x} = 0. \tag{4.120}$$

Considering the latter Equation (4.120), one can obtain the formulation of the optimal control law expressed in terms of the first and second order derivatives of the cost-to-go function w.r.t. W^f:

$$\widehat{\mathbf{x}} = -\frac{\mathcal{J}_{W^f}}{W^f \mathcal{J}_{W^fW^f}} \mathbf{\Sigma}^{-1}(\mu - \mathbf{e}r_f). \tag{4.121}$$

Inserting the expression of Equation (4.121) into the HJB Equation (4.118), one obtains the following PDE expressed in terms of the cost-to-go function:

$$\mathcal{J}_t + W^f r_f \mathcal{J}_{W^f} - \frac{{\mathcal{J}_{W^f}}^2}{2\mathcal{J}_{W^fW^f}} (\mu - \mathbf{e}r_f)^\top \mathbf{\Sigma}^{-1}(\mu - \mathbf{e}r_f) = 0. \tag{4.122}$$

The following Ansatz introduces the scalar function $h(t) : \mathbb{R} \mapsto \mathbb{R}$ subjected to the terminal condition given by $h(T) = 1$:

$$\mathcal{J}(W^f,t) = h(t)\frac{{W^f}^{1-\gamma}}{1-\gamma}. \tag{4.123}$$

The relevant derivatives of the value function are consequently obtained as:

$$\begin{cases} \mathcal{J}_t(W^f,t) = \dot{h}(t)\frac{{W^f}^{1-\gamma}}{1-\gamma}, \\ \mathcal{J}_{W^f}(W^f,t) = h(t){W^f}^{-\gamma}, \\ \mathcal{J}_{W^fW^f}(W^f,t) = -\gamma h(t){W^f}^{-(\gamma+1)}, \end{cases} \tag{4.124}$$

allowing to express the solution of the optimization problem ruling the investments of the fundamentalist trader as:

$$\widehat{\mathbf{x}} = \frac{1}{\gamma}\mathbf{\Sigma}^{-1}(\mu - \mathbf{e}r_f) \tag{4.125}$$

This expression is solves the generalized wealth allocation problems faced by any trader subjected to the necessity to maximize an expected CRRA utility function w.r.t. her investments in $M-1$ coupon and dividend paying assets. The rest is allocated in the risk-free asset according to the boundary condition expressed in Equation (4.75).

Market clearing process of the extended market model

The developments provided below shed lights into the market clearing process achieved at each time step by the Walrasian auctioneer. As has been expressed in 4.1.2, Equation (4.64) is the condition at the origin of the updates of the assets' prices. The expressions of the excess demands for each asset A^i with $i \in \{2, ..., M\}$ respectively formulated by the fundamentalist and chartist traders are:

$$\Delta D^{f,i}_{t-1\to t} = W^f_{t-1}\left(A^i_t + \frac{B^i_t}{P^i_t}\right)\left[\left(1 - \sum_{k=2}^{M} x^{f,k}_{t-1}\right)(1+r_f) + \sum_{k=2}^{M} x^{f,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - x^{f,i}_{t-1} W^f_{t-1}\frac{P^i_t}{P^i_{t-1}}, \quad (4.126)$$

$$\Delta D^{c,i}_{t-1\to t} = W^c_{t-1} x^{c,i}_t\left[\left(1 - \sum_{k=2}^{M} x^{c,k}_{t-1}\right)(1+r_f) + \sum_{k=2}^{M} x^{c,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - x^{c,i}_{t-1} W^c_{t-1}\frac{P^i_t}{P^i_{t-1}}, \quad (4.127)$$

where both A^i_t and B^i_t are defined according to Equations (4.58) and (4.59). For the sake of clarity, let $\left(1-\sum_{k=2}^{M} x^{f,k}_{t-1}\right)(1+r_f)$ be represented by Π_{t-1} and Φ_{t-1} account for $\left(1-\sum_{k=2}^{M} x^{c,k}_{t-1}\right)(1+r_f)$. The condition of the Walrasian auctioneer is refined hereafter, considering Equations (4.126) and (4.127):

$$W^f_{t-1}\left(A^i_t + \frac{B^i_t}{P^i_t}\right)\left[\Pi_{t-1} + \sum_{k=2}^{M} x^{f,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - x^{f,i}_{t-1} W^f_{t-1}\frac{P^i_t}{P^i_{t-1}} + W^c_{t-1} x^{c,i}_t\left[\Phi_{t-1} + \sum_{k=2}^{M} x^{c,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - x^{c,i}_{t-1} W^c_{t-1}\frac{P^i_t}{P^i_{t-1}} = 0. \quad (4.128)$$

One can further factorize the terms accounting for the ratios of the prices evaluated at t and $t-1$ by isolating P_t as follows:

$$W^f_{t-1}\left(A^i_t + \frac{B^i_t}{P^i_t}\right)\left[\Pi_{t-1} + \sum_{k=2}^{M} x^{f,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] + W^c_{t-1} x^{c,i}_t\left[\Phi_{t-1} + \sum_{k=2}^{M} x^{c,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - P^i_t\left(\frac{x^{c,i}_{t-1} W^c_{t-1} + x^{f,i}_{t-1} W^f_{t-1}}{P^i_{t-1}}\right) = 0. \quad (4.129)$$

Let the parameter β_i associated with asset A^i be introduced as:

$$\beta_i = \frac{x^{c,i}_{t-1} W^c_{t-1} + x^{f,i}_{t-1} W^f_{t-1}}{P^i_{t-1}}, \quad (4.130)$$

Equation (4.129) can be consequently expressed as:

$$W^f_{t-1}\left(P^i_t A^i_t + B^i_t\right)\left[\Pi_{t-1} + \sum_{k=2}^{M} x^{f,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] + P^i_t W^c_{t-1} x^{c,i}_t\left[\Phi_{t-1} + \sum_{k=2}^{M} x^{c,k}_{t-1}\left(\frac{P^k_t + d^k_t}{P^k_{t-1}}\right)\right] - {P^i_t}^2 \beta_i = 0, \quad (4.131)$$

before being arranged in the following form:

$$P_t^i\Bigg[W_{t-1}^f A_t^i \Pi_{t-1} + W_{t-1}^c x_t^{c,i}\Phi_{t-1} + W_{t-1}^f A_t^i \sum_{k=2}^{M} x_{t-1}^{f,k}\left(\frac{d_t^k}{P_{t-1}^k}\right) + W_{t-1}^c x_t^{c,i}\sum_{k=2}^{M} x_{t-1}^{c,k}\left(\frac{d_t^k}{P_{t-1}^k}\right)$$
$$+\sum_{k=2}^{M}\frac{P_t^k}{P_{t-1}^k}\left(x_{t-1}^{f,k}W_{t-1}^f A_t^i + x_{t-1}^{c,k}W_{t-1}^c x_t^{c,i}\right)\Bigg] + W_{t-1}^f B_t^i \Pi_{t-1} + W_{t-1}^f B_t^i \sum_{k=2}^{M} x_{t-1}^{f,k}\left(\frac{P_t^k}{P_{t-1}^k}\right)$$
$$+ W_{t-1}^f B_t^i \sum_{k=2}^{M} x_{t-1}^{f,k}\left(\frac{d_t^k}{P_{t-1}^k}\right) - {P_t^i}^2\beta_i = 0. \quad (4.132)$$

The parameter α_{ij} associated with the pair of assets i and j is introduced as follows:

$$\alpha_{ij} = \frac{x_{t-1}^{f,j}W_{t-1}^f A_t^i + x_{t-1}^{c,j}W_{t-1}^c x_t^{c,i}}{P_{t-1}^j} \quad (4.133)$$

and the parameter ζ_i is defined such that:

$$\begin{aligned}\zeta_i =& W_{t-1}^f A_t^i \Pi_{t-1} + W_{t-1}^c x_t^{c,i}\Phi_{t-1} + W_{t-1}^f A_t^i \sum_{k=2}^{M} x_{t-1}^{f,k}\left(\frac{d_t^k}{P_{t-1}^k}\right) + W_{t-1}^c x_t^{c,i}\sum_{k=2}^{M} x_{t-1}^{c,k}\left(\frac{d_t^k}{P_{t-1}^k}\right)\\ =& W_{t-1}^f A_t^i\left(1-\sum_{k=2}^{M} x_{t-1}^{f,k}\right)(1+r_f) + W_{t-1}^c x_t^{c,i}\left(1-\sum_{k=2}^{M} x_{t-1}^{c,k}\right)(1+r_f)\\ &+\sum_{k=2}^{M}\left[\frac{d_t^k}{P_{t-1}^k}\left(x_{t-1}^{f,k}W_{t-1}^f A_t^i + x_{t-1}^{c,k}W_{t-1}^c x_t^{c,i}\right)\right]. \quad (4.134)\end{aligned}$$

Equation (4.132) can hence be written in the following compact form:

$$P_t^i\left(\zeta_i + \sum_{k=2}^{M} P_t^k \alpha_{ik}\right) + \sum_{k=2}^{M}\frac{P_t^k}{P_{t-1}^k}W_{t-1}^f B_t^i x_{t-1}^{f,k}$$
$$+ W_{t-1}^f B_t^i \Pi_{t-1} + W_{t-1}^f B_t^i \sum_{k=2}^{M} x_{t-1}^{f,k}\left(\frac{d_t^k}{P_{t-1}^k}\right) - {P_t^i}^2\beta_i = 0. \quad (4.135)$$

Successively, let χ_{ij} be defined as follows:

$$\chi_{ij} = \frac{W_{t-1}^f B_t^i x_{t-1}^{f,k}}{P_{t-1}^k} \quad (4.136)$$

and λ_i be introduced by the subsequent expression:

$$\lambda_i = W_{t-1}^f B_t^i \Pi_{t-1} + W_{t-1}^f B_t^i \sum_{k=2}^{M} x_{t-1}^{f,k}\left(\frac{d_t^k}{P_{t-1}^k}\right) \quad (4.137)$$
$$= W_{t-1}^f B_t^i\left[\left(1-\sum_{k=2}^{M} x_{t-1}^{f,k}\right)(1+r_f) + \sum_{k=2}^{M} x_{t-1}^{f,k}\frac{d_t^k}{P_{t-1}^k}\right]. \quad (4.138)$$

Inserting these latter parameters in Equation (4.135) further enables to obtain the following expression:

$$P_t^i\zeta_i + P_t^i\sum_{k=2}^{M} P_t^k\alpha_{ik} + \sum_{k=2}^{M} P_t^k\chi_{ik} - {P_t^i}^2\beta_i + \lambda_i = 0. \quad (4.139)$$

By isolating the nonzero terms when inserting the Kronecker delta in the summations, one can finally express Equation (4.139) as:

$$P_t^{i\,2}\big(\alpha_{ii}-\beta_i\big)+P_t^i\big(\zeta_i+\chi_{ii}\big)\sum_{k=2,k\neq i}^{M}P_t^kP_t^i\alpha_{ik}+\sum_{k=2,k\neq i}^{M}P_t^k\chi_{ik}+\lambda_i=0. \tag{4.140}$$

This concludes the presentation of the solvable set of $M-1$ nonlinear equations expressed in terms of the $M-1$ asset prices remaining to be evaluated at each time step of the simulations of the model.

Initialization of the parameters of the Fixed Income Market

Table (4.4) below introduces the initial values of the parameters associated with the Fixed Income market framework. As one could anticipate, the initial values of the modified and Macaulay durations approximate the maturities of the bonds.

Parameter	**Initial value associated with each asset**
YTM_0	$(0.025, 0.025, 0.025, 0.025, 0.025, 0.025, 0.025, 0.03)$
MV_0	$(1.00019, 1.00085, 1.00220, 1.000457, 1.00856, 1.01487, 1.02442, 1)$
$\text{MV}_0\vert_{\text{YTM}\uparrow}$	$(0.999214, 0.996099, 0.992972, 0.991099, 0.991015, 0.993365, 0.998971, 0.967742)$
$\text{MV}_0\vert_{\text{YTM}\downarrow}$	$(1.00118, 1.00563, 1.01152, 1.01825, 1.02645, 1.03691, 1.05061, 1.03448)$
ApproxModDur_0	$(0.980583, 4.76207, 9.25537, 13.5134, 17.5685, 21.4533, 25.2023, 33.3704)$
ApproxMacDur_0	$(1, 4.88112, 9.48676, 13.8513, 18.0078, 21.9896, 25.8324, 34.3715)$
ApproxConv_0	$(1.92309, 27.6518, 97.5912, 205.408, 347.599, 521.555, 725.624, 2224.69)$
DF_0	$0.980392, 0.915282, 0.845022, 0.776775, 0.710366, 0.645617, 0.582376,$ n.a.)
ISR_0	$0.0199010, 0.0177832, 0.0169104, 0.0169114, 0.0171721, 0.0175788, 0.0181027,$ n.a.)

Table 4.4: Initialization of the parameters of the Fixed Income market. The values correspond respectively to each asset in terms of increasing maturities, starting with the one-year zero-coupon risk free bond and ending with the perpetual bond.

4.2 Multi-asset financial bubbles in an agent-based model with noise traders' herding described by an n-vector Ising model

We introduce a multi-asset model of a financial market in which rational fundamentalists and trend following noise traders co-exist. The interactions and opinion formation of the noise traders are described by an $O(n)$ model. The n components of the generalized Ising model represent the different assets in which the traders can invest. Rejection-less transition probabilities are derived to describe realistic investment decisions at the "micro" level. The agent-based model is validated by testing for several characteristics of financial markets such as volatility clustering and fat-tails of the distribution of returns. Furthermore, the model is prone to develop bubbles and crashes. We distinguish three different regimes depending on the inverse temperature controlling the traders' propensity to herd and imitate others. In the subcritical regime, the traders' opinions are idiosyncratic and no bubbles emerge. Around the critical value, asynchronous bubbles emerge, triggered by an actual phase transition of the underlying $O(n)$ model. Above the critical value, synchronous bubbles emerge controlled by the time-varying external field representing the price momenta of the risky assets.

This section is based on the working paper Cividino, D., Westphal, R., & Sornette, D. (2020). Multi-asset financial bubbles in an agent-based model with noise traders' herding described by an n-vector Ising model.

4.2.1 Introduction

One of the most important concepts in complex systems theory is the emergence of highly non-trivial collective phenomena from the interactions between a large number of agents. The presence of repetitive interactions is the crucial element governing these phenomena. A clear analogy can be visualised between the interactions of the spins in a ferromagnetic material which tend to align their orientations while the temperature tends to push the system towards a disordered state, with the social imitation between agents which tends to polarize the class of agents towards a common preference, while stochastic idiosyncratic opinions among agents favor lack of consensus. These two systems illustrate the ubiquitous fight between order and disorder, be it in a ferromagnetic material or in a group of investors, which leads to a rich phenomenology and dynamics associated with phase transitions between different regimes.

In the 1970s, Weidlich (1971) introduced the idea to describe decision making within social groups and specifically their polarization as a physical ensemble of interacting spins. The idea to use the Ising model to represent opinion dynamics was further developed for example by Galam et al. (1982) who applied it to a strike process in a plant containing satisfied and dissatisfied workers and by Grabowski and Kosiński (2006) who took the spacial location of individuals on a complex network into account. Another application of Ising-based decision models is the voter model by Holley and Liggett (1975), in which the opinion of a voter is a binary variable stochastically changing under the influence of its neighbors' opinions. Roehner et al. (2004) have shown that, under partial information, the rational optimization of expected payoffs under a utility function that considers cultural norm, as well as herding, can be described by the Ising model. Many more Ising-like models have been developed to describe collective behavior of animal and human societies. See e.g. Sornette (2014) for a review and references therein.

The approaches to model opinion formation in social systems are also useful to understand the investment decisions of traders in a financial market. Neoclassical economic theory is based on the assumption of agents' rationality and helped to describe many macroeconomic phenomena. However, the assumptions of rational representative agents and general equilibrium are hard pushed to explain extreme events such as bubbles and crashes. For instance, the dynamical stochastic general

equilibrium (DSGE) models used by central banks to inform their monetary policies were impotent during the great financial crisis of 2008, as bubbles and crashes were by construction assumed impossible. This realisation has motivated renewed interest in agent-based models (Gualdi, Bouchaud, et al., 2015; Gualdi, Tarzia, et al., 2015), which emphasise the existence of many interactive decision makers with bounded rationality and subjected to limited and possibly asymmetric information. To better understand financial markets, it is crucial to embrace the fact that the world economy is a constantly evolving multi-agent complex system, that can be studied with agent-based models (ABM). In ABMs, the asset prices are endogenously defined by the agents' investment decisions. De Long et al. (1990, 2007) show that irrational traders with stochastic beliefs, so-called noise traders, can create endogenous financial bubbles from positive feedback in an ABM. Especially, the interaction of heterogeneous agents with heterogeneous beliefs can reproduce some of the characteristic features of financial markets known as "stylized facts". For example, Brock and Hommes (1997) introduced an ABM in which traders switch between predictors of the future return based on the past performance of these predictors. In Lux and Marchesi (1999), the traders switch between being rational fundamentalists and noise traders by comparing the expected returns of the strategies. Furthermore, the noise traders switch between optimistic and pessimistic mood influenced by the other traders' opinions. The model can reproduce some of the "stylized facts" such as excess kurtosis and volatility clustering.

However, the vast majority of ABMs have been concerned with modeling the dynamics of one risky asset, for instance a financial index or the stock of a firm traded in an organised exchange market, coexisting with a riskless asset such as a treasury bill. But a fundamental characteristic of investing is the possibility to diversify one's wealth among many assets. The typical investor, especially the large institutional investors, mutual funds, pension funds and the like that dominate the markets in terms of asset value under management, is mainly focused on optimising his diversified portfolio. Diversification of investments over the whole universe of assets has led to the rich literature on asset pricing, starting with the Capital Asset Pricing model, the Arbitrage Pricing Theory, and the whole "industry" of factor models (R. Campbell et al., 2016).

There are some ABMs that model several assets. Xu et al. (2014) developed an ABM in which both fundamentalists and noise traders maximise their expected utility under a CRRA utility function differing only in the construction of the expected return. Similarly, Chiarella et al. (2007) introduced a model of heterogeneous agents that maximise a CARA utility function and build their expectation of the future return based on past observations. Here, our goal is to extend this literature and address some of the limitations of previous works. In the aforementioned ABMs, the traders individually optimise their investment strategy without considering the other traders' decisions. However, phenomena such as bubbles and crashes occur due to the imitation of traders. In order to describe realistic price dynamics, we consider the opinion formation in a group of investors. In particular, our goal is to derive a multi-asset market model of interacting agents that is prone to develop bubbles and crashes in order to understand bubble formation among multiple assets. We are especially interested in the formation of time-synchronous and asynchronous bubbles as the synchronous emergence of a bubble among multiple assets creates systemic risk and the crash of one asset can trigger the other assets to crash as well, potentially resulting in a severe crash of the whole stock market.

Our model is an extension of the ABM with Ising-like characteristics introduced by Kaizoji et al. (2015). Fundamentalists and noise traders co-exist in this ABM of a financial market without switching between the strategies. The rational fundamentalists maximize their expected utility whereas the noise traders invest based on momentum following and social imitation. The noise trader class is implemented as an Ising model in which each spin represents one trader invested either in the risky or in the risk-free asset. The noise traders' decision to switch between the assets is influenced by the other traders' opinion, their idiosyncratic opinion, as well as the price momentum, which plays a role analogous to an external field of the Ising model. The collective opinion of the noise traders exhibit a phase transition that underpins the emergence of bubbles in the asset's price.

Building on the ability of the Agent-based model by Kaizoji et al. (2015) to create endogenous super-exponential bubbles in one asset, while being able to reproduce the most important features of financial markets, we present a multi-asset extension in which the social imitation and momentum following of traders are represented by an $O(n)$ model. We propose this formulation of the interactions between agents deciding to invest among several assets as an $O(n)$ model because it appears as the most elegant and rich description that naturally generalise the Ising model.

The paper is organised as follows. In the next section, the extension of the original market model by Kaizoji et al. (2015) to a multi-asset framework consisting of one risk-free and multiple risky assets is introduced. This includes the generalization of the dividend process, the wealth dynamics, and the traders' decision process. A special emphasis is put on the noise traders' stochastic dynamics for the investment derived from an $O(n)$ model. The price equation is derived from the market-clearing conditions according to Walras' theory of general equilibrium (Walras, 1954). In section 4.2.3, the time series resulting from the traders' interactions are analyzed and the "stylized-facts" of financial markets are tested. Section 4.2.4 examines the emergence of bubbles among the risky assets focusing on three different regimes of the noise traders' underlying $O(n)$ model-like structure. Section 4.2.5 concludes.

4.2.2 The agent-based market model

The market model evolves according to discrete-time dynamics, where each time-step represents one trading day, and is constituted of two classes of agents, the fundamentalist and the noise traders. They invest according to different strategies into one risk-free and n risky assets. Their investment decisions enter in the price equations, which govern the dynamics of the time-varying prices of the risky assets. The latter quantities present super-exponential growths, implying the presence of financial bubbles. In the following, we explain in detail the various components of the model.

The assets and the wealth dynamics

The model features one risk-free asset, representing a zero-coupon government bond with constant unitary price, yielding a constant rate of return r_f, and n risky assets, representing n stocks with time-varying prices $P_{k,t}$, paying stochastic dividends $d_{k,t}$ with $k \in \{1, \dots, n\}$.

Allowing for possible correlations of their dynamics, we describe their time evolution through n multiplicative growth processes

$$\begin{cases} d_{1,t} = (1 + r_t^{d,1}) d_{1,t-1} \\ d_{2,t} = (1 + r_t^{d,2}) d_{2,t-1} \\ \quad \vdots \\ d_{n,t} = (1 + r_t^{d,n}) d_{n,t-1} \end{cases} \tag{4.141}$$

where the stochastic growth factors follow a multivariate normal distribution

$$(r_t^{d,1}, r_t^{d,2}, \dots, r_t^{d,n}) \sim \mathcal{N}(\vec{\mu},\, \Sigma_d), \tag{4.142}$$

with mean $\vec{\mu} = (r_{d,1}, r_{d,2}, \dots, r_{d,n})$ and covariance matrix Σ_d. The stochastic dividend processes represent the impact of the real economy on the stocks' value. Consequently, the covariance matrix provides a framework to analyze the impact of correlations coming from the real economy on the price formation mechanism and the synchronization of bubbles. In the following, the variances and correlations of the dividend processes are small in order to focus on the endogenous dynamics resulting from the traders' decisions and interactions. The returns of the risky assets consist of the dividend yields $\vec{y}_t^{\,d}$ and the price returns $\vec{r}_t$. The dividend yields are defined as

$$\vec{y}_t^{\,d} = \left(\frac{d_{1,t}}{P_{1,t-1}}, \dots, \frac{d_{n,t}}{P_{n,t-1}} \right) \tag{4.143}$$

where $P_{k,t}$ is the price of risky asset k at time t and the price returns are defined as

$$\vec{r}_{p,t} = \left(\frac{P_{1,t}}{P_{1,t-1}} - 1, \ldots, \frac{P_{n,t}}{P_{n,t-1}} - 1 \right). \tag{4.144}$$

Consequently, the vector containing the total returns of the risky assets is defined as

$$\vec{r}_t^{tot} = \vec{y}_t^d + \vec{r}_t. \tag{4.145}$$

The traders' investment decisions are described in terms of the fraction of their wealth they invest into each asset. Thus, the portfolio of a trader i is constituted by n risky fractions $(x^i_{1,t}, \ldots, x^i_{n,t})$ and one risk-free fraction $x^i_{r_f,t}$. Borrowing and short-selling are not admitted in the market model, hence the wealth fractions including the risk-free fraction are constraint to $x^i_{k,t} \in [0, 1]$. Moreover, the fractions must sum to one at each time step t

$$x^i_{r_f,t} + \sum_{k=1}^{n} x^i_{k,t} = 1. \tag{4.146}$$

The wealth of trader i evolves according to

$$W^i_t = W^i_{t-1} \left[1 + r_f + \sum_{k=1}^{n} x^i_{k,t-1} r_t^{k,excess} \right], \tag{4.147}$$

where the quantity $r_t^{k,excess}$ represents the excess return of the risky asset k with respect to the risk-free return r_f

$$r_t^{k,excess} = \left(\frac{d_{k,t}}{P_{k,t-1}} + \frac{P_{k,t}}{P_{k,t-1}} - 1 - r_f \right). \tag{4.148}$$

Fundamentalist traders

The fundamentalists are rational risk-averse traders who at each time-step maximize the expected constant relative risk aversion (CRRA) utility function of the future wealth in terms of the risky fractions and for a given level of risk. The derivation of the fundamentalists follows the original description for one risky asset in Kaizoji et al. (2015) whereas the multi-asset extension follows Chiarella et al. (2009). At each time step, each fundamentalist trader constructs its portfolio $(x^f_{1,t}, \ldots, x^f_{n,t})$ solving the maximization problem

$$\underset{(x^f_{1,t}, \ldots, x^f_{n,t})}{\arg\max} \; E_t[U(W^f_{t+1}(x^f_{1,t}, \ldots, x^f_{n,t})], \tag{4.149}$$

where U represents the CRRA utility function with constant risk aversion γ

$$U(W) = \begin{cases} \log(W) & \gamma = 1 \\ \frac{W^{1-\gamma}}{1-\gamma} & \gamma \neq 1 \end{cases} \tag{4.150}$$

Each fundamentalist trader solves the same optimization problem. Hence, the investment impact can be considered at the aggregate level through a representative agent, whose wealth is equal to the sum of all the fundamentalists' wealth. The maximization problem (4.149) has been solved in Xu et al. (2014), here we report the final solution. The fundamentalist portfolio allocation strategy condenses into the equation

$$\begin{pmatrix} x^f_{1,t} \\ \vdots \\ x^f_{n,t} \end{pmatrix} = \frac{1}{\gamma} \mathrm{Cov}^{-1} \begin{pmatrix} E_{r,1} + \frac{d_{1,t}(1+r^{d,1})}{P_{1,t}} - r_f \\ \vdots \\ E_{r,n} + \frac{d_{n,t}(1+r^{d,n})}{P_{n,t}} - r_f \end{pmatrix} \tag{4.151}$$

where Cov^{-1} is the inverse matrix of the expected covariances of the future price returns estimated by the fundamentalist traders. Furthermore, the fundamentalists build an expectation $(E_{r,1}, \ldots, E_{r,n})$ of the price return of the risky assets. These expectations could in principle depend on time, but are assumed time-independent in the following for simplicity.

Equation (4.151) shows that the mean-variance trade-off present in the original model characterizes also the multi-asset framework.

Noise traders

As in Kaizoji et al. (2015), the noise traders' intrinsically stochastic investment strategy is driven by social imitation and trend following, see also Lux and Marchesi (1999). The central feature characterizing the noise traders class is its Ising-like structure which models the competition between the ordering force of social imitation and the disordering impact of idiosyncratic opinion. This very Ising-like structure explains the emergence of the bubbles and governs their dynamics. Indeed, these highly non-trivial collective phenomena emerge through a polarization phenomenon which constitutes an actual phase transition from a disordered state to an ordered one. The following section extends the Ising-like structure of the noise traders class described in Kaizoji et al. (2015) to a multi-asset framework. Specifically, we introduce an $O(n)$ model on the fully connected graph with an external field of price momenta to model the noise traders class.

Each of the N noise traders is associated with a spin vector

$$\vec{S}_i = (s_{i1}, \ldots, s_{in}) \in \mathbb{S}^{n-1} \tag{4.152}$$

representing its portfolio allocation. The positive components of the spin vector represent investments in the risky assets and the negative components represent investments in the risk-free asset. More precisely, the risky fraction invested in a, if s_{ia} is non-negative is

$$x_{ia} = s_{ia}^2 \quad \text{if } s_{ia} \geq 0. \tag{4.153}$$

The sum of all the negative components squared represents at the aggregate level the risk-free fraction

$$\sum_{a:s_{ia}<0}^{n} s_{ia}^2 = x_{r_f}. \tag{4.154}$$

This definition ensures that the condition of wealth conservation

$$\sum_{a=1}^{n} x_{ia} = 1. \tag{4.155}$$

is always satisfied.

Each noise trader, i.e. spin, interacts with all the others. Additionally, a vectorial external field $\vec{H}$ acts on each spin, modeling the trend following attitude of the noise investors. In particular, each component of the vector $\vec{H}$ corresponds to the price momentum H_k associated with the respective risky asset k, defined as an exponential moving average of the past price changes, constituting an indicator of the asset performance

$$H_{k,t} = \theta H_{k,t-1} + (1-\theta)\left(\frac{P_{k,t}}{P_{k,t-1}} - 1\right), \tag{4.156}$$

where $\theta \in [0, 1)$ controls the characteristic time window of memory of the past returns. The model is characterized by the Hamiltonian

$$\mathcal{H}(\{\vec{S}_1, \ldots, \vec{S}_N\}) = -\frac{1}{2N} \sum_{i \neq j=1}^{N} \vec{S}_i \cdot \vec{S}_j - \sum_{i=1}^{N} \vec{H} \cdot \vec{S}_i, \tag{4.157}$$

where we assume a uniform constant unit interaction strength.

The standard Boltzmann weight fully describing the statistical properties of the model, is taken as the equilibrium distribution

$$P(\{\vec{S}_1, \dots, \vec{S}_N\}) = e^{-\kappa \mathcal{H}}, \tag{4.158}$$

where the herding propensity parameter κ, first introduced in Kaizoji et al. (2015), corresponds to the inverse temperature β of the standard $O(n)$ model and governs the relative importance of the common investment preferences shared by the noise traders class with respect to the idiosyncratic opinion of each agent.

Derivation of the transition probabilities

We are interested in modeling the investments' dynamics of the noise traders class and not just their average properties, hence in the following the transition probabilities characterizing the stochastic dynamics of the model are derived. Our goal is to construct a Markov-chain Monte Carlo (MCMC) having (4.158) as its equilibrium distribution (to discuss?), which defines realistic dynamics for the traders' investments. The latter constitutes a crucial point of the following derivation, indeed we cannot rely on standard methods to generate a stochastic dynamics for the $O(n)$ model unless they give rise to a realistic description of the investment strategy of the noise traders from the finance point of view. No standard method was found that well fitted the task, hence in this and in the following two sections we derive an original method to generate a stochastic dynamics for the $O(n)$ model that allows for a realistic investment description from the finance perspective.

Each trading day, each noise trader updates its investment decision based solely on the information available up to the previous trading day. We start from the discrete-time master equation governing the evolution of the time-dependent conditional probability distribution

$$P(\vec{S}_l^t \mid \{\vec{S}_1^{t-1}, \dots, \vec{S}_{l-1}^{t-1}, \vec{S}_{l+1}^{t-1}, \dots, \vec{S}_N^{t-1}\}) \tag{4.159}$$

of a single spin $\vec{S}_l^t$ given all the others. In the following $\{\vec{S}_1^{t-1}, \dots, \vec{S}_{l-1}^{t-1}, \vec{S}_{l+1}^{t-1}, \dots, \vec{S}_N^{t-1}\}$ will be considered as fixed parameters, each having a specific value on the $(n-1)$-sphere $\mathbb{S}^{n-1}$. To lighten the notation we denote

$$P(\vec{S}_l^t = \vec{A} \mid \{\vec{S}_1^{t-1}, \dots, \vec{S}_{l-1}^{t-1}, \vec{S}_{l+1}^{t-1}, \dots, \vec{S}_N^{t-1}\}) = P(\vec{A}, t). \tag{4.160}$$

The discrete-time master equation for the conditional probability to have $\vec{S}_l^t = \vec{A} \in \mathbb{S}^{n-1}$ reads

$$\begin{aligned} &\frac{P(\vec{A}, t) - P(\vec{A}, t - \Delta t)}{\Delta t} = \\ &\int\limits_{\vec{B} \in \mathbb{S}^{n-1}} W(\vec{B} \to \vec{A}) P(\vec{B}, t - \Delta t) - W(\vec{A} \to \vec{B}) P(\vec{A}, t - \Delta t), \end{aligned} \tag{4.161}$$

where we assume a unit time increment $\Delta t = 1$, corresponding to one trading day. In order to derive the transition rates characterizing the MCMC, we set the stationarity condition using the standard detailed balance rule. In this case, it reads

$$\frac{W(\vec{A} \to \vec{B})}{W(\vec{B} \to \vec{A})} = \frac{P(\vec{B})}{P(\vec{A})} \quad \forall \vec{A}, \vec{B} \in \mathbb{S}^{n-1}. \tag{4.162}$$

Using the definition of conditional probability we can easily switch to the joint probabilities

$$\begin{aligned} \frac{P(\vec{B})}{P(\vec{A})} &= \frac{P(\vec{S}_l = \vec{B} \mid \{\vec{S}_1, \dots, \vec{S}_{l-1}, \vec{S}_{l+1}, \dots, \vec{S}_N\}}{P(\vec{S}_l = \vec{A} \mid \{\vec{S}_1, \dots, \vec{S}_{l-1}, \vec{S}_{l+1}, \dots, \vec{S}_N\}} \\ &= \frac{\frac{P(\vec{S}_l = \vec{B}, \{\vec{S}_1, \dots, \vec{S}_{l-1}, \vec{S}_{l+1}, \dots, \vec{S}_N\}}{P(\{\vec{S}_1, \dots, \vec{S}_{l-1}, \vec{S}_{l+1}, \dots, \vec{S}_N\})}}{\frac{P(\vec{S}_l = \vec{A}, \{\vec{S}_1, \dots, \vec{S}_{l-1}, \vec{S}_{l+1}, \dots, \vec{S}_N\}}{P(\{\vec{S}_1, \dots, \vec{S}_{l-1}, \vec{S}_{l+1}, \dots, \vec{S}_N\})}}. \end{aligned} \tag{4.163}$$

Simplifying the denominators we get the ratio between the joint probabilities, which we can explicitly compute as

$$\begin{aligned}&\frac{P(\vec{S}_l=\vec{B},\{\vec{S}_1,\dots,\vec{S}_{l-1},\vec{S}_{l+1},\dots,\vec{S}_N\}}{P(\vec{S}_l=\vec{A},\{\vec{S}_1,\dots,\vec{S}_{l-1},\vec{S}_{l+1},\dots,\vec{S}_N\}}\\&=\frac{\frac{1}{Z}e^{-\kappa(-\frac{1}{2N}\sum_{i\neq j\neq l}\vec{S}_i\cdot\vec{S}_j-\frac{1}{N}\sum_{i\neq l}\vec{S}_i\cdot\vec{B}-\sum_{i\neq l}\vec{H}\cdot\vec{S}_i-\vec{H}\cdot\vec{B})}}{\frac{1}{Z}e^{-\kappa(-\frac{1}{2N}\sum_{i\neq j\neq l}\vec{S}_i\cdot\vec{S}_j-\frac{1}{N}\sum_{i\neq l}\vec{S}_i\cdot\vec{A}-\sum_{i\neq l}\vec{H}\cdot\vec{S}_i-\vec{H}\cdot\vec{A})}}.\end{aligned}\tag{4.164}$$

Simplifying the common factors at numerator and denominator we get (minus a constant term of order $\frac{1}{N}$ negligible in our case of large N)

$$\begin{aligned}\frac{P(\vec{B})}{P(\vec{A})}&=e^{-\kappa(-\frac{1}{N}\sum_i\vec{S}_i\cdot\vec{B}-\vec{H}\cdot\vec{B}+\frac{1}{N}\sum_i\vec{S}_i\cdot\vec{A}-\vec{H}\cdot\vec{A})}\\&=e^{\kappa(\frac{\sum_i\vec{S}_i\cdot(\vec{B}-\vec{A})}{N}+\vec{H}\cdot(\vec{B}-\vec{A}))}.\end{aligned}\tag{4.165}$$

Hence the ratio of the transition probabilities is

$$\frac{W(\vec{A}\to\vec{B})}{W(\vec{B}\to\vec{A})}=\frac{P(\vec{B})}{P(\vec{A})}=e^{\kappa(\frac{\sum_i\vec{S}_i\cdot(\vec{B}-\vec{A})}{N}+\vec{H}\cdot(\vec{B}-\vec{A}))}.\tag{4.166}$$

There are different possibilities to split the expression in two well-defined transition probabilities. As already mentioned, the probabilities need to define a realistic dynamics. For example the standard Metropolis-Hastings rule (Metropolis et al., 1953), that assumes the transition probabilities are composed of a uniform move proposal probability $\mathcal{P}$ and a move acceptance probability $\mathcal{A}$

$$W(\vec{A}\to\vec{B})=\mathcal{P}(\vec{A}\to\vec{B})\mathcal{A}(\vec{A}\to\vec{B}),\tag{4.167}$$

would lead to a very bad modeling of the real actions of the traders. Indeed, in our setup the acceptance probability would read

$$\mathcal{A}(\vec{A}\to\vec{B})=\begin{cases}1\quad\text{if }(\frac{\sum_i\vec{S}_i\cdot(\vec{B}-\vec{A})}{N}+\vec{H}\cdot(\vec{B}-\vec{A}))>0\\e^{\kappa(\frac{\sum_i\vec{S}_i\cdot(\vec{B}-\vec{A})}{N}+\vec{H}\cdot(\vec{B}-\vec{A}))}\quad\text{otherwise.}\end{cases}\tag{4.168}$$

The uniform proposal probability together with equation (4.168) define a totally unrealistic dynamics where at each time-step, each noise trader chooses at random a new portfolio of investments to switch to and its trade decision is arbitrarily accepted or rejected according to a probability $\mathcal{A}(\vec{A}\to\vec{B})$.

In order to obtain a realistic stochastic dynamics we need a rejection-less MCMC. We impose

$$\mathcal{A}(\vec{A}\to\vec{B})=\mathcal{A}(\vec{B}\to\vec{A})=1\quad\forall\vec{A},\vec{B}\in\mathbb{S}^{n-1}.\tag{4.169}$$

Hence, from the detailed balance condition it follows

$$\frac{\mathcal{P}(\vec{A}\to\vec{B})}{\mathcal{P}(\vec{B}\to\vec{A})}=\frac{P(\vec{B})}{P(\vec{A})}=e^{\kappa\left(\frac{\sum_i\vec{S}_i\cdot(\vec{B}-\vec{A})}{N}+\vec{H}\cdot(\vec{B}-\vec{A})\right)}.\tag{4.170}$$

This condition is satisfied setting

$$\mathcal{P}(\vec{A}\to\vec{B})=e^{\kappa_t\left(\frac{\sum_i\vec{S}_i\cdot\vec{B}}{N}+\vec{H}\cdot\vec{B}\right)}.\tag{4.171}$$

Normalizing we finally get the following transition probabilities

$$W(\vec{A}\to\vec{B})=\frac{e^{\kappa\left(\frac{\sum_i\vec{S}_i\cdot\vec{B}}{N}+\vec{H}\cdot\vec{B}\right)}}{\int_{\vec{K}\in\mathbb{S}^{n-1}}e^{\kappa\left(\frac{\sum_i\vec{S}_i\cdot\vec{K}}{N}+\vec{H}\cdot\vec{K}\right)}}.\tag{4.172}$$

We first notice that the transition probability from state $\vec{A}$ to state $\vec{B}$ is independent of the initial state $\vec{A}$, a typical property of mean-field models, indeed it can be easily showed that our fully-connected $O(n)$ model is equivalent to the correspondent mean-field version in the large N regime, in which we are.

Most important, equation (4.172) defines realistic dynamics from the point of view of the investment strategy of the noise traders. We show it in the next section, discussing the connection between the Ising-like model we have constructed for our agents and the framework of decision theory.

Decision-theoretic interpretation of the transition probabilities The form of the probability distribution (4.172) coincides with a continuous version of the Logit probability distribution for the arrival states. The Logit distribution is throughout used and studied in the discrete choice theory. McFadden has shown in (McFadden, 1981), that the Logit probability distribution actually models individuals who maximize a utility function which has an implicit random idiosyncratic part. Indeed, if each agent makes its choice maximizing

$$s^* = \arg\max_s \{\beta u_s + \eta_s\}, \tag{4.173}$$

where s represents the possible choices from a finite set, u_s is the deterministic part of the utility function, η_s is a random variable and β plays the role of the inverse temperature, McFadden proved that $P(s^* = s)$ coincides with the Logit distribution if the random variable η_s is distributed according to the Gumbel distribution, with cumulative distribution function

$$F_\eta(x) = e^{-e^{-\frac{x-\mu}{\lambda}}}. \tag{4.174}$$

In our case the space of choices is continuous, the choices live on the hypersphere $\mathbb{S}^{n-1}$, $\vec{B} \in \mathbb{S}^{n-1}$. Nonetheless, the result of McFadden has been extended to the the case of continuous space of choices, see for example Ben-Akiva et al. (1985). Hence, modeling the investment decisions with the multivariate distribution (4.172), we are actually modeling traders whose decision process boils down to the solution of the maximization problem

$$\vec{B}^* = \arg\max_{\vec{B}\in\mathbb{S}^{n-1}} \{\kappa u_{\vec{B}} + \eta_{\vec{B}}\}, \tag{4.175}$$

where the deterministic utility function is given by

$$u_{\vec{B}} = \left(\frac{\sum_i \vec{S}_i}{N} + \vec{H}\right) \cdot \vec{B}. \tag{4.176}$$

Having established the connection between the transition rates characterizing the noise traders' investment strategy and the Logit distribution we have unveiled the bridge between the present Ising-like model of the noise traders class and the framework of decision theory in which we can understand the decision process undertaken by each agent. This is, once again, an example of the deep connection between the two approaches in modeling social systems, as also discussed in Sornette (2014).

In light of this, all the modeling of the "macro" noise traders class can be understood starting from the "micro" stochastic decision process of each noise investor. Due to their herding behavior, they tend to align their portfolio allocations to the average portfolio allocation of the class.

The herding behavior is modeled by the deterministic part of the utility optimization in equation (4.176), which can be considered as the rational part of the choice process undertaken by each agent. In absence of clear information, imitation of others can be the unique rational strategy possible. The Gumbel random variable instead models the random idiosyncratic part, specific to each trader, which enters in the decision process.

In particular, formula (4.175) clarifies how the herding propensity κ governs the relative importance of the deterministic part of the utility function common to all the traders, hence pushing the

traders to polarize towards the same investments, with respect to the random part modeled by the Gumbel distribution representing the importance of the idiosyncratic opinion of each individual.

We utilize this connection to construct an algorithm to simulate the noise traders' investment decisions. To generate the dynamics of our model we have to sample realizations of the random vector $\vec{B} = (B_1, \dots, B_n)$ from the multivariate probability distribution (4.172). Unfortunately, constructing a sampling procedure is difficult due to the high dimensionality of the distribution and to the non-trivial relation between the components of the random vector, $\sum_{k=1}^{n} B_k^2 = 1$.

Nevertheless, discretizing the space of choices $\mathbb{S}^{n-1}$ results in a discrete Logit distribution and instead of sampling the choices directly from the probability density function we can model each decision as the discrete version of the maximization problem (4.175).

At the beginning of the simulation, we discretize the space of choices $\mathbb{S}^{n-1}$ once and for all to preserve the detailed balance condition. Then at each time-step, the trading decision of each noise trader, i.e. the reallocation of its portfolio, is generated according to the following algorithm.

Algorithm 1: Simulation of the noise traders' investments (decision-theoretic approach)

for *each noise trader* i **do**
  for $p \in \{1, \dots, N_{points}\}$ **do**
    generate i.i.d. Gumbel RV η_p;
  end
  $\vec{B}_i^* = \arg\max_{\vec{B}_p} \{\kappa_t (\frac{\sum_i \vec{S}_i}{N} + \vec{H}) \cdot \vec{B}_p + \eta_p\}$
end

We comment again that the deterministic quantities entering the decision process refer to the previous trading day and are common to all traders. Unfortunately, despite the intuitive interpretation of the actual decision process of each agent, the method suffers from the curse of dimensionality. The computational cost of the simulation is exponential in the number N_{points} used to discretize the hypersphere and clearly a large number of points is needed to have a good discretization and a realistic simulation. Moreover, the problem aggravates for increasing dimensionality n and we are interested in building an algorithm well scalable to a large number of risky assets n, in order to have a realistic description of financial markets.

To overcome the computational problems of the present method, in the next section we derive a different algorithm to generate the stochastic dynamics of the noise traders' investments.

Symmetry-based approach

To construct an efficient method to sample from the probability density function (4.172), we exploit its symmetry property. The method is constituted of two steps. In the first, we sample an angle theta from a univariate distribution and then in the second we sample uniformly at random from a particular hypersphere $\mathbb{S}^{n-2}$, a subset of the original space of choices $\mathbb{S}^{n-1}$. We start from the non-normalized form of the transition rates

$$W(\vec{A} \to \vec{B}) \propto e^{\kappa_t \left(\frac{\sum_i \vec{S}_i}{N} + \vec{H}\right) \cdot \vec{B}}. \tag{4.177}$$

We explicate the symmetry property of the dot product entering the density function, expressing it in terms of the angle θ between the vector

$$\vec{M} = \frac{\sum_i \vec{S}_i}{N} + \vec{H} \tag{4.178}$$

and the vector $\vec{B}$. It is always possible to define a unique angle between two vectors in any inner product space and in particular in the Euclidean space $\mathbb{R}^n$ to which our vector belongs to. Hence we express the dot product as

$$\kappa_t \vec{M} \cdot \vec{B} = \kappa_t \|\vec{M}\| \cos\theta, \tag{4.179}$$

where we used that the norm of $\vec{B}$ is exactly equal to one. Taking advantage of this result and discarding a constant factor independent of θ, we can rewrite the probability distribution as

$$P(\theta) \propto e^{\kappa_t \|\vec{M}\| \cos\theta} (\sin\theta)^{n-2}, \tag{4.180}$$

which now effectively depends only on one variable

$$\theta \in [0, 2\pi). \tag{4.181}$$

The sets of equiprobable choices, i.e. equiprobable vectors $\vec{B}$, are defined by the conditions

$$\begin{cases} \frac{1}{\|\vec{M}\|} \vec{M} \cdot \vec{B} = \cos\theta \\ \|\vec{B}\| = 1 \end{cases} \tag{4.182}$$

The first condition defines an hyperplane in $\mathbb{R}^n$. In fact it can be written as

$$m_1 b_1 + m_2 b_2 + \ldots + m_n b_n = \cos\theta, \tag{4.183}$$

where $m_1, \ldots, m_n$ are fixed coefficients. The second condition instead enforces the choice vectors to belong to the hypersphere $\mathbb{S}^{n-1}$. We now use the geometric fact that the intersection of a hypersphere $\mathbb{S}^{n-1}$ and a n-dimensional hyperplane, is still a hypersphere, yet of one less dimension. Indeed, the system (4.182) defines a hypersphere $\mathbb{S}^{n-2}$ in $\mathbb{R}^n$, with center

$$\vec{C} = \cos\theta \frac{1}{\|\vec{M}\|} \vec{M} \tag{4.184}$$

and radius

$$r = \sqrt{1 - \cos^2\theta} = \sin\theta. \tag{4.185}$$

This fact will be crucial in constructing the sampling algorithm.

First of all, we focus on the angle θ, which represents for a noise trader the reallocation of its portfolio. Before sampling it, we take into account the status-quo bias which could push the noise trader to maintain its present portfolio composition through a Bernoulli random variable.

Each trading day, before proceeding with its trading decisions, the noise trader first decides if to perform any trading moves at all or to just be inactive on the financial market for that day and hold its previous portfolio allocation position which he considers solid and profitable.

The trader decides to be active in the financial market and modifies its portfolio allocation with a probability

$$P(active) = \min\left\{1, \frac{1}{t_h} e^{\kappa \|\vec{M}\|}\right\}. \tag{4.186}$$

The probability to be inactive and hold the previous portfolio allocation is clearly

$$P(holding) = 1 - P(active). \tag{4.187}$$

The parameter t_h represents the average number of trading days the noise trader keeps its asset in absence of herding behaviour. In this way, we can directly control the trading frequency and the intensity of the oscillating behavior of the time series characterizing the resulting market dynamics.

In case the trader decides to be active, we have to sample an angle θ representing its portfolio reallocation. Since we have proven that the complicated multivariate distribution (4.177) actually reduces to the univariate distribution (4.180), we can directly sample from the latter with rejection sampling, having effectively overcome the problem related to the curse of dimensionality. The rejection sampling method is presented in the following algorithm.

Algorithm 2: Rejection sampling from the univariate $P(\theta)$ distribution

while $(u * e^{\kappa_t \|\vec{M}\|}) > e^{\kappa_t \|\vec{M}\| \cos\theta} (\sin\theta)^{n-2}$ **do**
  $\theta = Uniform(0,1) * 2\pi$;
  $u = Uniform(0,1)$;
  (the uniform RVs are sampled i.i.d.)
end
return θ;

After having sampled an angle θ, we have to choose a vector uniformly at random from the equiprobable set defined by that angle. Fortunately, the set is just a hypersphere $\mathbb{S}^{n-2}$ and we can rely on an efficient algorithm to perform the sampling. Indeed, due to the spherical symmetry property of the multivariate normal distribution, the normalized random vector whose components are sampled in an i.i.d. manner from the standard normal distribution $\mathcal{N}(0,1)$

$$\vec{\mathcal{B}}_{unnorm} = \big(\mathcal{N}_1(0,1), \ldots, \mathcal{N}_{n-1}(0,1)\big) \tag{4.188}$$

$$\vec{\mathcal{B}}^*_{n-1} = \frac{\vec{\mathcal{B}}_{unnorm}}{\|\vec{\mathcal{B}}_{unnorm}\|} \tag{4.189}$$

is uniformly distributed on $\mathbb{S}^{n-2}$. To immerse the vector in $\mathbb{R}^n$ we have to add one extra zero component, for example at the beginning of the vector, effectively increasing its dimensionality by one. The new vector is

$$\vec{\mathcal{B}}^*_n = \big(0, \vec{\mathcal{B}}^*_{n-1}\big). \tag{4.190}$$

Moreover, the hypersphere has to be translated and its radius rescaled according to (4.184) and (4.185). Finally, in order to correctly represent the intersection between the higher dimensional $\mathbb{S}^{n-1}$ hypersphere and the hyperplane, the $\mathbb{S}^{n-2}$ hypersphere needs to be rotated in such a way that the unit versor, corresponding to the extra component added in (4.190), is rotated to the direction of the normalized vector $\frac{1}{\|\vec{M}\|}\vec{M}$.

We have to construct the orthogonal matrix R representing the rotation of the unit versor

$$\vec{X} = (1, 0, 0, \ldots, 0), \tag{4.191}$$

to the direction of the vector $\frac{1}{\|\vec{M}\|}\vec{M}$. The matrix R has to satisfy

$$\frac{1}{\|\vec{M}\|}\vec{M} = R\vec{X}. \tag{4.192}$$

In two or three dimensions, such a rotation is given by the standard matrices containing sine and cosine functions. In the general case of n dimensions, finding an efficient and numerical stable algorithm is not an easy task. We use an approach based on Givens rotations and to construct the rotation we thoroughly refer to Zhelezov (2017).

Eventually, the choice vector $\vec{B}$ sampled from the distribution (4.172), representing the noise trader's portfolio reallocation, is given by

$$\vec{B}^* = \sin\theta R\vec{\mathcal{B}}^*_n + \cos\theta \frac{1}{\|\vec{M}\|}\vec{M}. \tag{4.193}$$

Summarizing, the sampling algorithm is as follows.

Algorithm 3: Simulation of the noise traders' investments (symmetry-based approach)

if $Uniform(0,1) < \min\left\{1, \frac{1}{t_h} e^{\kappa \|\vec{M}\|}\right\}$ **then**
  sample angle θ from $P(\theta)$ with algorithm 2;
  sample uniformly a vector $\vec{\mathcal{B}}^*_{n-1} \in \mathbb{S}^{n-2}$;
  $\vec{\mathcal{B}}^*_n = (0, \vec{\mathcal{B}}^*_{n-1}) \in \mathbb{S}^{n-1}$;
  construct the rotation matrix R following (Zhelezov, 2017);
  $\vec{B}^* = \sin\theta R \vec{\mathcal{B}}^*_n + \cos\theta \frac{1}{\|\vec{M}\|} \vec{M}$;
else
  $\vec{B}^*$ = previous time-step portfolio;
end
return $\vec{B}^*$;

Market clearing and price equations

The market price is set according to a Walrasian auction, i.e. at each time-step the supply and the demand must equilibrate (Walras, 1954).

Setting the aggregate excess demand to zero in the original case of only one risky asset leads to a second-order equation in the unknown P_t. The equation can be solved explicitly giving a unique physical solution, i.e. a positive price, which represents the new price of the risky asset. Extending the model to n risky assets, the equilibrium condition has to hold simultaneously for each asset.

Defining the excess demand from time $t-1$ to t for each risky asset k for the trader i with $i \in \{f, n\}$ as

$$\Delta D^{i,k}_{t-1\to t} = W^i_t x^i_{k,t} - W^i_{t-1} x^i_{k,t-1} \frac{P_{k,t}}{P_{k,t-1}} \tag{4.194}$$

together with the risky fractions defined in (4.151) for the fundamentalists and in (4.153) for the noise traders, the equilibrium condition translates into the system

$$\begin{cases} \Delta D^{f,1}_{t-1\to t} + \Delta D^{n,1}_{t-1\to t} = 0 \\ \Delta D^{f,2}_{t-1\to t} + \Delta D^{n,2}_{t-1\to t} = 0 \\ \quad \vdots \\ \Delta D^{f,n}_{t-1\to t} + \Delta D^{n,n}_{t-1\to t} = 0. \end{cases} \tag{4.195}$$

The system (4.195) is a non-linear system in the n unknowns $P_{1,t}, \ldots, P_{n,t}$, where each equation is a polynomial equation of degree $n+1$ in all the unknowns. The system is solved numerically with an iterative method, based on the Hybrid algorithm proposed in Powell (1968) and Powell (1970), derived from the classical Newton–Raphson algorithm. Using the prices at the previous time-step as initial condition for the numerical solver, the method consistently converges to the correct physical solution for a wide range of parameters close to the real market quantities.

Parameters

The market parameters and initial values are chosen such that each time-step represents a typical trading day. In particular, such that the standard deviation of daily returns resembles the 1-2% observed in real markets following the derivation in Kaizoji et al. (2015). The complete set of parameters is reported in table 4.5. All rates such as the risk-free return r_f, the initial dividend $d_{i,0}$, the dividend growth rate $r_{d,i}$, and the expected returns $E_{r,i}$, as well as variances are reported as daily values. Fundamentalists and noise traders are initialized with equal wealth $W^n_0 = 10^9$ and the simulations are conducted over $T = 5000$ time-steps which corresponds to 20 years assuming that

	Parameters	
Assets	$N_{stocks} = 4$	$r_f = 4 \times 10^{-5}$
	$r_{d,i} = 1.6 \times 10^{-4}\ \forall i$	$d_{i,0} = 1.6 \times 10^{-4}\ \forall i$
	$P_{i,0} = 1\ \forall i$	$\Sigma^d_{i,i} = 1.6 \times 10^{-5}\ \forall i$
Fundamentalist traders	$W_0^f = 10^9$	$E_{r,i} = 1.6 \times 10^{-4}\ \forall i$
	$\Sigma^f_{i,i} = 0.0004\ \forall i$	
Noise traders	$W_0^n = 10^9$	$N = 1000$
	$\theta = 0.99$	$H_{i,0} = 1.6 \times 10^{-4}\ \forall i$
Market	$T = 5000$	

Table 4.5: Set of parameters and initial values used in the simulations.

one year contains approximately 250 trading days. No correlation among the dividend processes is assumed. Regarding the fundamentalist traders, the expected covariance matrix implementation is divided into a vector of expected variances and a matrix of expected correlations. The variances are

$$\Sigma^f_{i,i} = 0.0004 \quad i = 1, \ldots, 4 \tag{4.196}$$

and the correlation matrix is set equal to

$$C^f = \begin{pmatrix} 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{pmatrix} \tag{4.197}$$

The initial investment decisions for fundamentalists and noise traders are as follows

$$\vec{x}_0^f = (0.075, 0.075, 0.075, 0.075) \tag{4.198}$$

$$\vec{x}_0^n = (0.125, 0.125, 0.125, 0.125), \tag{4.199}$$

where each component represents the investment into one of the $n = 4$ risky assets. The remaining fraction of each traders' wealth is invested into the risk-free asset.

Finally, an average holding time of ten trading days is imposed with the parameter $t_h = 10$.

The constant risk aversion is endogenously computed at the beginning of the simulation from the initial conditions as

$$\gamma = \frac{E_{r,1} + \frac{d_{1,0}(1+r^{d,1})}{P_{1,0}} - r_f}{\mathrm{Cov}_{1,1} x^f_{1,0} + \ldots + \mathrm{Cov}_{1,n} x^f_{n,0}}, \tag{4.200}$$

which constitutes a natural generalization of the original model's formula in Kaizoji et al. (2015), also adopted in Damiani (2019) and Kopp (2020).

We are interested in analyzing the impact of both constant herding propensity κ and time-varying κ_t, in several ranges of values that will be specified from time to time. In particular, the time-varying herding propensity parameter models the impact of a changing geopolitical and economical situation on the tendency to herd of the noise traders and is defined by an Ornstein-Uhlenbeck stochastic process as

$$\kappa_t = \kappa_{t-1} + \eta_\kappa(\mu_\kappa - \kappa_{t-1}) + \sigma_\kappa \nu_t, \tag{4.201}$$

with $\nu_t \sim \mathcal{N}(0, 1)$. The mean reversion strength η_κ and the standard deviation σ_κ are explicitly indicated for each different simulation.

The core of the model is implemented in *C++*. Each part of the model (e.g. the fundamentalist traders, the risky assets, the price equations) is implemented by one specific class, following an object-oriented programming paradigm. To have reproducible results, a pseudo-random number generator with a random seed specified as a run-time parameter is used. The results of the simulation are stored in a database using the *HDF5* high-performance data software library. Then all the analysis of the data and the plotting are performed with *Python*.

4.2.3 Time series analysis and stylized facts of financial markets

Time series analysis

In this section, we present the time series resulting from a simulation characterized by the set of parameters introduced in the previous section. Figure 4.15 shows characteristic price time series of the four assets resulting from the traders' investment decisions. In this simulation, the herding propensity follows an Ornstein-Uhlenbeck process to represent time-varying susceptibility to herding and momentum following observed empirically. The mean value $\mu_k = 0.98 \cdot \kappa_c$ is defined slightly below the critical value $\kappa_c = 4$, such that κ_t transiently fluctuates into the ordered regime. The mean reversion strength $\eta = 0.013$ and the standard deviation $\sigma_\kappa = 0.25\kappa_c\sqrt{2\eta}$ are defined such that κ_t returns from two standard deviations above the mean to the subcritical regime in approximately $\Delta T = 250$ time steps, see Kaizoji et al. (2015) for the derivation.

The prices contain volatile and stable regimes while demonstrating a long-term growth rate that is similar for all four assets. The long-term growth rate is equal to the growth rate of the dividend process as verified for the two asset model in Westphal and Sornette (2020a). Figure 4.16 exemplary presents one of the assets in more detail. The comparison of the price time series with the κ_t process shows that the noise traders create bubbles by polarizing their opinion when κ_t is above the critical value. This is for example the case between $t = 4000$ and $t = 4500$ where the increase of the noise traders' fraction of wealth invested in the asset results in a bubble that grows over a time interval of two years. When κ_t is below the critical value, for example between $t = 1000$ and $t = 1500$, the noise traders' opinions are disordered, the price is stable and dominated by the growth of the asset. Furthermore, the return time series contains regimes with a large amplitude and regimes with a small amplitude of returns. This tendency of large absolute values of the return to cluster together is known as volatility clustering.

The price momentum used in the noise trader's investment decision is calculated as the exponential moving average of returns. Thus, it exhibits similar peaks and troughs as the price time series but lags a few time steps behind. The lag is controlled by the memory length $\sim 1/(1-\theta)$ the noise trader uses to calculate the momentum. Similarly, the fraction of wealth the noise trader invests in the asset contains peaks and valleys. The invested fraction controls the price changes and the price changes render positive feedback on the risky fraction via the momentum. The fundamentalists invest countercyclically by investing proportionally to the dividend-price ratio. This means that they decrease their wealth fraction invested in the asset during a bubble and return to their normal risky fraction during the crash. Consequently, the fundamentalists become wealthier than the noise traders in the long-term. During bubbles, the noise traders' wealth increases due to their larger exposure to the risky asset. However, during a crash, they lose most of their transiently acquired wealth while the fundamentalists are less affected by the crash and can even buy the asset when it is undervalued.

Stylized facts of the financial markets

Financial time series feature the presence of ubiquitous statistical properties independent of the details of the series itself (Cont, 2001). These emerging empirical properties have been observed across a wide range of instruments, markets, and time periods and they constitute the so-called stylized facts of the financial markets.

Figure 4.15: Price time series resulting from a simulation with four risky and one risk-free assets featuring an Ornstein-Uhlenbeck kappa process κ_t, stochastically fluctuating near the critical value κ_c. The simulation parameters are given in table 4.5. Several bubbles are identifiable as super-exponential growths of the prices.

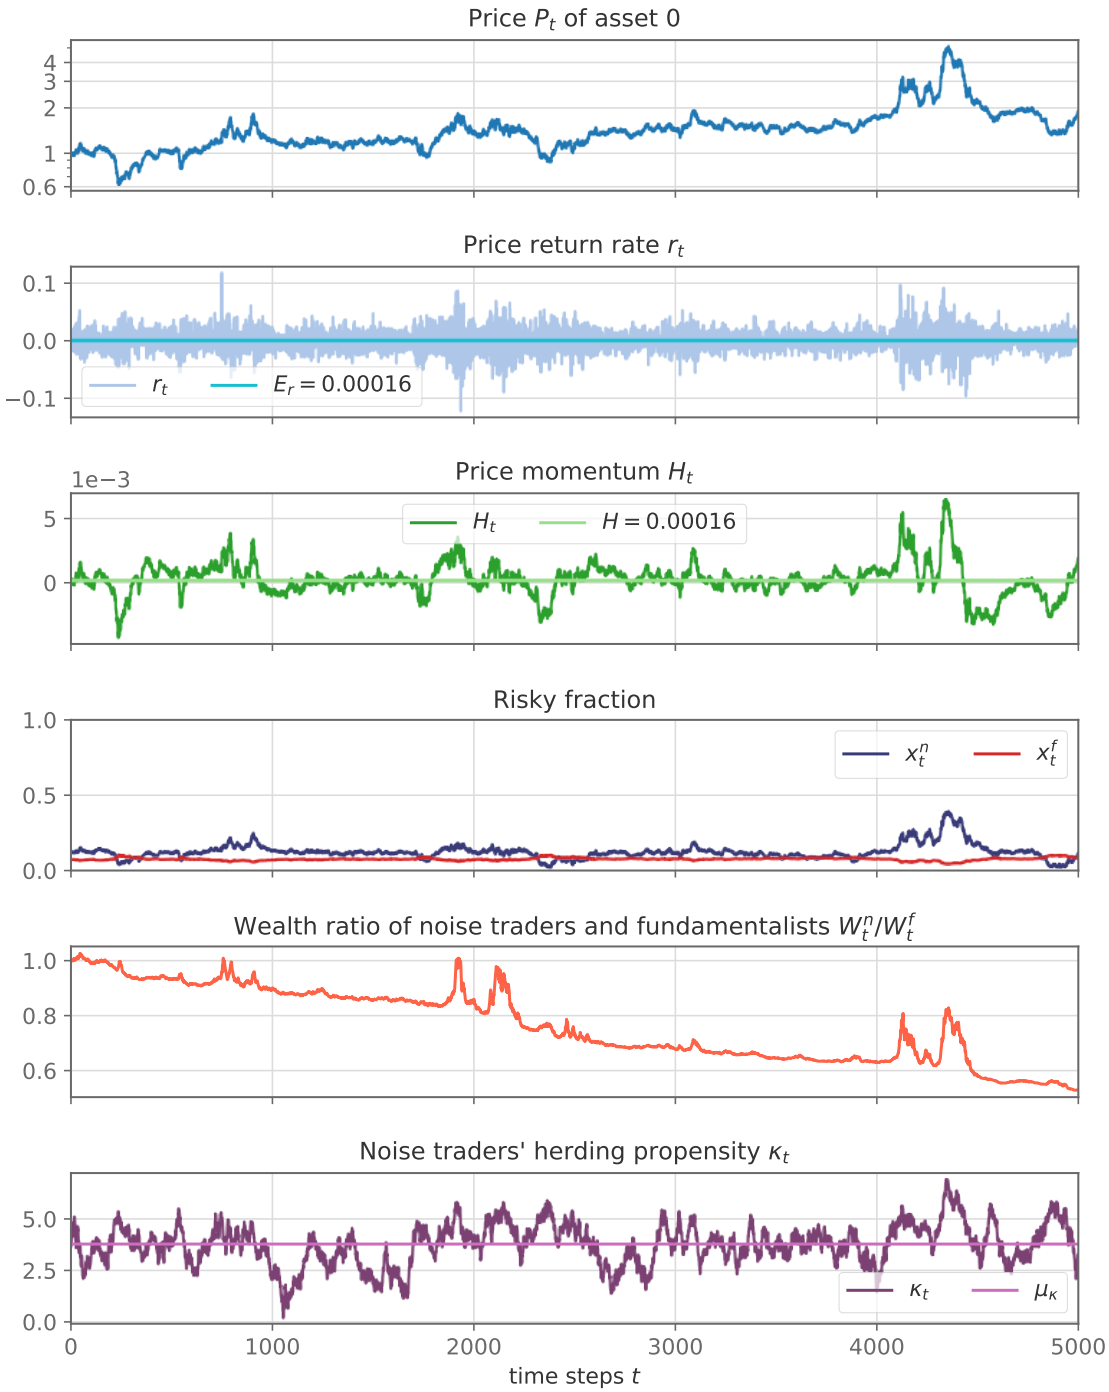

Figure 4.16: The figure shows the detailed time series characterizing the risky asset 0. It shows the price, the actual return and the return expected by the fundamentalists, the price momentum, the fraction of their wealth that fundamentalists and noise traders invest into the asset, the wealth ratio of noise traders and fundamentalists over time and the herding propensity kappa. The market is simulated over $T = 5000$ time steps with the parameters given in 4.5. For a comparison with the other assets see figure 4.15.

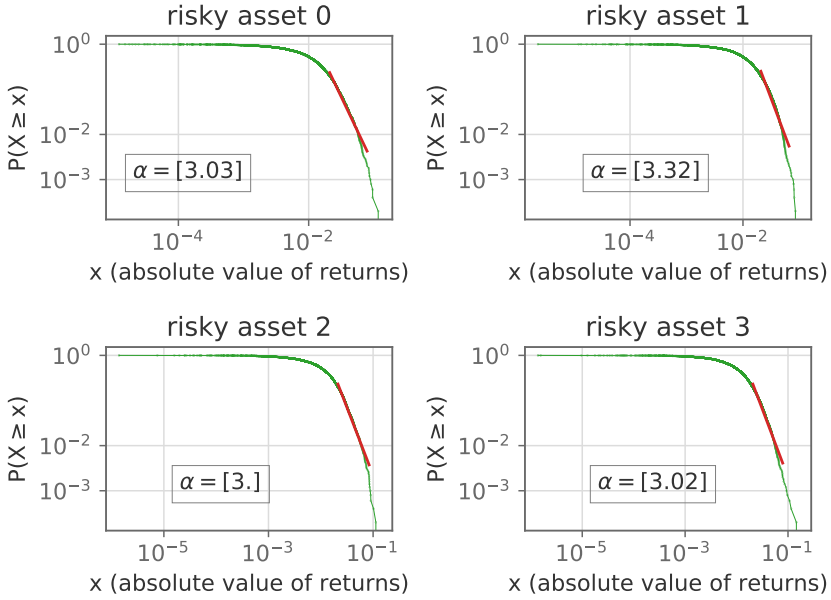

Figure 4.17: Log-log plot of the complementary cumulative distribution functions of the absolute returns of the four risky assets from the simulation presented in figure 4.15. The exponents are found fitting data from the last 20th percentile of the cumulative distribution, disregarding the largest ten values.

To test the validity of the time series generated by the simulation of our model, it is important to check if the model can reproduce some of these stylized facts.

Here we focus in particular on two of them, the fat-tailedness of assets' absolute returns and the long memory in the autocorrelation of the same quantities. The time series analyzed are the ones from the simulation presented in the previous section.

The fat-tailedness of absolute returns

In this section, we compare the decay of the distribution of assets' absolute returns resulting from the simulation of the market model to the observed leptokurtic behavior of their empirical counterparts (Cont, 2001, 2007). This leptokurtic trait can be understood from the point of view of extreme value theory. Indeed, the empirical distribution of absolute returns has fatter tails with respect to the Gaussian distribution. This means that rare events, represented by remarkably high or low returns (trading days characterized by booms or crashes), happen more frequently with respect to what the standard Normal distribution would predict. The relatively frequent presence of bubbles and crashes represents one distinctive feature of the financial markets.

The empirical fat-tail decay of the distribution

$$p(x) \sim x^{-1-\alpha} \tag{4.202}$$

is characterized by an exponent α in the range $(2, 4)$. As shown in figure 4.17, the fitted parameter from the simulated time series shown in figure 4.15 falls in this range of values.

The long memory in the autocorrelation of absolute returns

The daily returns are not independent random variables. In the financial markets periods of tranquility alternate to periods of high volatility. A common observation characterizing the financial time series is the concentration of high price's exuberance in definite time windows (Cont, 2001, 2007).

This clustering phenomenon goes under the name of volatility clustering. Its presence can be analyzed by looking for patterns in the time evolution of the absolute returns. An effective tool to accomplish this task is constituted by the autocorrelation function. Indeed, the presence

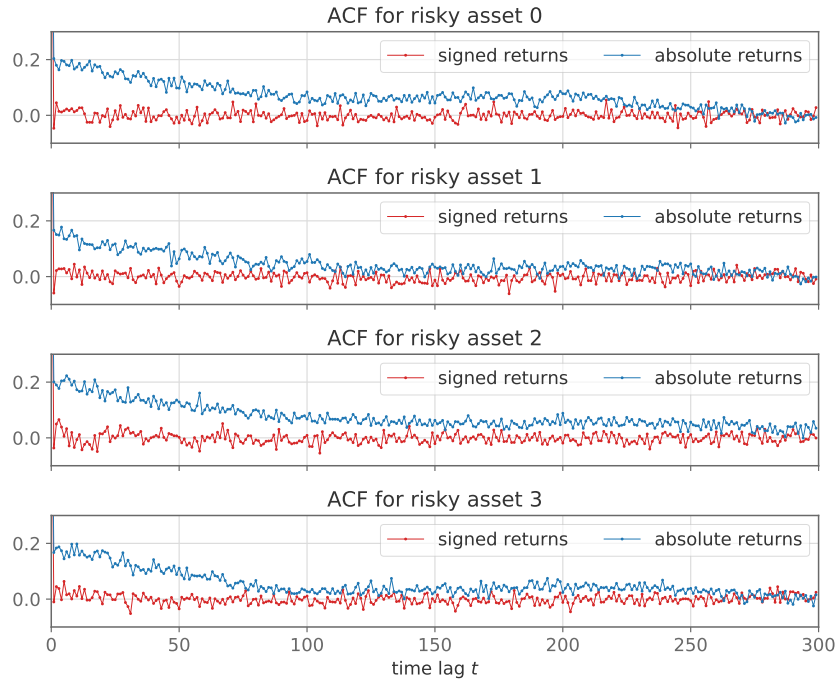

Figure 4.18: Autocorrelation function of signed and absolute returns of the four risky assets from the simulation presented in figure 4.15. The autocorrelation functions are computed for the data after the 500th trading day in order to exclude possible misleading contributions due to the initial conditions.

of long memory in the autocorrelation function is an indication of time inhomogeneity in the time distribution of the returns, characteristic of the volatility clustering phenomenon. The signed returns

$$\vec{r}_{p,t} = \left(\frac{P_{1,t}}{P_{1,t-1}} - 1, \ldots, \frac{P_{n,t}}{P_{n,t-1}} - 1 \right). \tag{4.203}$$

are empirically characterized by a fast-decaying autocorrelation, while the absolute returns $|\vec{r}_{p,t}|$ have instead an autocorrelation with a longer memory. In figure 4.18 we check if our ABM is able to reproduce this empirical fact. Indeed, the emerging autocorrelation functions present manifestly this behavior.

4.2.4 Emergence and time synchronization of bubbles among the risky assets

In this section, we employ the extended market model endowed with the $O(n)$ noise traders class to investigate the phenomenon of the emergence of bubbles and, in particular, their synchronous or asynchronous character.

Can we understand the mechanism triggering and governing the bubbles in the risky assets' prices? Will the assets develop faster-than-exponential growths in their prices at the same time or will the market be dominated by alternating peaks in different risky assets? Is there a correlation in the return time series of different risky assets and, if it exists, can we understand how this endogenous correlation emerges?

The analysis of the time series shows the presence of three distinct regimes characterizing the dynamics of the market model. The three regimes depend on the range of values of the herding propensity κ and are related to the regimes of the underlying $O(n)$ model structure.

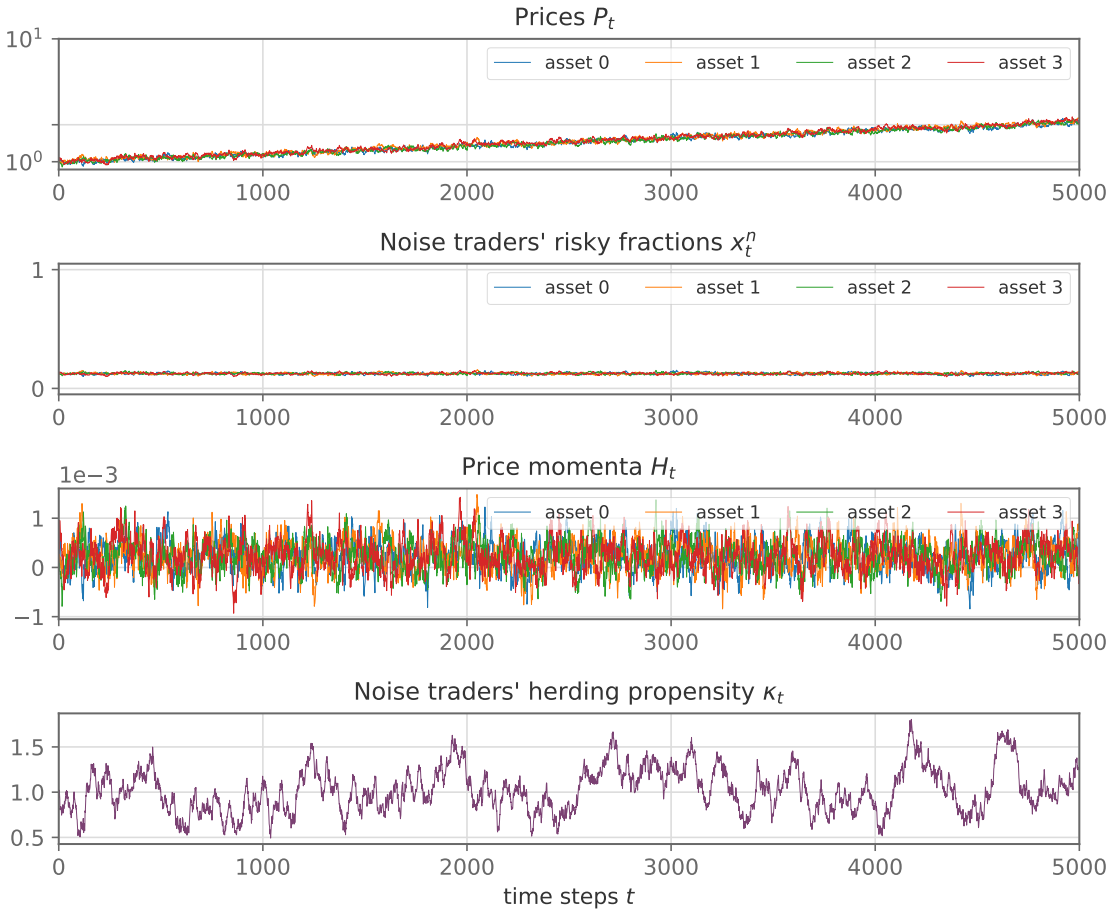

Figure 4.19: The figure presents the time series resulting from a simulation with four risky and one risk-free assets featuring an Ornstein-Uhlenbeck kappa process κ_t, characterized by a mean value $\mu_k = 0.98 \cdot 1$, far below the critical value $\kappa_c = 4$. The mean reversion strength $\eta = 0.013$ and the standard deviation $\sigma_\kappa = 0.25 \cdot 1 \cdot \sqrt{2\eta}$ are defined analogously to section 4.2.3, following Kaizoji et al. (2015). The other parameters coincide with the ones presented in table 4.5. The market is clearly in the disordered regime dominated by the idiosyncratic opinion. No bubble is present.

The subcritical market regime

The first regime corresponds to small values of κ far from the critical value $\kappa_c = n$. In this subcritical regime, the market model does not produce super-exponential bubbles, as evident in figure 4.19. This is true for both constant κ and Ornstein-Uhlenbeck κ_t processes, provided that the latter moves stochastically in a range of values far from the critical point.

The reason for this behavior has to be searched in the noise traders class. Indeed small values of κ correspond to a larger impact of the random component peculiar of each agent, with respect to the deterministic utility function common to all the traders. This is clear in formula (4.175), which we report here for convenience

$$\vec{B}^* = \underset{\vec{B}\in\mathbb{S}^{n-1}}{\arg\max}\{\kappa u_{\vec{B}} + \eta_{\vec{B}}\}. \tag{4.204}$$

A small herding propensity represents a noise traders class dominated by the random idiosyncratic opinion, hence the polarization of the class is not present or too small to exert any relevant effect on the market. There is no emerging collective phenomenon of the investors capable of resulting in super-exponential growths of the prices.

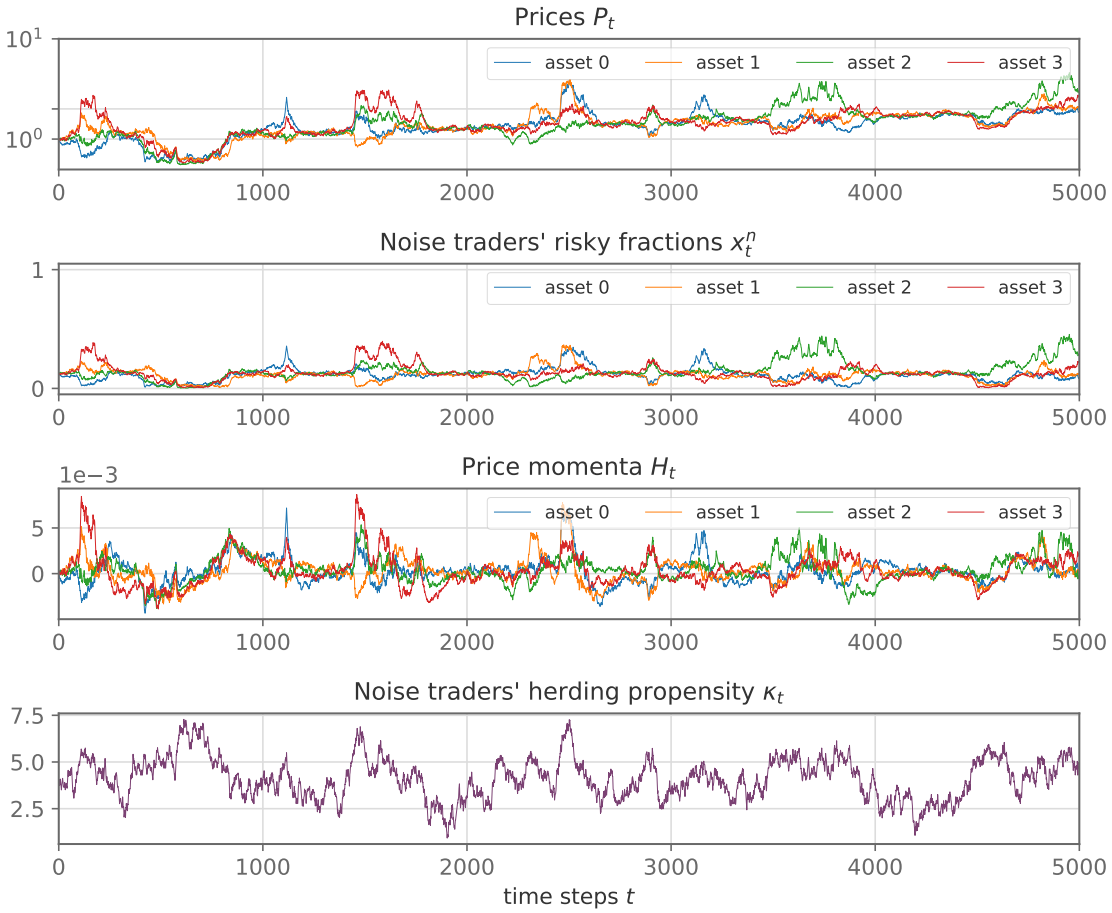

Figure 4.20: The figure presents the time series resulting from a simulation with four risky and one risk-free assets featuring an Ornstein-Uhlenbeck process κ_t, characterized by a mean value $\mu_k = 0.98\kappa_c$, with $\kappa_c = 4$. The market is simulated with the same set of parameters used for the simulation in figure 4.15, see table 4.5 and section 4.2.3. Tranquil periods of steady exponential growth of the prices alternate with volatile periods dominated by the emergence of super-exponential bubbles, mostly asynchronous in time.

The critical market regime

The second regime corresponds to values of the herding propensity near the critical one $\kappa \approx \kappa_c$. In this regime we observe the emergence of clear bubbles if the simulation is characterized by a time-varying Ornstein-Uhlenbeck kappa process κ_t, stochastically fluctuating near the critical threshold, as presented in figure 4.20.

Phase transition and emergence of bubbles

It is the noise traders class that triggers the emergence of bubbles and governs their dynamics. The social imitation attitude of these investors is able to trigger a polarization phenomenon of their investment decisions, constituted by an actual phase transition of the underlying $O(n)$ model. This also motivates the critical value $\kappa_c = n$, characteristic of a fully connected $O(n)$ model.

Despite the noise traders' positive feedback on the price momentum and their interaction with the fundamentalists, their strategy clearly exhibits the underlying statistical properties of the $O(n)$ model. Figure 4.21 shows that the average opinion $\frac{\sum_i \vec{s}_i}{N}$ of the noise traders class presents the well-known picture of a second-order phase transition, in turn characterizing the $O(n)$ model.

This interaction-driven phase transition is controlled by the time-varying herding propensity

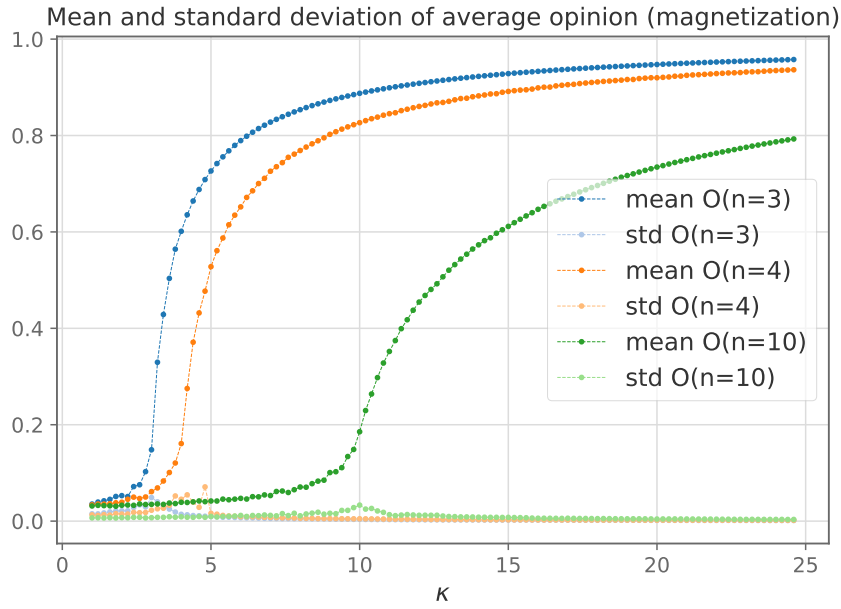

Figure 4.21: Mean and standard deviation over all the trading days of the average opinion of the noise traders class (norm of the average spin vector of that trading day $\|\frac{\sum_i \vec{S}_i^t}{N}\|$). The figure reports the values from the constant kappa simulations of three different market model characterized by a different number of risky assets $n = 3, 4, 10$ and hence by three different $O(n)$ models.

κ_t, which stochastically enters and exits the ordered regime, triggering the bubbles. This effect is further amplified by the price momentum H_t.

When the bubbles develop in one or more risky assets, the price momenta of these assets increase following the price dynamics, pushing more and more noise traders to invest into them, hence creating a self-reinforcing loop.

When the herding propensity reverts to the sub-critical regime, the polarization of the noise traders class starts to decrease, the idiosyncratic opinion starts to gain back importance and the noise traders start to move to the other assets. This selling phenomenon decreases the prices of the assets towards which the investments were polarized, as an effect the respective price momenta become negative pushing more and more traders to sell the falling prices' assets. The bubbles burst and the prices return to their fundamental values or overshoot below the fundamental value before returning to it. The mechanism just explained characterizes the positive bubbles, nevertheless, in an analogous way, the noise traders can also polarize towards the risk-free asset, creating a negative bubble. These alternating regimes of pessimistic and exuberant mood are described for example in Sornette (2017) and R. J. Shiller (2006).

From this analysis, we can also motivate our choice of the detailed balance condition in deriving the stochastic dynamics of the noise traders class, in contrast to the more general global balance rule. Indeed choosing the detailed balance condition corresponds to restricting all the sources of non-equilibrium to the stochastic wandering of the herding propensity parameter κ_t. Since we are interested in the out-of-equilibrium effects deriving from a change in herding propensity κ_t of the noise traders, in turns modeling the changes in the geopolitical and economical situation, we assume the system is constantly pushed out of equilibrium solely by the parameter κ_t.

The asynchronous bubbles in the critical regime

In this regime, the market model features tranquil periods with subcritical values of κ_t, where all the prices stay close to their fundamental values characterized by steady exponential growth. These alternate with volatile periods with values of the herding propensity above the critical threshold, which are instead dominated by the emergence of super-exponential bubbles in all the risky assets.

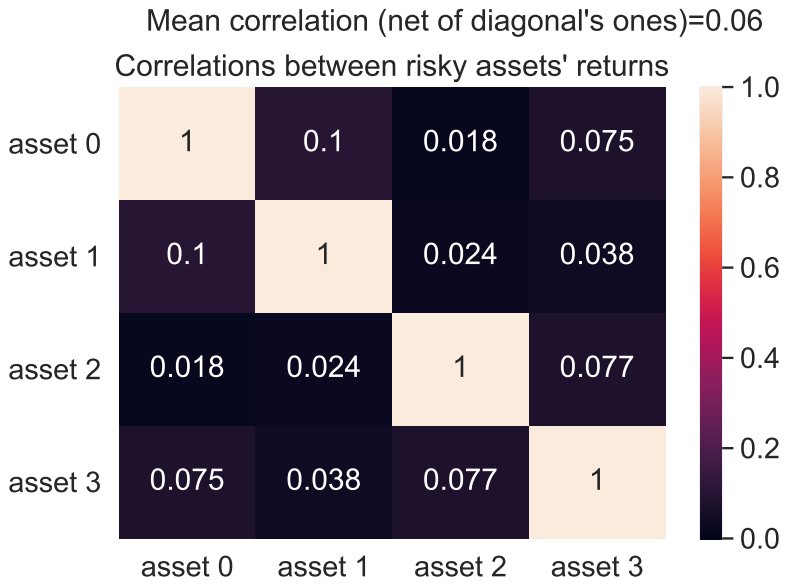

Figure 4.22: Spearman's rank correlation coefficients between the prices' returns from the simulation presented in figure 4.20.

The bubbles originating in the latter periods are mostly asynchronous among the different assets as observable at a qualitative level from figure 4.20. Moving to the quantitative ground, the lack of time synchronization is confirmed by the small values of the correlations among the assets' returns presented in figure 4.22. We use Spearman's rank correlation coefficients to quantify the relation between the returns because it measures the monotonic relationship between the variables. Pearson correlation coefficients only measure linear relations and are not suitable for random variables with fat-tails such as the return of this time-series. Neither the fundamentalists' strategy nor the noise traders' one introduces correlation among the asset. In the case of the fundamentalists, this is clear from the equation governing their portfolio allocation strategy (4.151) and from the fact that in these simulations we always assume that no correlation is expected between the risky assets by this type of agents

$$C^f = \begin{pmatrix} 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{pmatrix} \tag{4.205}$$

Clearly, non-zero values in this matrix would introduce some correlation in the prices. Studying the impact of the expected correlations by the fundamentalists on the realized correlations constitutes an interesting further direction of analysis. Nevertheless, here we are mainly interested in the correlations introduced by the noise traders, hence we always set to zero the non-diagonal elements of matrix (4.205). For the latter agents, we know that the $O(n)$ model does not introduce correlation among different components of the spin vector, hence explaining the asynchronous character of the emergence of the bubble in this regime.

The supercritical market regime

The third and last regime is instead characterized by a large value of the herding propensity, far above κ_c. The noise traders class is deeply inside the ordered phase, being dominated by a common investment preference. Remarkably, also this regime presents bubbles in the prices. The bubbles

Figure 4.23: The figure presents the time series resulting from a simulation with four risky and one risk-free assets featuring an Ornstein-Uhlenbeck process κ_t, characterized by a mean value $\mu_k = 0.98 \cdot 20$, far above the critical value $\kappa_c = 4$. The noise traders class is completely polarized. The mean reversion strength $\eta = 0.013$ and the standard deviation $\sigma_\kappa = 0.25 \cdot 20 \cdot \sqrt{2\eta}$ are defined analogously to section 4.2.3, following (Kaizoji et al., 2015). The other parameters coincide with the ones presented in table 4.5. A clear pattern emerges, with a first bubble in one asset's price which triggers a cascade of synchronous bubbles in the other prices.

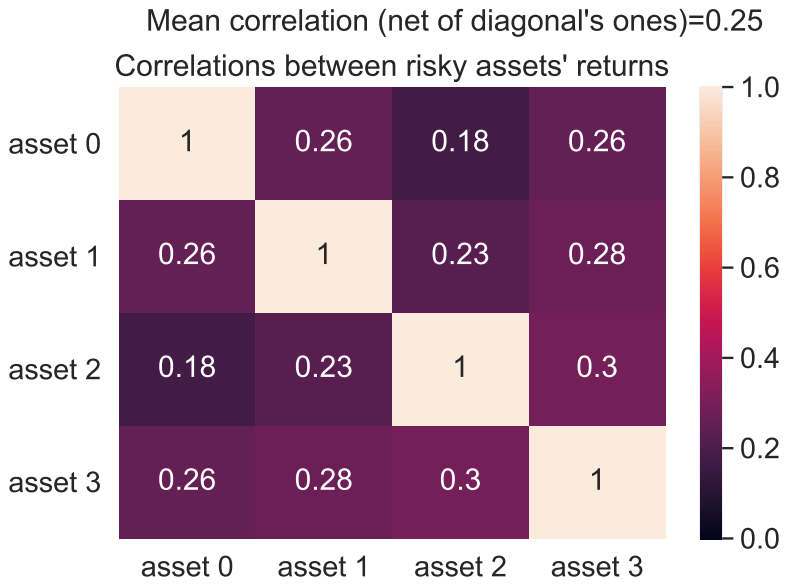

Figure 4.24: Spearman's rank correlation coefficients between the prices' returns from the simulation presented in figure 4.23.

characterizing these large values of herding propensity are of an intrinsically different nature with respect to the bubbles of the critical regime, these last resulting from the phase transition of the underlying $O(n)$ model. Moreover, at odds with the second regime, we observe synchronous bubbles, which are mainly driven by the trend-following attitude of the traders.

We present in figure 4.23 the time series from a setup with an Ornstein-Uhlenbeck κ_t wandering around a mean reversion level $\mu_k = 0.98 \cdot 20$, deeply inside the ordered supercritical phase. An important point is that the phenomena we are going to analyze in this regime do not depend on the nature of the kappa process. Indeed, both constant kappa and Ornstein-Uhlenbeck one lead to the same situation, provided κ, being time-varying or not, is always deep inside the ordered phase. At odds with the critical regime with $\kappa_t \approx \kappa_c$, the time-varying nature of the herding propensity is no more important being only the supercritical nature of noise traders class the very element governing this regime.

Two questions arise. First, what is the mechanism governing this type of bubbles since we know that the explanation formulated for the bubbles near criticality does no longer apply? Second, can we understand how this mechanism introduces the time synchronization feature in the emergence of the bubbles?

Mexican hat potential and bubbles in the supercritical regime

The supercritical $O(n)$ model is characterized by the well-known Mexican hat potential. The continuous phase transition from disorder to order of the $O(n)$ model is characterized by the smooth deformation of the paraboloid, constituting the subcritical potential of the system, into the Mexican hat. The single minimum located in the origin for the paraboloid, transform into a degenerate valley of minima in the supercritical case. This phase transition is the underlying mechanism triggering and governing the bubbles in the second regime.

In the supercritical regime, the situation is different and it is the very presence of the degenerate valley of minima to govern the bubbles. Indeed, even if we are far from the critical point, the system is characterized by a diverging susceptibility in the directions perpendicular to the non-zero magnetization vector. The noise traders class reacts with a collective behavior in response to small changes in the external field of price momenta. The price momenta can tilt the common investments' preferences of the class at the "macro" level.

This emerging collective behavior governed by the price momenta is ultimately the mechanism governing the super-exponential bubbles in this regime. Now it is clear why we stated that in the critical regime the bubbles are governed by the social imitation attitude, while in the supercritical one they are dominated by the trend-following attitude. In the first case, it is the transition to the ordered phase triggering the bubbles, in the second instead the class is already polarized and it is the tilting effect of the external field to drive them. Indeed, these are fundamentally different mechanisms. We now move to the second question, why are the resulting bubbles synchronous?

The synchronous bubbles in the supercritical regime

We start presenting in figure 4.24, the realized returns' correlations from the time series in figure 4.23. The Spearman correlation coefficients are between 0.18 and 0.3, which is significantly larger than the correlations in the previous section. For each of the three bubbles, the asset that starts to grow the bubble is not correlated to the others. The remaining three assets are positively correlated during the bubble. The situation is fundamentally different from the second regime's. The bubbles are indeed mainly synchronous and the average positive correlations quantify this behavior.

The synchronous character of their appearance is unexpected. Indeed, on one hand, we know that the fundamentalists' strategy cannot introduce correlation among the prices since the expected correlations are set to zero and hence the system (4.151) defining their allocation strategy is constituted by independent equations for the different risky fractions. Moreover, we again explicitly state for clarity that as assumed throughout the present work no correlation has been introduced between the dividends processes.

Furthermore, we know that the components of the spin vector are uncorrelated, hence neither the noise traders' strategy can introduce correlation among the prices. Then, where does the correlation come from?

There is only one component of the model we have not considered, the price equations. Indeed, the price equations are a set of non-linear coupled equations for the prices. The correlation can only come from the coupled character of these equations. Yet, the complex structure of the non-linear system makes an analytical study of the correlation introduced by the coupling between the equations unfeasible.

Nonetheless, the large κ regime can give us a clue in understanding the origin of the positive correlation. Indeed, for large values of the herding propensity, a clear pattern emerges in the time series, as represented in figure 4.23.

The first bubble develops in one asset, triggered by a stochastic fluctuation and then amplified by the positive feedback of the momentum. Its price starts to grow super-exponentially. Then on average the investors, both fundamentalist and noise traders, get richer. Since their strategy leads to a different response to the bubble event, their increase in wealth is different. The noise traders, pushed by the social imitation and trend following attitudes, invest more on the asset undergoing the bubble the more its price rises. The fundamentalist instead decreases their exposition on the asset the more its price rises and reallocate their wealth to the other assets, following their risk-averse strategy. Hence is understandable how the noise traders get richer with respect to the fundamentalist investors during the bubbles. Nevertheless, both the investors get richer when a bubble develops.

Since the traders' wealth increases, in order for the Walrasian equilibrium to be still satisfied, the other assets' prices must increase. This is evident from the market clearing condition characterizing the Walrasian equilibrium

$$\Delta D^{f,k}_{t-1\to t} + \Delta D^{n,k}_{t-1\to t} = 0 \quad \forall k \tag{4.206}$$

where $\Delta D^{f,k}_{t-1\to t}$ and $D^{n,k}_{t-1\to t}$ represent respectively the aggregate excess demands of each group for risky asset k. Expressing the equation in a more explicit form

$$W^f_t x^f_{k,t} - W^f_{t-1} x^f_{k,t-1} \frac{P_{k,t}}{P_{k,t-1}} \qquad + W^n_t x^n_{k,t} - W^n_{t-1} x^n_{k,t-1} \frac{P_{k,t}}{P_{k,t-1}} = 0 \tag{4.207}$$

and assuming constant all the quantities not explicitly depending on the other assets' prices, in particular on the one undergoing the bubble, we observe that the effect of the bubble of a specific

asset on the other assets' price equations is constituted solely by an increase of both the wealth of the fundamentalist and noise traders.

The only effect, emerging from this differential analysis, of a bubble of asset $i \neq k$ on the price equation (4.207) of any other asset k is an increase of the quantities W_t^f and W_t^n. The only way in which the equilibrium equation could be still satisfied is that the price $P_{k,t}$ is larger than the previous time-step's one $P_{k,t-1}$. We have indeed shown that an increase in one price triggers a positively correlated increase in all the other prices. If the increase is strong enough to have a relevant impact on the price momenta associated with the other assets, as in the case of a super-exponential bubble in the supercritical regime, the increase of the price momenta triggers the emergence of synchronous bubbles in all the other risky assets, through the tilting mechanism explained before.

Summarizing, since the investment strategies of both the classes of traders do not depend on the magnitude of their total wealth but only on the characteristic of the assets, the increase of wealth from which the traders benefit during the super-exponential growth of a bubble, cascades into the synchronous emerging of bubbles in the whole risky assets market.

This explains the pattern in figures 4.23 and constitutes the mechanism at the origin of the positive correlation among the risky assets.

4.2.5 Conclusion

We derived a market model with two types of agents who trade n risky and one risk-free assets. The model is a multi-asset extension of an agent-based model with fundamentalists and noise traders introduced by Kaizoji et al. (2015). The fundamentalists allocate their portfolio according to a maximization of the expected utility under a CRRA utility function. Tliographyhe noise traders' investment decision is described by an $O(n)$ model in which the n components represent the different assets in which the traders can invest. This allows us to define realistic stochastic dynamics while having control over the statistical properties of the model. The price momenta influence the traders' investment decisions in the form of the external field. The stochastic dynamics is completely specified by a discrete-time Markov chain, defined by the possible states, and the transition rates among them. We derived rejection-less transition probabilities in order to describe a realistic behavior at the "micro" level of the single investor.

The price at each time-step is defined by the traders' demand and supply for each of the assets. The resulting price time series exhibits bubbles and crashes and reproduces several "stylized facts" of financial markets such as volatility clustering or fat-tails of the return distribution. The bubbles emerge when the noise traders polarize their opinions towards one or more assets describing an actual phase transition of the underlying $O(n)$ model.

Furthermore, the model was applied to understand the synchronous and asynchronous emergence of bubbles among multiple assets. Three regimes were found in which the mechanism responsible for creating bubbles differed. The regimes are defined by the inverse temperature representing the traders' propensity to herd. The critical value of the inverse temperature is n in the $O(n)$ model. In the disordered regime with a herding propensity smaller than n, no polarization and consequently no bubbles emerged. Around the critical value, with a herding propensity transiently fluctuating above the critical value, asynchronous bubbles emerged in various assets when phase transitions occurred. The bubbles are largely driven by the polarization of the traders' opinion. In the third regime, where the herding propensity is far above the critical value and the noise traders class is already polarized, synchronous bubbles still emerge even if not triggered by an actual phase transition. They are instead driven by the traders' momentum following. Small random price fluctuations can trigger the noise traders to herd into that asset. The overall increase in wealth makes the traders buy more of the other assets even if they keep the fraction invested into each asset constant. The increase in momentum initiates the noise traders to shift their wealth from the risk-free to the risky assets creating a time-synchronous bubble in all of the assets. The global trend that emerges from this analysis is that the synchronization of bubbles increases with the increase of the range of

value of the herding propensity parameter κ_t. When the tendency to herd of the investors is large, the risk of a crash spreads from one asset to all of the other assets by triggering the growth of synchronized bubbles. This creates systemic risk for an economy because the simultaneous crash of a whole financial market or asset class can severely threaten an economy.

The model can be extended in several directions. In the present work, we focused mostly on the derivation of the model, with the construction of an efficient algorithm to simulate it, together with the identification of relevant regimes and the study of the characteristic time series. A deeper analytical study of the theoretical properties of the model, absent in the present work, is of great interest. For example in the interesting case of the number of risky assets $n \to \infty$. Moreover, the more realistic case of non-uniform social imitation strength could be addressed, together with the exploration of the challenging field of the application of less constrained rules to derive the stochastic dynamics of the noise traders' investments from the $O(n)$ model, for example through the global balance rule (Bouchaud, 2013).

The realistic multi-asset market model that reproduces "stylized facts" of financial markets can be applied to test the market impact of various portfolio optimization strategies in future research. Furthermore, the role of contrarian traders modeled by a negative herding propensity in stabilizing markets or triggering crashes could be analyzed. The market model also provides a framework to test policies intended to decrease systemic risk or to prevent bubbles and crashes. It can help to quantify the effect bursting one bubble has on other assets.

Contributions

D.C. and R.W. contributed equally to this work.

References

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American economic review*, *84*(2), 406–411.

Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R., & Tayler, P. (1996). Asset pricing under endogenous expectations in an artificial stock market. *The economy as an evolving complex system II*, *27*.

Axelrod, R. (2006). Agent-based modeling as a bridge between disciplines. *Handbook of computational economics*, *2*, 1565–1584.

Ayres, R. U. (2020). A brief history of financial booms and bubbles. In *On capitalism and inequality* (pp. 113–123). Springer.

Baghestanian, S., Lugovskyy, V., & Puzzello, D. (2015). Traders' heterogeneity and bubble-crash patterns in experimental asset markets. *Journal of Economic Behavior & Organization*, *117*, 82–101.

Barro, R., & Kamphuis, R. (1988). The stock market and the macroeconomy, implications of the october 1987 crash. *Black Monday and the future of financial markets*.

Ben-Akiva, M., Litinas, N., & Tsunokawa, K. (1985). Continuous spatial choice: the continuous logit model and distributions of trips and urban densities. *Transportation Research Part A: General*, *19*(2), 119–154.

Bernanke, B., & Gertler, M. (2000). Monetary policy and asset price volatility. *National bureau of economic research*(No. w7559).

Bernanke, B. S., & Gertler, M. (2001). Should central banks respond to movements in asset prices? *American Economic Review*, *91*(2), 253–257.

Bernstein, M. A. (1987). *The great depression: delayed recovery and economic change in america, 1929-1939*. Cambridge University Press.

Bertocchi, M., Giacometti, R., & Zenios, S. A. (2005). Risk factor analysis and portfolio immunization in the corporate bond market. *European Journal of Operational Research*, *161*(2), 348–363.

Biswas, S., Hanson, A., & Phan, T. (2020). Bubbly recessions. *American Economic Journal: Macroeconomics*, *12*(4), 33-70.

Black, F. (1986). Noise. *The Journal of Finance*, *16*(3), 529–543.

Blot, C., Hubert, P., & Labondance, F. (2018). Monetary policy and asset price bubbles. *Sciences Po OFCE Working Paper, 2018-11-13*, *37*.

Borghesi, C., & Bouchaud, J.-P. (2007). Of songs and men: a model for multiple choice with herding. *Quality & quantity*, *41*(4), 557–568.

Bouchaud, J.-P. (2008). Economics needs a scientific revolution. *Nature*, *455*(7217), 1181–1181.

Bouchaud, J.-P. (2013). Crises and collective socio-economic phenomena: simple models and challenges. *Journal of Statistical Physics*, *151*(3-4), 567–606.

Boyd, S., Busseti, E., Diamond, S., Kahn, R., Koh, K., Nystrup, P., & Speth, J. (2016). Multi-period trading via convex optimization. *Foundations and Trend in Optimization*, *3*(1), 1-76.

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

Brock, W. A., & Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control*, *22*(8-9), 1235–1274.

Campbell, G., & Turner, J. (2010). 'the greatest bubble in history': Stock prices during the british railway mania. *MPRA*(21820).

Campbell, J. Y., Lo, A. W., MacKinlay, A. C., & Whitelaw, R. F. (1998). The econometrics of financial markets. *Macroeconomic Dynamics*, *2*(4), 559–562.

Campbell, R., Liu, Y., & Zhu, H. (2016). ...and the cross-section of expected returns. *The Review of Financial Studies*, *29*, 5-68.

Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of modern physics*, *81*(2), 591.

Cavagna, A., Garrahan, J. P., Giardina, I., & Sherrington, D. (1999). Thermal model for adaptive competition in a market. *Physical Review Letters*, *83*(21), 4429.

Cecchetti, S., Genberg, H., Lipsky, J., & Wadhwani, S. (2000). Asset prices and central bank policy. *Geneva Report on the World Economy*, *2*, CEPR and ICMB.

Cecchetti, S., Genberg, H., & Wadhwani, S. (2002). Asset prices in a flexible inflation targeting framework, in hunter, w.c., kaufman, g.g. and pomerleano, m. (eds), Asset Price Bubbles: The Implications for Monetary, Regulatory and International Policies, MIT Press. , 427-444.

Chakraborti, A., Toke, I. M., Patriarca, M., & Abergel, F. (2011). Econophysics review: Ii. agent-based models. *Quantitative Finance*, *11*(7), 1013–1041.

Challet, D., & Zhang, Y.-C. (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical Mechanics and its Applications*, *246*(3-4), 407–418.

Challet, D., & Zhang, Y.-C. (1998). On the minority game: Analytical and numerical studies. *Physica A: Statistical Mechanics and its applications*, *256*(3-4), 514–532.

Chiarella, C., Dieci, R., & Gardini, L. (2006). Asset price and wealth dynamics in a financial market with heterogeneous agents. *Journal of Economic Dynamics and Control*, *30*(9-10), 1755–1786.

Chiarella, C., Dieci, R., & He, X.-Z. (2007). Heterogeneous expectations and speculative behavior in a dynamic multi-asset framework. *Journal of Economic Behavior & Organization*, *62*(3), 408–427.

Chiarella, C., Dieci, R., & He, X.-Z. (2009). *Heterogeneity, market mechanisms, and asset price dynamics*. Elsevier.

Cieslak, A., & Povala, P. (2016). Information in the term structure of yield curve volatility. *The Journal of Finance*, *71*(3), 1393–1436.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, *1*, 223–236.

Cont, R. (2007). Volatility clustering in financial markets: empirical facts and agent-based models. In *Long memory in economics* (pp. 289–309). Springer.

Coolen, A. C. (2005). *The mathematical theory of minority games: Statistical mechanics of interacting agents (oxford finance series)*. Oxford University Press, Inc.

Cox, J. C., Ingersoll Jr, J. E., & Ross, S. A. (2005). A theory of the term structure of interest rates. , 129–164.

Dai, Q., & Singleton, K. J. (2002). Expectation puzzles, time-varying risk premia, and affine models of the term structure. *Journal of financial Economics*, *63*(3), 415–441.

Damiani, E. (2019). Equilibrium model of fundamentalist and noise traders in a multi-asset framework. *Master's thesis, Politecnico di Torino*.

Dawood, M., Horsewood, N., & Strobel, F. (2017). Predicting sovereign debt crises : An early warning system approach. *Journal of Financial Stability*, *28*, 16–28.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Positive Feedback Investment Strategies and Destabilizing Rational Speculation. *The Journal of Finance*, *45*(2), 379–395.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (2007). Noise trader risk in financial markets. *The Journal of Political Economy*, *98*(4), 703–738.

De Grauwe, P., Dewachter, H., & Embrechts, M. (1995). *Exchange rate theory: chaotic models of foreign exchange markets*. Blackwell.

De Long, J. B., & Magin, K. (2006). *A short note on the size of the dot-com bubble* (Tech. Rep.). National Bureau of Economic Research.

Demirer, R., Demos, G., Gupta, R., & Sornette, D. (2019). On the predictability of stock market bubbles: evidence from lppls confidence multi-scale indicators. *Quantitative Finance*, *19*(5), 843-858.

Demos, G., & Sornette, D. (2019). Comparing nested data sets and objectively determining financial bubbles' inceptions. *Physica A: Statistical Mechanics and its Applications*, *524*, 661–675.

Dieci, R., & He, X. Z. (2018). *Heterogeneous agent models in finance* (Vol. 4). Elsevier B.V.

Drehmann, M., & Juselius, M. (2014). Evaluating early warning indicators of banking crises : Satisfying policy requirements. *International Journal of Forecasting*, *30*, 759–780.

Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *The Journal of Finance*, *57*(1), 405–443.

Duffie, D., & Kan, R. (1996). A yield-factor model of interest rates. *Mathematical finance*, *6*(4), 379–406.

Duffy, J., & Ünver, M. U. (2006). Asset price bubbles and crashes with near-zero-intelligence traders. *Economic theory*, *27*(3), 537–563.

Eckrot, A., Jurczyk, J., & Morgenstern, I. (2016). Ising model of financial markets with many assets. *Physica A: Statistical Mechanics and its Applications*, *462*, 250–254.

Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, *25*(2).

Fama, E. (1991). Efficient capital markets: Ii. *Journal of Finance*, *46*(5), 1575-1617.

Farmer, J. (2002). Market forces, ecology and evolution. *Indust. Corp. Change*, *11*, 895-953.

Farmer, J. D., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, *460*(7256), 685–686.

Fedyk, Y., Heyerdahl-Larsen, C., & Walden, J. (2013). Market selection and welfare in a multi-asset economy. *Review of Finance*, *17*(3), 1179–1237.

Filimonov, V., & Sornette, D. (2013). A stable and robust calibration scheme of the log-periodic power law model. *Physica A: Statistical Mechanics and its Applications*, *392*(17), 3698-3707.

Forró, Z., Woodard, R., & Sornette, D. (2015). Using trading strategies to detect phase transitions in financial markets. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, *91*(4), 1–5.

Frankel, J. A., & Froot, K. A. (1986). *The dollar as speculative bubble: a tale of fundamentalists and chartists* (Tech. Rep.). National Bureau of Economic Research.

Galam, S., Gefen, Y., & Shapir, Y. (1982). Sociophysics: A new approach of sociological collective behaviour. i. mean-behaviour description of a strike. *Journal of Mathematical Sociology*, *9*(1), 1–13.

Gali, J., Giusti, G., & Noussair, C. (2020). Monetary policy and asset price bubbles: A laboratory experiment. *Barcelona GSE Working Paper Series Working Paper n 1184, May*.

Garber, P. M. (1989). Tulipmania. *Journal of political Economy*, *97*(3), 535–560.

Garber, P. M. (1990). Famous first bubbles. *Journal of Economic perspectives*, *4*(2), 35–54.

Gerlach, J.-C., Demos, G., & Sornette, D. (2019). Dissection of bitcoin's multiscale bubble history from january 2012 to february 2018. *Royal Society open science*, *6*(7), 180643.

Gordon, R. J. (2005). *The 1920s and the 1990s in mutual reflection* (Tech. Rep.). National Bureau of Economic Research.

Grabowski, A., & Kosiński, R. (2006). Ising-based model of opinion formation in a complex network of interpersonal interactions. *Physica A: Statistical Mechanics and its Applications*,

361(2), 651–664.

Greenspan, A. (2004). Risk and uncertainty in monetary policy. *Remarks at the Meetings of the American Economic Association, San Diego, California, January 3rd*.

Grinold, R. C., & Kahn, R. N. (2000). Active portfolio management.

Gualdi, S., Bouchaud, J.-P., Cencetti, G., Tarzia, M., & Zamponi, F. (2015). Endogenous crisis waves: Stochastic model with synchronized collective behavior. *Physical Review Letters*, *114*, 088701.

Gualdi, S., Tarzia, M., Zamponi, F., & Bouchaud, J.-P. (2015). Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, *50*, 29-61.

Gürkaynak, R. S. (2008). Econometric tests of asset price bubbles: taking stock. *Journal of Economic surveys*, *22*(1), 166–186.

Harras, G., Tessone, C., & Sornette, D. (2012). Noise-induced volatility of collective dynamics. *Phys. Rev. E*, *85*, 011150.

Haruvy, E., & Noussair, C. N. (2006). The effect of short selling on bubbles and crashes in experimental spot asset markets. *The Journal of Finance*, *61*(3), 1119–1157.

Heath, D., Jarrow, R., & Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, 77–105.

Hegselmann, R. (2017). Thomas C. Schelling and James M. Sakoda: The intellectual, technical, and social history of a model. *Journal of Artificial Societies and Social Simulation*, *20*(3), 1-152.

Holley, R. A., & Liggett, T. M. (1975). Ergodic theorems for weakly interacting infinite systems and the voter model. *The annals of probability*, 643–663.

Hommes, C., & LeBaron, B. (2018). *Computational economics: Heterogeneous agent modeling* (Vol. 4). Elsevier B.V.

Hommes, C. H., & Wagener, F. (2009). Bounded rationality and learning in complex markets. In *Handbook of economic complexity* (pp. 87–123).

Huang, Y., Miao, J., & Wang, P. (2019). Saving China's Stock Market? *IMF Economic Review*, *67*(2), 349–394.

Hwang, S., & Rubesam, A. (2015). The disappearance of momentum. *The European Journal of Finance*, *21*(7), 584-607.

Ide, K., & Sornette, D. (2002). Oscillatory finite-time singularities in finance, population and rupture. *Physica A: Statistical Mechanics and its Applications*, *307*(1-2), 63-106.

Ikeda, D. (2019). Monetary policy, inflation and rational asset price bubbles. *SSRN*(2938298).

Jiang, Z. Q., Zhou, W. X., Sornette, D., Woodard, R., Bastiaensen, K., & Cauwels, P. (2010). Bubble diagnosis and prediction of the 2005-2007 and 2008-2009 Chinese stock market bubbles. *Journal of Economic Behavior and Organization*, *74*(3), 149–162.

Johansen, A., & Sornette, D. (2002). Large stock market price drawdowns are outliers. *Journal of Risk*, *4*(2), 69-110.

Johansen, A., & Sornette, D. (2010). Shocks, crashes and bubbles in financial markets. *Brussels Economic Review (Cahiers economiques de Bruxelles)*, *53*(2), 201-253.

Johansen, A., Sornette, D., & Ledoit, O. (1999). Predicting financial crashes using discrete scale invariance. *Journal of Risk*, *1*(4), 5–32.

Johnson, N. F., Hui, P. M., Jonson, R., & Lo, T. S. (1999). Self-organized segregation within an evolving population. *Physical Review Letters*, *82*(16), 3360.

Juster, F. T., Lupton, J. P., Smith, J. P., & Stafford, F. (2006). The decline in household saving and the wealth effect. *Review of Economics and statistics*, *88*(1), 20–27.

Kaizoji, T., Leiss, M., Saichev, A., & Sornette, D. (2015). Super-exponential endogenous bubbles in an equilibrium model of fundamentalist and chartist traders. *Journal of Economic Behavior and Organization*, *112*(June 2013), 289–310.

Kaldor, N. (1961). *Capital accumulation and economic growth*. Springer.

Kim, G.-r., & Markowitz, H. M. (1989). Investment rules, margin, and market volatility. *Journal of Portfolio Management*, *16*(1), 45.

Kindleberger, C. P., & Aliber, R. Z. (2011). *Manias, panics and crashes: a history of financial crises*. Palgrave Macmillan.

Kirman, A. (2012). Can artificial economics help us unterstand real economies? *Revue de l'OFCE*, *5*, 15–41.

Kopp, A. (2020). Equilibrium model of a fixed income market with fundamentalist and chartist traders. *Master's thesis, ETH Zürich*.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, *53*(6), 1315–1335.

Lebergott, S. (1957). Annual estimates of unemployment in the united states, 1900-1954. In *The measurement and behavior of unemployment* (pp. 211–242). NBER.

Lempérière, Y., Deremble, C., Seager, P., Potters, M., & Bouchaud, J.-P. (2014). Two centuries of trend following. *The Journal of Investment Strategies*, *3*(3), 41-61.

Lera, S., & Sornette, D. (2017). Evidence of a bimodal US GDP growth rate distribution: A wavelet approach. *Quantitative Finance and Economics*, *1*(1), 26-43.

Li, Y., Riolo, R., & Savit, R. (2000a). Evolution in minority games.(i). games with a fixed strategy space. *Physica A: Statistical Mechanics and its Applications*, *276*(1-2), 234–264.

Li, Y., Riolo, R., & Savit, R. (2000b). Evolution in minority games.(ii). games with variable strategy spaces. *Physica A: Statistical Mechanics and its Applications*, *276*(1-2), 265–283.

Ljungqvist, L., & Sargent, T. J. (2018). *Recursive macroeconomic theory*. MIT press.

Lux, T., & Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, *397*(6719), 498–500.

Lux, T., & Marchesi, M. (2000). Volatility clustering in financial markets: a microsimulation of interacting agents. *International journal of theoretical and applied finance*, *3*(04), 675–702.

Lux, T., & Westerhoff, F. (2009). Economics crisis. *Nature Physics*, *5*(1), 2–3.

Majewski, A. A., Ciliberti, S., & Bouchaud, J.-P. (2020). Co-existence of trend and value in financial markets: Estimating an extended chiarella model. *Journal of Economic Dynamics and Control*, *112*, 103791.

Marrison, C. (2002). *The fundamentals of risk measurement*. McGraw Hill Professional.

McFadden, D. (1981). Econometric models of probabilistic choice. *Structural analysis of discrete data with econometric applications*, *198272*.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, *21*(6), 1087–1092.

Mishkin, F. (2011). How should central banks respond to asset-price bubbles? the 'lean' versus 'clean' debate after the gfc. *Bulletin of the Reserve Bank of Australia*, *June quarter*, 59-70.

Mishkin, F. S. (2011). Over the cliff: From the subprime to the global financial crisis. *Journal of Economic Perspectives*, *25*(1), 49–70.

Mishkin, F. S., & White, E. N. (2002). *Us stock market crashes and their aftermath: implications for monetary policy* (Tech. Rep.). National bureau of economic research.

Nakamoto, S. (2019). *Bitcoin: A peer-to-peer electronic cash system* (Tech. Rep.). Manubot.

Nelson, C. R., & Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of business*, 473–489.

Patzelt, F., & Pawelzik, K. R. (2012). Unstable price dynamics as a result of information absorption in speculative markets. *arXiv preprint arxiv:1211.6695*.

Powell, M. J. (1968). *A fortran subroutine for solving systems of nonlinear algebraic equations* (Tech. Rep.). Atomic Energy Research Establishment, Harwell, England (United Kingdom).

Powell, M. J. (1970). A hybrid method for nonlinear equations. *Numerical methods for nonlinear algebraic equations*.

Roehner, B., Sornette, D., & Andersen, J. V. (2004). Response functions to critical shocks in social sciences: An empirical and numerical study. *International Journal of Modern Physics C*,

15(06), 809–834.

Romer, C. D. (2003). Great depression. *Encyclopædia Britannica*, 225.

Rosser, J. B. (2000). *From catastrophe to chaos: A general theory of economic discontinuities: Mathematics, microeconomics and finance* (Vol. 1). Springer Science & Business Media.

Roubini, N. (2006). Why central banks should burst bubbles. *International Finance*, *9*(1), 87-107.

Sakoda, J. (1949). Minidoka: An analysis of changing patterns of social behavior. *PhD thesis, University of California*.

Sakoda, J. (1971). The checkerboard model of social interaction. *Journal of Mathematical Sociology*, *1*(1), 119-132.

Samanidou, E., Zschischang, E., Stauffer, D., & Lux, T. (2007). Agent-based models of financial markets. *Reports on Progress in Physics*, *70*(3), 406.

Satinover, J., & Sornette, D. (2007). " illusion of control" in time-horizon minority and parrondo games. *The European Physical Journal B*, *60*(3), 369–384.

Schatz, M., & Sornette, D. (2019). Inefficient bubbles and efficient drawdowns in financial markets. *Swiss Finance Institute Research Paper*(18-49).

Schelling, T. C. (1971). Dynamics Model of Segregation. *Journal of Mathematical Sociology*, *1*(2), 143–186.

Schmitt, N., & Westerhoff, F. (2017). On the bimodality of the distribution of the s&p 500's distortion: Empirical evidence and theoretical explanations. *Journal of Economic Dynamics and Control*, *80*, 34–53.

Sharpe, W. F. (1994). The Sharpe Ratio. *The Journal of Portfolio Management*, *21*(1), 49–58.

Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, *71*, 421-436.

Shiller, R. J. (2006). Irrational exuberance, 2nd edition. *Crown Business*.

Shiller, R. J. (2012). *The subprime solution: How today's global financial crisis happened, and what to do about it*. Princeton University Press.

Shin, H. S. (2013). Procyclicality and the search for early warning indicators. *International Monetary Fund*, *13*(258).

Smith, D. J. (2014). *Bond math,+ website: The theory behind the formulas*. John Wiley & Sons.

Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica: Journal of the Econometric Society*, 1119–1151.

Sornette, D. (1994). Sweeping of an instability: an alternative to self-organized criticality to get powerlaws without parameter tuning. *Journal de Physique I*, *4*(2), 209–221.

Sornette, D. (2003). *Why stock market crash: critical events in complex financial systems* (No. 1). Princeton University Press.

Sornette, D. (2009). Dragon-kings, black swans and the prediction of crises. *International Journal of Terraspace Science and Engineering*, *2*(1), 1-18 (`http://ssrn.com/abstract=1470006`).

Sornette, D. (2014). Physics and financial economics (1776-2014): Puzzles, ising and agent-based models. *Reports on Progress in Physics*, *77*(6), 1–76.

Sornette, D. (2017). *Why stock markets crash: critical events in complex financial systems* (Vol. 49). Princeton University Press.

Sornette, D., & Cauwels, P. (2014). 1980–2008: The illusion of the perpetual money machine and what it bodes for the future. *Risks*, *2*(2), 103–131.

Sornette, D., & Cauwels, P. (2015a). Financial bubbles: mechanisms and diagnostics. *Review of Behavioral Economics*, *2*(3), 279-305.

Sornette, D., & Cauwels, P. (2015b). Managing risk in a creepy world. *Journal of Risk Management in Financial Institutions*, *8*(1), 83–108.

Sornette, D., Demos, G., Zhang, Q., Cauwels, P., & Zhang, Q. (2015). Real-time prediction and post-mortem analysis of the shanghai 2015 stock market bubble and crash. *The Journal of*

Investment Strategies, *4*(4), 77–95.

Sornette, D., & Johansen, A. (1997). Large financial crashes. *Physica A: Statistical Mechanics and its Applications*, *245*(3-4), 411–422.

Sornette, D., & Johansen, A. (2001). Significance of log-periodic precursors to financial crashes. *Quantitative Finance*, *1*(4), 452-471.

Sornette, D., Johansen, A., & Bouchaud, J.-P. (1996). Stock market crashes, precursors and replicas. *J. Phys. I France*, *6*(1), 167-175.

Sornette, D., & Ouillon, G. (2012). Dragon-kings: mechanisms, statistical methods and empirical evidence. *Eur. Phys. J. Special Topics*, *205*, 1-26.

Sornette, D., Wheatley, S., & Cauwels, P. (2019). The fair reward problem: the illusion of success and how to solve it. *arXiv preprint arxiv:1902.04940*.

Sornette, D., & Woodard, R. (2010). Financial bubbles, real estate bubbles, derivative bubbles, and the financial and economic crisis. In *Econophysics approaches to large-scale business data and financial crisis* (pp. 101–148). Springer.

Stigler, G. J. (1964). Public regulation of the securities markets. *The Journal of Business*, *37*(2), 117–142.

Su, Y., Yip, Y., & Wong, R. W. (2002). The impact of government intervention on stock returns evidence from Hong Kong. *International Review of Economics and Finance*, *11*(3), 277–297.

Taylor, N. (2014). The rise and fall of technical trading rule success. *Journal of Banking & Finance*, *40*, 286-302.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of financial economics*, *5*(2), 177–188.

Wadhwani, S. (2008). Should monetary policy respond to asset price bubbles? revisiting the debate. *National Institute Economic Review*, *206*, 25-34.

Walras, L. (1954). *Elements of pure economics*. Routledge.

Weidlich, W. (1971). The statistical description of polarization phenomena in society. *British Journal of Mathematical and Statistical Psychology*, *24*(2), 251–266.

Westphal, R., & Sornette, D. (2020a). How market intervention can prevent bubbles and crashes. *Swiss Finance Institute Research Paper*(12-74).

Westphal, R., & Sornette, D. (2020b). Market Impact and Performance of Arbitrageurs of Financial Bubbles in An Agent-Based Model. *Journal of Economic Behavior and Organization*, *171*, 1–23.

Wheatley, S., Sornette, D., Huber, T., Reppen, M., & Gantner, R. N. (2019). Are bitcoin bubbles predictable? combining a generalized metcalfe's law and the log-periodic power law singularity model. *Royal Society open science*, *6*(6), 180538.

Wyart, M., & Bouchaud, J.-P. (2007). Self-referential behaviour, overreaction and conventions in financial markets. *Journal of Economic Behavior & Organization*, *63*(1), 1–24.

Xu, H.-C., Zhang, W., Xiong, X., & Zhou, W.-X. (2014). Wealth share analysis with "fundamentalist/chartist" heterogeneous agents. In *Abstract and applied analysis* (Vol. 2014).

Zhang, Q., Sornette, D., Balcilar, M., Gupta, R., Ozdemir, Z. A., & Yetkiner, H. (2016). LPPLS bubble indicators over two centuries of the S&P 500 index. *Physica A: Statistical Mechanics and its Applications*, *458*, 126–139.

Zhang, Q., Zhang, Q., & Sornette, D. S. (2016). Early warning signals of financial crises with multi-scale quantile regressions of log-periodic power law singularities. *PLoS ONE*, *11*(11).

Zhelezov, O. I. (2017). N-dimensional rotation matrix generation algorithm. *American Journal of Computational and Applied Mathematics*, *7*(2), 51–57.

Zhou, W. X., & Sornette, D. (2003). 2000-2003 real estate bubble in the uk but not in the usa. *Physica A: Statistical Mechanics and its Applications*, *329*(1-2), 249–263.

Zhou, W. X., & Sornette, D. (2006). Is there a real-estate bubble in the us? *Physica A: Statistical Mechanics and its Applications*, *361*(1), 297–308.

Zhou, W. X., Sornette, D., Hill, R. A., & Dunbar, R. I. (2005). Discrete hierarchical organization of social group sizes. *Proc. Royal Soc. London*, *272*, 439–444.