ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Effects of controlling or exploiting financial bubbles on market dynamics, in the framework of an agent-based model

Master Thesis

S. Drouard

September 1st, 2020

Advisors: Prof. Dr. D. Sornette, R. Westphal

Department of Management, Technology and Economics, ETH Zürich

Abstract

Speculative bubbles are abnormally fast surges in asset prices, followed by sharp descents, that cause tremendous stress on financial markets. Prevention of these events can be enforced by monetary authorities such as central banks. This thesis investigates the effects of asset price targeting based monetary policies on the transient super exponential growth of asset prices referred to as price bubbles. To that end, three strategies are studied in the framework of an agent-based model. This model, introduced in Kaizoji, Leiss, Saichev, and Sornette (2015), features a risk-free asset and a risky asset, as well as two types of traders, and can reproduce empirical characteristics of price bubbles. In this context, arbitrageurs using a bubble detection tool based on the exponential moving average of the returns to their own profit are shown in Westphal and Sornette (2020a) to reduce the occurrence of bubbles. This same detection tool is used in a model implementing open market operations, and a model featuring interest rate modification. A comparison is drawn among these applications. The agents of regulatory policy, representing central banks, use so-called lean against the wind policies to stabilize the price paths. Such price targeting policies are shown to have hindering effects on positive price bubbles, especially with early detection. In certain proportions, Sharpe ratio effective arbitrageurs are shown to be more efficient at preventing bubbles than monetary policies. Additionally, arguments are drawn towards some moderate positive correlation between optimizing arbitrageurs towards the highest reduction in the number of bubbles, and optimizing towards the highest Sharpe ratio of the strategy. This indicates that the agent-based model follows the invisible hand argument of Smith (1761) to some extent.

Acknowledgements

I would like to address my sincere appreciation and thanks to Prof. Dr. Didier Sornette for giving me the opportunity to write this Master's thesis at the Chair of Entrepreneurial Risks, for providing invaluable advice on my research process and bringing enlightening perspective to my work.
I also want to express my sincere gratitude to Rebecca Westphal, for her continuous guidance and support of my efforts, as well as abundant feedback throughout the course of the completion of this present thesis.
I also want to thank Prof. Dr. Farkas, the program director of the Quantitative Finance Master's at UZH and ETH, for his support of this research choice, as well as all the professors of this program for providing me with an understanding of quantitative finance in general.
Finally, I want to thank my family and friends for the continuous support they have provided throughout the course of my studies.

Contents

Contents **v**

Introduction **1**

1 Initial market model: two types of agents **9**
1.1 Market investing setup - assets 10
1.2 Fundamentalist traders 11
1.3 Noise traders 14
1.4 Market clearing 17
1.4.1 Market clearing conditions 18
1.4.2 Derivation of the price 18
1.5 Behavior of the two-asset market model 20

2 A three-agent market model: the Dragon Rider trader **29**
2.1 Description of the three-agent model 30
2.1.1 Assets 30
2.1.2 The Dragon Rider/Slayer strategy 31
2.2 Price derivation 36
2.3 Dragon Rider grid exploration 36

3 A second three-agent market model: a Dragon Slayer trader **53**
3.1 Description of the three-agent model 54
3.1.1 Market setup 54
3.1.2 Dragon Slayer strategy 54
3.1.3 Price derivation 59
3.2 Dragon Slayer results 59
3.2.1 Choosing the weight w 59
3.2.2 Wealth variation 59
3.2.3 Optimization of smoothing window m 63
3.2.4 Sharpe ratio distribution 63

3.2.5 Comparison with Dragon Rider 63

4 Regulation through interest rate modification **67**
4.1 Market setup 67
4.1.1 Basic market setup 68
4.1.2 Dragon Slayer scheme 68
4.2 Dragon Slayer results 72
4.2.1 Effects of varying a 73
4.2.2 Effects of varying w 73
4.2.3 Effects of varying m 73
4.2.4 Summary results 74

Conclusion **79**

A Appendix **81**

Bibliography **87**

Introduction

Our focus in this Master's thesis is to evaluate the effect of asset price targeting based monetary policies on financial bubbles, using a simple agent-based model of a financial market. We will now provide some background information on asset price bubbles and monetary policy, as well as background information on the use of agent-based models in finance.
Let us first specify what financial bubbles are. Positive asset price bubbles occur when the asset prices, as they are traded on the market, are suspected to be significantly above their intrinsic value. This behavior has been referred to as "irrational exuberance" by the then Chairman of the Federal Reserve Alan Greenspan in a 1996 speech, as well as Nobel prize winner Robert Shiller in Shiller (2000). The reasons behind price bubbles are still unclear, and much of the research cited below uses agent-based models to better understand their inception. Historically, bubbles are often identified by the subsequent crash, i.e. when a particularly sharp and fast collapse in prices follows a spectacular growth. Note that this spectacular growth has been theorized as a symptom of a bubble. More specifically, the formation of super exponential growth of the asset price has been deemed characteristic of empirical market bubbles (as studied in Sornette, Woodard, and Zhou (2009), Jiang et al. (2010) and Yan, Rebib, Woodard, and Sornette (2012)).
We will now clarify the notion of monetary policies. Traditionally, real regulatory financial institutions have used varied degrees of influence over the market in order to alleviate damage from possible market crashes. We can first divide the regulatory financial institutions into two categories: monetary authorities like central banks or currency boards, whose tools are a set of monetary policies they can adopt, or governments who employ fiscal policy. The effects of fiscal policy, that is to say the manipulations of the level of taxes leading to modifications of the interest rate, are outside the scope of this study. Thus, we will now focus on central banks, and the typical goals and instruments they have when they apply monetary policies.
Central banks are several, and are specific to geographical regions. The

goals they aim for, and the instruments they employ, vary among them. The Compendium of monetary policy frameworks and market operations Committee (2019) (BIS Markets Committee, 2019), lists 15 central banks and their policies, for regions including Eurosystem, the United States, Switzerland and China. The Compendium highlights that the philosophy and frequency of monetary policy decisions can vary slightly among different banks. However, the line of questioning and scope of the problem of the present study are general enough to be indifferent to geography. For the purposes of this introduction and some examples through the report, we will often display the empirical evidence related to European Central Bank monetary policies, or the Federal Reserve monetary policies in the United States.
As related in Committee (2019) (BIS Markets Committee, 2019), central banks aim at price stabilization, which is actualized in terms of inflation. Some central banks also include economic growth targets, or employment targets, as secondary objectives in their policies, but this varies among them. The policy target of the Governing Council concerning Eurosystem, of instance, is exclusively based on a price stability policy. More specifically, it is defined as a maintained 2% (or slightly lower) yearly increase of the Harmonised Index of Consumer Prices for the euro area. In comparison, the Federal Reserve is said to have a "dual mandate" of maximum employment and price stability.
Let us briefly mention that in economic theory, the relationship between inflation and interest rate is such that the expected inflation rate π is determined by the difference between the nominal interest rate i and the real interest rate r: $\pi = i - r$. This is the Fisher equation, first introduced in Fisher (1930). It helps conceptualize inflation. As we will see, central banks have means to affect the nominal interest rate, and the amount of money in circulation, and as such, have effects on inflation.
Having established the goals and targets of the central banks, we will now specify the instruments they have at their disposal to maintain their objectives for the market. Their understanding is useful for us later, in order to gauge the realism of the regulatory agents in the model we will use.

The instruments employed by central banks can be separated into standard and non standard tools – also called conventional and unconventional tools. While conventional monetary policies are routinely used, unconventional monetary policies are put in place under extraordinary circumstances, such as a major economic crisis or a risk of deflation, when the usual tools are insufficient. Note that in Europe, the emergence of unconventional monetary policies is relatively recent. It can be traced back to the financial crisis in 2009, according to the European Central Bank European Central Bank (2020).
Let us now specify what conventional monetary instruments the central

banks use. These typically encompass three types of policies: open market operations, reserve requirements, and the discount window, or standing facilities (as referred to in Europe). Firstly, open market operations consist in the purchase or sale of financial instruments. Traditionally, this involved mostly government issued bonds, but this practice has now extended to company-issued bonds. For instance, buying bonds increases the amount of circulating cash, and leads to banks reducing the cost of loans - which decreases the nominal interest rates. Intuitively, such measures encourage spending and fuels economic growth, while keeping inflation low. The sale of bonds has the opposite effect, making open market operations a way to maintain inflation and growth targets. Secondly, reserve requirements refer to the fact that central banks have jurisdiction over what amount of reserves of cash the commercial banks should hold. For instance, before the Covid-19 related crisis in March of 2020, the Federal Reserve set reserve requirements at 10% of the liabilities, according to the Board of Governors of the Federal Reserve System (2020b). Lowering the reserve requirements enables banks to take on comparatively more liabilities, which also encourages them to reduce the cost of loans, and decreases the nominal interest rates. Thus, manipulation of the reserve requirements also helps the central banks keep their targets. Thirdly, central banks also use the the discount window, also called standing facilities, to influence the nominal rate. This instrument refers to the fact that commercial banks have the possibility to borrow or lend to and from the central banks, with a rate set by the central banks themselves, called the discount rate or policy rate. This serves as an alternative to inter bank overnight lending, that occurs because the banks who have too little reserves at the end of the day borrow from others who may deal with excess reserves. The discount rate also influences the nominal rate, since it influences the cost of overnight lending among commercial banks. This in turn affects the interest rate at which commercial banks loan to businesses and individuals.

Mario Draghi, in 2012, famously stated that "the ECB is ready to do whatever it takes to save the euro". Unconventional instruments are extraordinary measures put in place when the state of the market is precarious, and cannot be controlled through regular monitoring anymore. Such states of emergency lead to crossing quantitative ceilings as well as the pioneering the use of new instruments. In our study, the fact that monetary institutions resort to extreme measures when bubbles crash justifies implementing monetary policies at very high levels of investment compared to the total volume of the market and studying their effectiveness.

For instance, in the recent Covid-19 related crisis, as of mid March 2020, the reserve requirement rates set by the Fed have gone from 10% of the liabilities, to 0%, according to the Board of Governors of the Federal Reserve System (2020b). However, in this example, the instrument is still conventional in nature. New instruments were put in place as well, such as Long Term

Refinancing Operations, which are long term loans allotted to commercial banks by the central banks. Asset Purchase Programs are also significant innovations by which central banks no longer only buy loans, but also financial assets, from either public or even private sector in times of crises. Post 2008 use of unconventional instruments and their implementation in Europe are further analyzed in Eser, Carmona Amaro, Iacobelli, and Rubens (2012). This is a foundation for our research, since we will study the effect of regulatory entities who trade on the market in high amounts, or have direct influence on the nominal interest rate. Finally, it is important to note that unconventional monetary policy has been accompanied with unprecedented increases in amounts injected in the financial system by central banks. According to the Board of Governors of the Federal Reserve System (2020a), between mid March and mid June of 2020, the amount held in securities by the Fed has gone from $3.9 trillion to $6.1 trillion. That last figure represents nearly 30% of the US Gross Domestic Product. Figure 0.1 (presented in a speech by Cœuré (2019)) shows that the crisis of 2008 and the subsequent recession has lead to a pattern of unprecedented investments by the European Central Bank, especially in public sector securities (in blue). Let

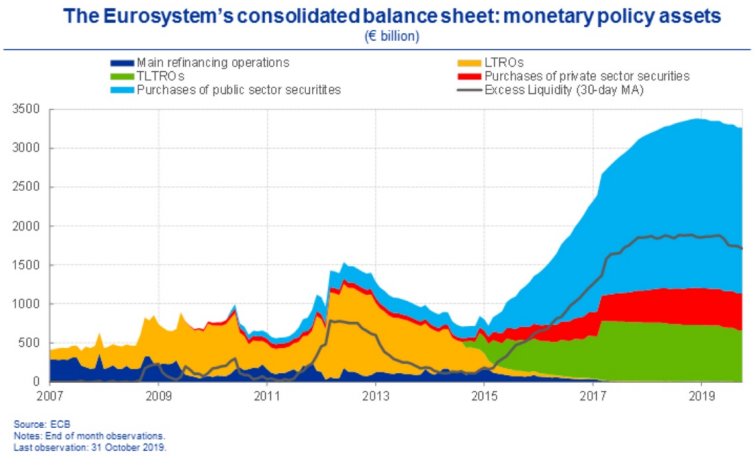

Figure 0.1: Evolution of the Eurosystem's balance sheet from 2007 to 2019, as showed in a 2019 speech by Board member of the European Central Bank Benoît Cœuré Cœuré (2019)

us also note that a common characteristic across central banks is that they do not systematically target asset prices, contrary to what we attempt in our study. When there are responses to price bubbles, they are traditionally rather reactive at the stage of the burst than at the beginning or even at the tipping point of the bubble. For instance, the then Federal Reserve Chairman Alan Greenspan stated in 2004, 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' ". This highlights a policy oriented towards only intervening after a bubble has crashed. Additional arguments in this direction were held in 2004 by the then Chairman Greenspan, including the risk of bubble bursting monetary policy being too strong and stifling genuine economic growth, as related in Roubini (2006). However, Roubini argues in the same exposé that in 2005, central banks in the United Kingdom, Australia and New Zealand had successfully burst their housing bubbles early on, encountering no such side effects. Furthermore, recent bubbles in China have been met with pre-emptive measures involving press communications to slow down a speculative bubble in formation, and not already burst, as detailed in Gatley (2020). A response to Greenspan's position has been developed in favor of central banks adopting asset price targeting, especially in anticipation of market crashes rather than in reaction to them, in paper Roubini (2006). In summary, in the analytical literature, asset price targeting has been put forth as optimal. Notably, theoretical work by B. Bernanke and Gertler (2000); B. S. Bernanke and Gertler (2001) initially argued that basing the reaction function of a monetary authority exclusively on asset prices (to the exclusion of the output gap and inflation measures), in case of exogenous deterministic bubbles, was optimal. This work was refuted to include output and inflation targeting in the reaction function for an optimal result, by A. J. Filardo (2001); A. J. Filardo et al. (2000). Still, the general conclusion is that asset price targeting is complementary to the inflation and growth targeting, and indispensable to help monetary authorities deal optimally with bubbles. Another important result by A. Filardo (2003); A. J. Filardo et al. (2000) is that in the case of endogenous bubbles, that is if bubbles can be affected by monetary policy, then it is optimal for monetary policies to try and suppress the bubble even as it is forming. We rely on this reasoning to support our choice to study asset price targeting-based monetary policies. This leads us to adopt asset price targeting as a main axis of study in the following report.

Our thesis studies asset pricing based monetary policies in the context of a general agent-based model, in the hopes of transferring our findings to reality. We will now highlight the use of agent-based models in the study of financial markets.
As explained in Axelrod (2006), agent-based models improve and encourage interdisciplinary work by addressing problems that are fundamental to many disciplines. He also argues that agent-based models are useful in situations where the math is intractable due to the overwhelming interaction of many agents. This has been widely shown in the context of agent-based

models in finance. In particular, agent-based models of markets have been shown to overcome the limitations of experimental market settings based research in Chan, LeBaron, Lo, and Poggio (1999). In this paper, agent-based models are shown to replicate the results of experimental research while alleviating the doubt cast on rationality of human subjects, as well as simplifying the measure of the effect of the risk aversion and learning abilities of the individual subjects on the experiment results. In our thesis, we will observe two-agent and three-agent models. As developed in Hommes (2006), the use of multi-agent models, mostly introduced in the 1980s, allows for more irrationality in the market and more realistic behavior. It is also explained in Samanidou, Zschischang, Stauffer, and Lux (2007) that multi agent-based models have been used with increasing levels of success to replicate otherwise unexplained emergence of bubbles and subsequent crashes, under rational one-agent models, starting with the market crash of the 1987 US stock market. Samanidou et al. (2007) provides a detailed review the advances of these agent-based models which will also be showcased in the next paragraph.

Now that we have established the background of the scope of our research, let us review the existing works in our sector, that is the study of asset price bubbles using agent-based models, and especially the study of the effect of regulatory policies in this context.
Agent-based models have the potential to explain macroscopic characteristics of financial markets through the establishment of microscopic rules for the agents, which makes them valuable in the understanding of the emergence of asset price bubbles and crashes. Notably, an early attempt to model market crashes was the model of Kim and Markowitz (1989), uses two types of agents, one that favors a constant proportion of cash in their portfolio, and the other that favors a minimal level of wealth in the portfolio at all times. This model does not tackle the stylized facts of financial markets. However, it already shows the destabilizing effect of some types of investors, in this case, the ones who favor guaranteeing a minimal level of wealth. Another approach by Lux and Marchesi (1999, 2000) allows the agents to switch between a role of rational trader or noise trader, influenced by its peers. This model reflects characteristics of real stock price patterns, especially volatility clustering, power-law distributed returns and self similarity in prices.
Our research is based on a variation of the agent-based model of Kaizoji et al. (2015). In essence, it is an heterogeneous agent model with two assets (risky and risk-free), and two types of agents: a fundamentalist traders and a chartist traders (or noise traders). While the fundamentalist traders optimize their portfolio according to a rational process of expected utility maximization, the noise traders choose their wealth allocation by following the market trend (momentum trading) and the behavior of fellow noise traders

(social imitation). This model is interesting for us because it replicates stylized facts of the market such as volatility clustering and especially accounts for the formation of the aforementioned super exponential growth of the risky asset price that has been deemed characteristic of empirical market bubbles. This gives us a good foundation for bubble control through regulation. Note that in Xu, Zhang, Xiong, and Zhou (2014), the authors also developed such a model, mainly differing in the dynamics attributed to the chartist traders, and improving slightly the accuracy of the fundamentalist's estimations in their investment choices, successfully reproducing fat-tailed return distribution of returns and long-term memory of volatility. Examples of a similar model with a higher dimension in the risky asset price vector is also developed in Chiarella, Dieci, and He (2009). However, we keep the model of Kaizoji et al. (2015) because of its well established ability to generate asset price bubbles.

Using agent-based models in attempts to design or test regulatory policies is natural given how tremendous the scope and impact of real-life experiments would be. A review by Mizuta (2016) highlights varied attempts to use artificial markets, or agent-based models, to discuss the effect of some financial market regulations., including limits on short selling, limitations on frequency of trading, or the Basel regulatory framework for instance. Testing of central bank intervention within agent based models, however, in terms of interest rate intervention or as an additional trading agent modeling open market operations, has been sparse. Note that in Westerhoff (2008), the author specifically tests the effect of "lean against the wind" central bank intervention through foreign currency operations, on a different agent-based market model than that of Kaizoji et al. (2015), that reproduces the stylized facts of financial markets such as fat tails, volatility clustering, and the emergence of bubbles and crashes. He finds some efficiency in reducing bubbles and stabilizing the market in such a strategy. However, the foreign currency intervention is limited to a "lean against the wind" strategy that is non specific to bubble detection, contrary to what we will attempt in our thesis. Lastly, Westphal and Sornette (2020b) is a foundation for our research. It is not testing the effect of regulatory policies on the market, but rather the effect of introducing in Kaizoji et al. (2015) a category of traders, called Dragon Riders, equipped with methods of bubble detection, on the market. Dragon Riders are agents who use their ability to predict extreme events to profit from this information. Another type of agent is Dragon Slayers: they are agents who use the same ability to detect extreme events in order to prevent them. The denomination stems from the work of Sornette (2009); Sornette and Ouillon (2012), who proposed the name "Dragon-King" to describe statistical outliers of strong significance, such as asset price bubbles and their subsequent crashes. Kaizoji et al. (2015) find that in certain proportions, such traders reduce the size and frequency of bubbles in the model. An additional foundation of our work is an internal report of Westphal and

Sornette (2020a), which studies a similar introduction but with another bubble detection tool, which showed even better control of bubbles. These two results informed the research in that they hint at possible effectiveness of both open market operations in this context, and the use of a bubble detection tool (introduced in Westphal and Sornette (2020a)). We will use the latter in several experiments, and use the model of Westphal and Sornette (2020a) as a benchmark to assess the effect of regulatory policies.
In this context, we chose to study possible strategies of central banks on the two-agent model of Kaizoji et al. (2015) by introducing an agent using open market operations to reduce the size and amount of bubbles, and later an agent performing interest rate modification with the same motivation. Both are compared to the impact of the bubble arbitraging traders of Westphal and Sornette (2020a). Our objective is to find effective and realistic policy using open market operations and interest rate regulation in the context of a simple agent-based model. We also aim to assess the effect these actors can have on the market and how they compare to the effect of an informed trading population from Westphal and Sornette (2020a).
While the gist of the specific hypotheses needed in the model are outlined in chapter 1, we should note that more generally, the findings we have are related to a discrete market, with two assets (a risky and risk-free one), and two or three types of traders. The frequency of trading is daily, with no transaction costs, no borrowing or short selling, and the price calculations rely on the equilibrium of supply and demand.
In chapter 1, we establish the setup and properties of the two-agent model which serves as foundation for the rest of the study. We then move on, in chapter 2, to exploring the effect of the bubble arbitraging traders of Westphal and Sornette (2020a), or Dragon Riders, on this model, performing additional grid explorations to optimize them in a self-interested fashion. In chapter 3, we introduce to the model a central bank incentivized type of traders, or open market operations performers, called Dragon Slayers. They use bubble detection tools in order to reduce the size and number of bubbles, and study their impact. Finally, we test direct interest rate modification, in another effort to reduce the size and amount of bubbles, in chapter 4.

Chapter 1

Initial market model: two types of agents

In this chapter, we will describe the foundation for our work: a slight variation of the market model with two types of agents described by Kaizoji et al. (2015), with the same asset dynamics as Westphal and Sornette (2020b). The motivation is the following: in chapter 2 and chapter 3 of the report, we will present the workings of a type of intervention to reduce price bubbles in an agent-based model. These two models are different extensions with three types of agents, of the model presented in this chapter. This makes the original market model, with two types of agents, the foundation of our work.

In the extended models, as we will detail below, we keep the framework of Kaizoji et al. (2015), with two assets, and two types of standard traders, while a third type of agent is added: the "Dragon Slayer". As mentioned in the introduction, this denomination is intended to describe agents who are chasing "Dragon King" events, in this case, asset price bubbles, in order to prevent them. On the one hand, we have two kinds of traders who represent the "consumer" types of market agents, who are motivated by wealth gain, and are exhibiting either rational or irrational behavior. This setup alone is capable of reproducing several stylized facts of financial markets, including fat-tail distribution of returns and volatility clustering. In this model, the risky asset price also exhibits transient faster-than-exponential bubble growth, as shown in Kaizoji et al. (2015).

On the other hand, in this context, it makes sense to include some regulatory entity equipped with some method of bubble detection, as well as some method of intervention, and measure the effectiveness of its actions. Characteristic dynamics of the two-asset market model will serve as a benchmark to measure the effect of various interventions.

We will now explain the workings of the market model with two types of agents, describing the assets, the behavior of the fundamentalist traders and

noise traders, and the price setting via market clearing.

1.1 Market investing setup - assets

This agent-based model allows for discrete time investment. At each point in time, the traders decide which proportion of their wealth they attribute to which asset. In this model, the traders can choose between two assets: a risk-free asset and a risky asset.
The parameters are chosen such that one time step corresponds roughly to a trading day, therefore each of the returns exemplified below are to be understood as daily parameters. The convention is that the number of trading days is 250, therefore we will sometimes have to annualize our data in the form of (daily return)$\times$250 in order to gain more perspective on the understanding of our results.
The risk-free asset represents a risk-free government bond or bank account, such that the return on this asset is a constant interest rate r_f at each step t. In this setup, the interest rate r_f is constant over time, but we will explore the effects of deliberate modification of r_f on market dynamics in chapter 4.
The risky asset represents an index or an investment fund. Its dynamics are such that the asset price p_t follows the market equilibrium of supply and demand. The asset also pays a dividend d_t at each time step. The dividend itself follows a stochastic growth process as follows:

$$d_t = d_{t-1}(1 + r_t^d) = d_0 \Pi_{k=1}^{t}(1 + r_k^d), \tag{1.1}$$

with a normally distributed growth rate r_t^d

$$r_t^d = r_d + \sigma_d u_t, \tag{1.2}$$

where u_t are i.i.d. standard Gaussian processes $\mathcal{N}(0,1)$, $r_d > 0$ is the mean growth rate, and σ_d is the standard deviation. In practice, σ_d is chosen small enough that $r_t^d > 0$ in the test cases. This dividend dynamic differs from Kaizoji et al. (2015) and replicates that of Westphal and Sornette (2020b).
This means that the value of the asset at time $t+1$ is

$$p_{t+1} = p_t \left(1 + r_{t+1} + \frac{d_{t+1}}{p_t}\right), \tag{1.3}$$

where the capital return of the risky asset is defined as $r_{t+1} = \frac{p_{t+1} - p_t}{p_t}$.
Finally, we do not allow short selling or borrowing for the traders involved. This corresponds to constraining the fraction of wealth invested in the risky asset to $[0,1]$. In this simple context, we can now observe the resulting behavior of two types of investors, as follows.

1.2 Fundamentalist traders

The first type of trader is the fundamentalist trader, designed to represent an educated investor with the intention of growing wealth. In the following, as in Kaizoji et al. (2015), we determine the risky fraction that the fundamentalist chooses at a given time $t-1$.
Note first that the investor is risk-adverse, as most real-life individual investors are. In order to understand that notion, consider the lottery $L = (1, 0.5; 0, 0.5)$, that is, the possibility of receiving 1 of arbitrary currency with probability 0.5, and 0 otherwise. Then there exists an amount lower than $\mathbb{E}[L] = 0.5$ that the risk-free agent would prefer receiving with certainty, instead of the lottery.
The preferences of this risk-adverse agent are represented using expected utility theory. This is due the Von Neumann–Morgenstern utility theorem, established in Morgenstern and Von Neumann (1953). Indeed, the foundation of expected utility theory is the fact that an agent with "rational" preferences (defined as a set of axioms in the aforementioned work) can be equipped with a function $u(x)$ of the outcome of owning a total wealth of x (in arbitrary monetary units), called utility function. Then, it is possible to use the expected value of this utility function over the event space as a representation of the preferred lottery of a decision maker.
In concrete terms, in the following example in a discrete event space $\Omega = \{\omega_1, ..., \omega_n\}$, for any set of lotteries $L_i = (l_i^1, p_1, ..., l_i^n, p_n), i \in \{1, ..., m\}$, the lotteries can be ranked as an order of preferences in decreasing order of the value of the expected utility

$$E[u(L_i)] = \sum_{k=1}^{n} p_k u(l_i). \tag{1.4}$$

While empirical evidence shows limits of this theory in some paradox situations, notably cited in the work of Allais (1953), it is very useful in our simplified setup, since it allows for a clear representation of the fundamentalist's behavior.
We also make the assumption that the fundamentalist traders all have identical preferences and information, so that we can aggregate their behavior into only one representative agent with one utility function, investing a fraction of the total wealth of the fundamentalists. In the following, we only consider this agent.
The chosen utility function for the fundamentalist trader is CRRA (constant relative risk aversion). This means, intuitively, that the risk aversion of the agent is inversely proportional to its wealth level. This is mathematically represented in the following way, as defined in Arrow (1965):

$$-x\frac{u''(x)}{u(x)} = \gamma, \tag{1.5}$$

where $\gamma > 0$ is a constant (it is positive because the agent is strictly risk adverse). An example of a function fulfilling property (1.5) is:

$$u(x) = \begin{cases} log(x) \text{ if } \gamma = 1 \\ \frac{x^{1-\gamma}}{1-\gamma} \text{ if } \gamma \neq 1 \end{cases} \tag{1.6}$$

This agent is now given the ability to choose, at each time step, the fraction x_t^f of its wealth invested in the risky asset. The remainder will be invested in the risk-free asset. The choice reflects an optimization of the expected utility for the next time step. This follows Kaizoji et al. (2015) and Westphal and Sornette (2020a). Now, let us compute the risky fraction that the fundamentalist chooses at a given time $t-1$.
Firstly, let us look at the expected utility maximization problem. The fundamentalist solves for x_{t-1}^f in the following:

$$max_{x_{t-1}^f} E_{t-1}[u(W_t^f)], \tag{1.7}$$

where W_t^f is the fundamentalist's wealth at time t. That is, this trader maximizes the expectation of his utility at time t.
Here, we assume, as in Kaizoji et al. (2015), that the wealth variations $W_t^f - W_{t-1}^f$ at each time step are small, enabling us to use a Taylor expansion of order 2 to solve the problem. This implies, as we will see below, that r_f, r_t and $\frac{d_t}{p_{t-1}}$ have to be small compared to 1. Problem (1.14) becomes

$$max_{x_{t-1}^f} \left[u(W_{t-1}^f) + u'(W_{t-1}^f) E_{t-1}[W_t^f - W_{t-1}^f] + \frac{1}{2} u''(W_{t-1}^f) E_{t-1}[(W_t^f - W_{t-1}^f)^2] \right] \tag{1.8}$$

Secondly, in order to obtain an explicit formula, we can write the wealth of the fundamentalist as:

$$W_t^f = p_t n_t^f + P_t N_t^f, \tag{1.9}$$

where n_t^f and N_t^f are the number of risky and risk-free assets held, respectively, and p_t and P_t are the prices of the risky an risk-free assets respectively. In this context, we know that the risky fraction of the fundamentalist is $x_t^f = \frac{p_t n_t^f}{W_t^f}$, and we also have $1 - x_t^f = \frac{P_t N_t^f}{W_f^f}$. Therefore we can write the following wealth variation of the fundamentalist from $t-1$ to t. Note that as the traders choose their wealth allocation n_{t-1}^f and N_{t-1}^f at $t-1$, before knowing the state of the market at t.

$$\begin{aligned} W_t^f - W_{t-1}^f &= (p_t - p_{t-1}) n_{t-1}^f + d_t n_{t-1}^f + (P_t - P_{t-1}) N_{t-1}^f \\ &= W_{t-1}^f \left[x_{t-1}^f \left(\frac{p_t - p_{t-1}}{p_{t-1}} + \frac{d_t}{p_{t-1}} \right) + (1 - x_{t-1}^f) \left(\frac{P_t - P_{t-1}}{P_{t-1}} \right) \right] \\ &= W_{t-1}^f \left[x_{t-1}^f \left(r_t - r_f + \frac{d_t}{p_{t-1}} \right) + r_f \right] \end{aligned}$$

Therefore

$$W_t^f - W_{t-1}^f = W_{t-1}^f \left[x_{t-1}^f r_{excess,t} + r_f \right], \tag{1.10}$$

where the excess return $r_{excess,t}$ is defined as the difference between the risky asset return and the risk-free asset return (i.e. the risk-free rate r_f):

$$r_{excess,t} = r_t + \frac{d_t}{p_{t-1}} - r_f = r_t + \frac{d_{t-1}(1 + r_t^d)}{p_{t-1}} - r_f \tag{1.11}$$

and as previously, $r_t := \frac{p_t}{p_{t-1}} - 1$ is the capital return.
We can now rewrite the simplified expected utility maximization problem (1.8) as follows:

$$max_{x_{t-1}^f} \left[u(W_{t-1}^f) + u'(W_{t-1}^f) W_{t-1}^f x_{t-1}^f E_{t-1}[r_{excess,t}] + \frac{1}{2} u''(W_{t-1}^f)(W_{t-1}^f)^2 (x_{t-1}^f)^2 Var_{t-1}[r_{excess,t}] \right] \tag{1.12}$$

Solving gives us

$$\begin{aligned} x_{t-1}^f &= -\frac{u(W_{t-1}^f)}{W_{t-1}^f u''(W_{t-1}^f)} \frac{E_{t-1}[r_{excess,t}]}{Var_{t-1}[r_{excess,t}]} \\ &= \frac{1}{\gamma} \frac{E_{t-1}[r_{excess,t}]}{Var_{t-1}[r_{excess,t}]} \end{aligned} \tag{1.13}$$

From (1.11), if we denote $E_{r_t} := E_{t-1}[r_t]$, we can write

$$E_{t-1}[r_{excess,t}] = E_{r_t} + \frac{d_{t-1}(1 + r_d)}{p_{t-1}} - r_f \tag{1.14}$$

Now, using independence of the dividend return r_d and the capital return r_t, and denoting $\sigma^2 := Var_{t-1}[r_t]$, we can write the variance as follows

$$Var_{t-1}[r_{excess,t}] = \sigma^2 + \frac{d_{t-1}^2 \sigma_d^2}{p_{t-1}^2} \tag{1.15}$$

Where r_t and d_t are assumed to be uncorrelated, because the dividend policy is assumed to be set without consideration for the market price, as in Kaizoji et al. (2015). This is verified in Modigliani and Miller (1963, 1965) under symmetric information and bounded rationality. We also assume that the dividend d_t is very small compared to p_t and use first order approximation, therefore

$$x_{t-1}^f \approx \frac{E_{r_t} + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_f}{\gamma \sigma^2} \tag{1.16}$$

This is the equation that the model uses to simulate the behavior of fundamentalist traders. E_{r_t} and σ are taken to be exogenous variables, as we assume that fundamentalists are myopic, and do not learn in time but rather

choose their wealth allocation based on their fundamental valuation, derived from long-term historically computed growth rate and variance. As verified in Westphal and Sornette (2020b), the long-term growth rate of prices is approximately equal to r_d, the average growth rate of the dividends. Note that the fundamentalists only use constant σ as a best guess of what the variance will look like. Note that the long-term expected return of the risky asset E_{r_t} is higher than the interest rate r_f. The justification for this assumption is that the risk premium is supposed to be positive, otherwise the risk-adverse traders would never buy the risky asset. This is also coherent with real life behavior: there is always an equity premium, to the point where it cannot always be explained with risk aversion only, as exposed in Mehra and Prescott (1985). Therefore, the risky fraction stays positive. In the model, we need to ensure additionally that the risky fraction stays within $[0, 1]$, to prevent borrowing and short selling, as mentioned above 1.1.

1.3 Noise traders

In this section, we will describe how the behavior of the second category of traders is simulated. This type of trader is called a noise trader. They are also often called chartist traders, due to the fact that they analyze patterns in the price path. As explained in our introduction, and according to the review of Samanidou et al. (2007), such momentum trading agents were already studied in the 1950s, for instance in Baumol (1957), for their destabilizing effect. However, their use did not become popular at the time, because of arguments against their existence in real markets. Such agents were sometimes argued to naturally disappear in markets, as can be found in Friedman and Friedman (1953). However, momentum trading yields profitable strategies and as such, this argument was overcome. Thus, these traders gained traction again in the 1970s and 1980s, for their potential to help reproduce the contemporary price bubbles such as the dollar bubble, according to Samanidou et al. (2007).
Those traders decide how they invest based on the opinion of other noise traders, and the price behavior. Contrary to the fundamentalist traders, noise traders do not have the possibility to invest only part of their wealth in either asset. They invest fully in either the risky asset or the risk-free asset. Their decision is probabilistic, and each trader's behavior is independent. However, the probability of switching from one asset to another is dependent on the opinion of the other traders, and the behavior of the price, represented by the opinion index s_t and the price momentum H_t respectively.
The opinion index depends on the number of noise traders invested in either asset. We denote N_t^+ as the number of noise traders invested in the risky

asset, and N_t^- the number of noise traders invested in the risk-free asset. Then, we write the opinion index as

$$s_t = \frac{N_t^+ - N_t^-}{N_t^+ + N_t^-} \tag{1.17}$$

s_t is always in $[-1, 1]$. Note that $N_t^+ + N_t^- = N$ is a constant: the total number of noise traders doesn't change over time in our simulations. A positive value of the index indicates that noise traders are more invested in the risky asset, i.e. bullish opinion, whereas a negative value of s_t indicates that the risk-free asset seems more attractive to the noise traders, i.e. bearish opinion.
As for the price momentum, it is an exponential moving average of the price returns $r_t = \frac{p_t}{p_{t-1}} - 1$, defined as follows

$$H_t = \theta H_{t-1} + (1 - \theta) r_t \tag{1.18}$$

with $\theta \in [0, 1]$ a memory parameter. The noise traders keep a frame of memory of a duration around the order of $\frac{1}{1-\theta}$. The reason behind this order of magnitude is that the center of mass of the weights of an exponential moving average of memory parameter θ falls on the same day as a simple moving average of the previous n days (of the form $\frac{1}{n}\sum_{i=0}^{n-1} H_{t-i}$), when $n = \frac{1}{1-\theta}$. This means that when θ goes closer to 1, the memory goes to infinity.
The probability of switching from one asset to another changes at each time step. It is denoted written p_t^+ for the change from risky asset to risk-free asset, and p_t^- for the reverse. As stated previously, they depend on the opinion index and the momentum, but their influence is controlled by so-called herding propensity parameter, κ_t. This parameter is a time-dependent social imitation and momentum following strength parameter, so that we obtain the following formula

$$p_t^\pm = \frac{p_\pm}{2}(1 \mp \kappa_t(s_t + H_t)). \tag{1.19}$$

The expression of the transition probabilities is what causes this model to be analogous to a standard kinetic Ising model. As explained in Harras, Tessone, and Sornette (2012), the herding propensity H_t is analogous to a coupling strength in the Ising model, because it represents the influence of agents among themselves. The momentum parameter s_t is analogous to an external magnetic field.
The herding strength parameter is of crucial importance in defining the behavior of noise traders. Note first that s_t and H_t have the same multiplier κ_t in (1.19). As discussed in Kaizoji et al. (2015), one could envision two separate parameters for social imitation and momentum following respectively.

However, they are considered equivalent here for the sake of simplification.
Now, if $\kappa_t < 0$, there would be a tendency from the noise traders to shift from the risky asset to the risk-free asset when the risky asset growth is high (H_t high) and the opinion index is high, i.e. when many noise traders are invested in the risky asset. This contrarian behavior corresponds to agents going opposite direction from perceived trends, and is not representative of the herding effect that we want to include in the model.
$\kappa_t = 0$ corresponds to a state where noise trader behavior is not influenced by the opinion index or the price momentum. In that case, the probabilities of switching reduce to $\frac{p_\pm}{2}$. This means that on average, a noise trader will remain in the risky asset position $\sim \frac{2}{p_+}$ time steps, and in the risk-free asset position $\sim \frac{2}{p_-}$ time steps.
In our study, we keep $\kappa_t > 0$, to account for social imitation: the higher the proportion of noise traders invested in the risky asset, and the higher the price momentum, the more likely noise traders are to converge to and remain in the risky asset position, and vice-versa.
In this realm, we must also note that similarly to the Ising model, there is a critical value κ_c such that if $\kappa_t > \kappa_c$, the model enters a regime of spontaneous organization of the individual noise traders, causing phases of excess volatility and price bubbles. However, if $\kappa_t < \kappa_c$, the model is in a disordered state.
Now, the idea behind having a non constant parameter κ_t is to account for the realistic financial market characteristic of regime switching. Financial markets typically exhibit signs of distinct regimes and switches between them, in such ways that it has been shown that strategies of momentum following can be very successful for some stretches of time, but in other periods, i.e. under other regimes, they fall through. This phenomenon is further developed in Ang and Timmermann (2012). For the purposes of our work, we will examine the behavior of markets with a non constant parameter κ_t, defined as an Ornstein-Uhlenbeck process

$$\kappa_t = \kappa_{t-1} + \eta_\kappa(\mu_\kappa - \kappa_{t-1}) + \sigma_\kappa v_t, \tag{1.20}$$

where η_κ is the mean reversion strength, μ_κ is the mean value, the last term $\sigma_\kappa v_t$ is one of an i.i.d. set of Gaussian processes with standard deviation σ_κ, i.e. such that $v_t \overset{iid}{\sim} \mathcal{N}(0,1)$. Ornstein-Uhlenbeck processes are characteristically mean-reverting, here to the mean value μ_κ, with mean reversion strength η_κ, such that the deviations from the mean have a characteristic persistence time $\sim \frac{1}{\eta_\kappa}$. Parallels will occasionally be drawn with markets with constant κ, which do not account for regime switching and do not feature momentary super exponential growth, in order to gain better understanding of the effect of our strategies on bubbles, specifically, as opposed to a one-regime (constant κ) market.
This process allows for regime switching in temperament of the noise traders.

Said deviations are caused by the fluctuations of the last term.
Ornstein-Uhlenbeck processes have the following expected value

$$E[\kappa_t] = \mu_\kappa + (\kappa_0 - \mu_\kappa)e^{-\eta_\kappa t} \tag{1.21}$$

and following variance

$$Var[\kappa_t] = \frac{\sigma_\kappa^2}{2\eta_\kappa}(1 - e^{-\eta_\kappa t}) \tag{1.22}$$

This implies that the long-term dynamics of an Ornstein-Uhlenbeck process is that of a normal distribution

$$\kappa_t \sim_{t\to\infty} \mathcal{N}(\mu_\kappa, \frac{\sigma_\kappa^2}{2\eta_\kappa}) \tag{1.23}$$

As explained in Westphal and Sornette (2020b), the typical reversion time from a value $\kappa_0 > \kappa_c$ to a value $\kappa_c > \mu_\kappa$ is

$$\frac{1}{\eta} log\left(\frac{\kappa_0 - \mu_\kappa}{\kappa_c - \mu_\kappa}\right) \tag{1.24}$$

Since the noise traders all have identical behavior, we aggregate them into one representative trader with the following risky fraction, in $[0, 1]$.

$$\begin{aligned} x_t^n &= \frac{N_t^+}{N_t^+ + N_t^-} \\ &= \frac{1}{N_t^+ + N_t^-} \sum_{i=1}^{N_{t-1}^+} (1 - \xi_i(p_{t-1}^+)) + \frac{1}{N_t^+ + N_t^-} \sum_{j=1}^{N_{t-1}^-} (1 - \xi_j(p_{t-1}^-)) \end{aligned} \tag{1.25}$$

Note that the random variables $\xi(p)$ follow independent Bernoulli distributions, that is

$$\xi(p) = \begin{cases} 1 \text{ with probability } p \\ 0 \text{ with probability } 1 - p \end{cases} \tag{1.26}$$

Finally, we can write the aggregated wealth W_t^n for the noise traders. It has the same structure as the wealth of the fundamentalist trader:

$$W_t^n = W_{t-1}^n \left[1 + r_f + x_{t-1}^n \left(r_t + \frac{d_t}{p_{t-1}} - r_f\right)\right] \tag{1.27}$$

1.4 Market clearing

In order to derive the price dynamics for the model, we use equilibrium of supply and demand, as in Westphal and Sornette (2020b).

1.4.1 Market clearing conditions

We write the excess demand of the fundamentalist traders (or the representative fundamentalist trader) for the risky asset as

$$\Delta D_t^f = n_t p_t - n_{t-1} p_t. \tag{1.28}$$

It corresponds the variation of the amount invested in the risky asset by the representative fundamentalist trader between times $t-1$ and t. Further developing, we obtain

$$\begin{aligned}\Delta D_t^f &= x_t^f W_t^f - x_{t-1}^f W_{t-1}^f \frac{p_t}{p_{t-1}} \\ &= W_{t-1}^f \left(x_t^f \left[1 + r_f + x_{t-1}^f (r_t + \frac{d_t}{p_{t-1}} - r_f) \right] - x_{t-1}^f \frac{p_t}{p_{t-1}} \right).\end{aligned} \tag{1.29}$$

Similarly, we can write the excess demand of the representative noise trader as

$$\Delta D_t^n = W_{t-1}^n \left(x_t^n \left[1 + r_f + x_{t-1}^n (r_t + \frac{d_t}{p_{t-1}} - r_f) \right] - x_{t-1}^n \frac{p_t}{p_{t-1}} \right). \tag{1.30}$$

The structure is identical, due to the identical structures of wealth. The only difference between the agents is in how the risky fractions are computed (i.e. in their behavior). Now, the market clearing condition of equilibrium of supply and demand corresponds to a null sum of excess demands. This is in accordance to the theory of general equilibrium of Walras (1954). Here, it is written as the following condition:

$$\Delta D_t^f + \Delta D_t^n = 0. \tag{1.31}$$

This is the foundation of the following price setting.

1.4.2 Derivation of the price

In the framework we described, following the approaches of Kaizoji et al. (2015) as well as Westphal and Sornette (2020b), we compute the price of the risky asset p_t at each time step, using the fact that the sum of excess demands for each time step should be zero, if supply and demand are matched.
It follows that

$$\begin{aligned}0 = \;& \Delta D_t^f + \Delta D_t^n \\ = \;& W_{t-1}^f \left(x_t^f \left[1 + r_f + x_{t-1}^f (r_t + \frac{d_t}{p_{t-1}} - r_f) \right] - x_{t-1}^f \frac{p_t}{p_{t-1}} \right) \\ & + W_{t-1}^n \left(x_t^n \left[1 + r_f + x_{t-1}^n (r_t + \frac{d_t}{p_{t-1}} - r_f) \right] - x_{t-1}^n \frac{p_t}{p_{t-1}} \right),\end{aligned}$$

and using the approximation (1.16):

$$\begin{aligned}0 = \; & W_{t-1}^f \left(\frac{E_{r_{t+1}} + \frac{d_t(1+r_d)}{p_t} - r_f}{\gamma\sigma^2} \left[1 + r_f + x_{t-1}^f \left(r_t + \frac{d_t}{p_{t-1}} - r_f \right) \right] - x_{t-1}^f \frac{p_t}{p_{t-1}} \right) \\ & + W_{t-1}^n \left(x_t^n \left[1 + r_f + x_{t-1}^n (r_t + \frac{d_t}{p_{t-1}} - r_f) \right] - x_{t-1}^n \frac{p_t}{p_{t-1}} \right). \end{aligned} \tag{1.32}$$

Multiplying by p_t and developing the return $r_t = \frac{p_t}{p_{t-1}} - 1$ gives

$$\begin{aligned}0 = \; & W_{t-1}^f \left(\frac{E_{r_{t+1}} p_t + d_t(1+r_d) - r_f p_t}{\gamma\sigma^2} \left[1 + r_f + x_{t-1}^f (\frac{p_t}{p_{t-1}} - 1 + \frac{d_t}{p_{t-1}} - r_f) \right] - x_{t-1}^f \frac{p_t^2}{p_{t-1}} \right) \\ & + W_{t-1}^n \left(x_t^n p_t \left[1 + r_f + x_{t-1}^n \left(\frac{p_t}{p_{t-1}} - 1 + \frac{d_t}{p_{t-1}} - r_f \right) \right] - x_{t-1}^n \frac{p_t^2}{p_{t-1}} \right). \end{aligned} \tag{1.33}$$

We can now regroup the terms in powers of p_t to obtain a quadratic equation:

$$\begin{aligned}0 = \; & \frac{p_t^2}{p_{t-1}} \left(W_{t-1}^f x_{t-1}^f \left(\frac{E_{r_{t+1}} - r_f}{\gamma\sigma^2} - 1 \right) + W_{t-1}^n x_{t-1}^n (x_t^n - 1) \right) \\ & + p_t \left(W_{t-1}^f \left(x_{t-1}^f \frac{d_t(1+r_d)}{\gamma\sigma^2 p_{t-1}} + \frac{E_{r_{t+1}} - r_f}{\gamma\sigma^2} \left(1 + r_f + x_{t-1}^f \left(\frac{d_t}{p_{t-1}} - r_f - 1 \right) \right) \right) \right) \\ & + W_{t-1}^n x_t^n \left(1 + r_f + x_{t-1}^n \left(\frac{d_t}{p_{t-1}} - r_f - 1 \right) \right) \Big) \Big). \end{aligned}$$

We use the following notations from Westphal and Sornette (2020b) to simplify the expressions. A different decomposition will be given in chapter 4 to provide an understanding of the effect of interest rate r_f on the price.

$$v_1 := \frac{E_{r_{t+1}} - r_f}{\gamma\sigma^2}, v_2 := \frac{d_t}{p_{t-1}} - Rf - 1, v_3 = 1 + r_f, v_4 := \frac{d_t(1+r_d)}{\gamma\sigma^2}$$

This gives us the quadratic equation $a_t p_t^2 + b_t p_t + c_t = 0$, if we define a_t, b_t, c_t as follows

$$a_t = \frac{1}{p_{t-1}} \left(W_{t-1}^f x_{t-1}^f (v_1 - 1) + W_{t-1}^n x_{t-1}^n (x_t^n - 1) \right) \tag{1.34}$$

$$b_t = W_{t-1}^f \left(x_{t-1}^f \frac{v_4}{p_{t-1}} + v_1 (x_{t-1}^f v_2 + v_3) \right) + W_{t-1}^n x_t^n (x_{t-1}^n v_2 + v_3) \tag{1.35}$$

$$c_t = W_{t-1}^f v_4 (x_{t-1}^f v_2 + v_3) \tag{1.36}$$

If $a_t \neq 0$, the solutions to this equation are

$$p_t^{1,2} = \frac{-b_t \pm \sqrt{b_t^2 - 4a_t c_t}}{2a_t} \tag{1.37}$$

We seek a positive solution to the equation. In order to assess the sign of either solution (and what happens in the case of $a_t = 0$), we use the fact that both risky fractions x_t^n and x_t^f are maintained in $[0, 1]$, i.e. borrowing and short selling are not allowed, as mentioned in 1.1.
This implies, first, that $(x_t^n - 1) \leq 0$. Now, considering $(v_1 - 1)$, we can observe that $v_1 = x_t^f - \frac{d_t(1+r_d)}{p_t} < 1$ because $x_t^f \leq 1$. Therefore, since the other factors in equation (1.34) are either price or wealth, thus necessarily positive, we always have $a_t \leq 0$.
In order to show that b_t and c_t are strictly positive, we just have to note that for all $x \in [0, 1]$, $xv_2 + v_3 \geq (1 - x)(1 + r_f) > 0$. Since strictly risk adverse investors buy the risky asset, we also have a strictly positive risk premium, therefore $v_1 > 0$, and $v_3, v_4 > 0$. This leads to $b_t, c_t > 0$.
In the case of $a_t = 0$, $p_t = -\frac{c_t}{b_t}$ is not a positive solution. Therefore we only consider $a_t < 0$, and retain the positive solution

$$p_t = \frac{-b_t - \sqrt{b_t^2 - 4a_t c_t}}{2a_t} \tag{1.38}$$

1.5 Behavior of the two-asset market model

We will now summarize the main results of the two-asset market model. In essence, the model with Ornstein-Uhlenbeck herding propensity κ is able to account for several stylized facts of financial markets, including fat-tail distribution of returns and volatility clustering. The risky asset price also exhibits transient faster-than-exponential bubble growth, which is the basis for the following work on trading strategies and regulatory strategies, designed to hinder such growth.

Parameter choice

In order to present the behavior of the model, let us first state our parameter choices. They are designed to replicate real market behavior, where the risky asset represents an index or a fund, and the risk-free asset represents a bank account or a risk-free bond. The parameters are chosen according to the recommendations of Westphal and Sornette (2020b) and listed in table 1.1.

Notably, the wealth of the two initial traders is identical. Since they correspond to aggregated traders, this means that both trader populations have the same weight in the market. Another important point is that the fundamentalists trade under the belief that $E_{r_t} = r_d$, which is the long-term historical view that they form. We assume that they form it with good accuracy, following the line of thought of Lera and Sornette (2017).
Noise traders, as a group, are 1000, and they have a memory parameter θ corresponding to a window of time of $\sim \frac{1}{1-\theta} = 20$ trading days, that is around a month. The probabilities p_- and p_+ are the expected values of the switching probabilities in the case of zero herding propensity. Their values set according to Khort (2016). They are both around 0.2, indicating that the expected duration lasted in either state, in the absence of herding phenomena, is around $\sim \frac{1}{0.2} = 5$ days.
As for the assets themselves, we can observe that the standard deviation σ_d of the growth rate of the dividend is ten times lower than its expected value r_d. Since this is a Gaussian process, this relationship indicates that the growth rate does not take values very far from r_d. This allows the risky fraction of the fundamentalist to remain steady.

Parameter name	Explanation	Value
Market		
T	Length of the simulation	10000
Assets		
r_f	Interest rate (risk-free asset)	$0.01/250 = 4 \times 10^{-5}$
d_0	Initial dividend of the risky asset	$0.04/250 = 1.6 \times 10^{-4}$
r_d	Expected growth rate of the dividend	$0.04/250 = 1.6 \times 10^{-4}$
σ_d	Standard deviation of the dividend growth rate	1.6×10^{-5}
p_0	Initial price of the risky asset	1
σ	Expected standard deviation of the risky asset price p_t	$\sqrt{0.1/250} = 0.02$
N	Number of risky assets	1
Noise trader		
x_0^n	Initial risky fraction of the noise trader	0.5
W_0^n	Initial wealth of the noise traders	10^9
p_+	Basic switching probability from the risky asset to the risk-free asset	0.199375
p_-	Basic switching probability from the risk-free asset to the risky asset	0.200625
θ	Memory parameter	0.95
H_0	Initial momentum	1.6×10^{-4}
N^n	Number of noise traders	1000
Fundamentalist trader		
x_0^f	Initial risky fraction of the fundamentalist	0.3
W_0^f	Initial wealth of the fundamentalist	10^9
E_{r_t}	Expected return of the risky asset	1.6×10^{-4}
Social coupling strength		
κ_c	Critical social coupling strength	0.199375
κ_0	Initial social coupling strength	0.98×0.199375
μ_κ	Mean of the OU social coupling strength	0.98×0.199375
η_κ	Mean reversion of the OU social coupling strength	0.11
σ_κ	Standard deviation of the OU social coupling strength	0.001

Table 1.1: Parameters for the simulations. The interest rate, dividend and standard deviations are defined per time step, i.e. daily values.

Price behavior and bubbles, wealth

We will now summarize the main qualitative characteristics of this model, an extension of which is also analyzed in Westphal and Sornette (2020b), for cases of Ornstein-Uhlenbeck κ_t as well as constant κ, for reference.
As explained in Westphal and Sornette (2020b), and shown in Harras et al. (2012) and Kaizoji et al. (2015), under Ornstein-Uhlenbeck κ_t, the noise trader dynamics follow generalized standard Ising models. These models present characteristic shifts between disordered and ordered regimes. Essentially, noise traders go from phases where there is a great disparity among individuals, with a vanishing average demand, to phases where there is a dominating collective opinion one way or the other (positive or negative excess demand). The collectively organized regimes can be causes of super-exponential bubbles in this context. Transition periods are related to periods of high volatility.
In order to qualitatively examine the behavior of the market model, we provide two realizations of the market, illustrated in figures 1.1 and 1.2, under Ornstein-Uhlenbeck κ_t and constant κ_t respectively. These figures are based on the same random seed. Note that in accordance with the parameter values of table 1.1, the Ornstein-Uhlenbeck κ_t used in figure 1.1 has an expected value equal to the constant κ value used in figure 1.2. However, the standard deviation of the process is high enough that periods of collective organization of the noise traders occur, i.e. there are transient periods of such that $\kappa_t > \kappa_c$. This will be the case in all subsequent experiments as well. We show 5000 time steps, which corresponds to 20 trading years (in our scheme where one time step corresponds to a trading day), and they are set after 2000 initial time steps. This is to avoid the transient effects of the initialization.
The two figures show the path of the price of the risky asset in the top graph, the evolution of the risky fractions of the noise trader and fundamentalist trader in the middle graph, and at the bottom trace the evolution of the wealth of the two types of traders. Note that in figure 1.1, for Ornstein-Uhlenbeck κ_t, the experiment visibly creates positive price bubbles, approximately at times 2500 and 6800. However, in the constant κ case, such bubbles do not occur. Constant κ still accounts for clusters of high volatility, but no spontaneous collective organization of noise traders, when the social imitation strength reaches high enough levels. In the Ornstein-Uhlenbeck case, the bubbles correspond to an approximate doubling of the price at the maximum, followed with a sharp decrease, the whole process occurring over 100 time steps at the most, corresponding to around 5 months in trading days. In general, the noise traders have a high propensity to buy the risky asset, compared to the fundamentalists (for both κ behaviors). Note in the wealth section of figure 1.1, that this causes the price bubbles and crashes to be accompanied by especially pronounced wealth peaks and decreases for the

noise traders. The wealth of the fundamentalists is less volatile in both cases, and while the initial wealth of both is identical, at the end of the frame, the noise trader wealth is lower.

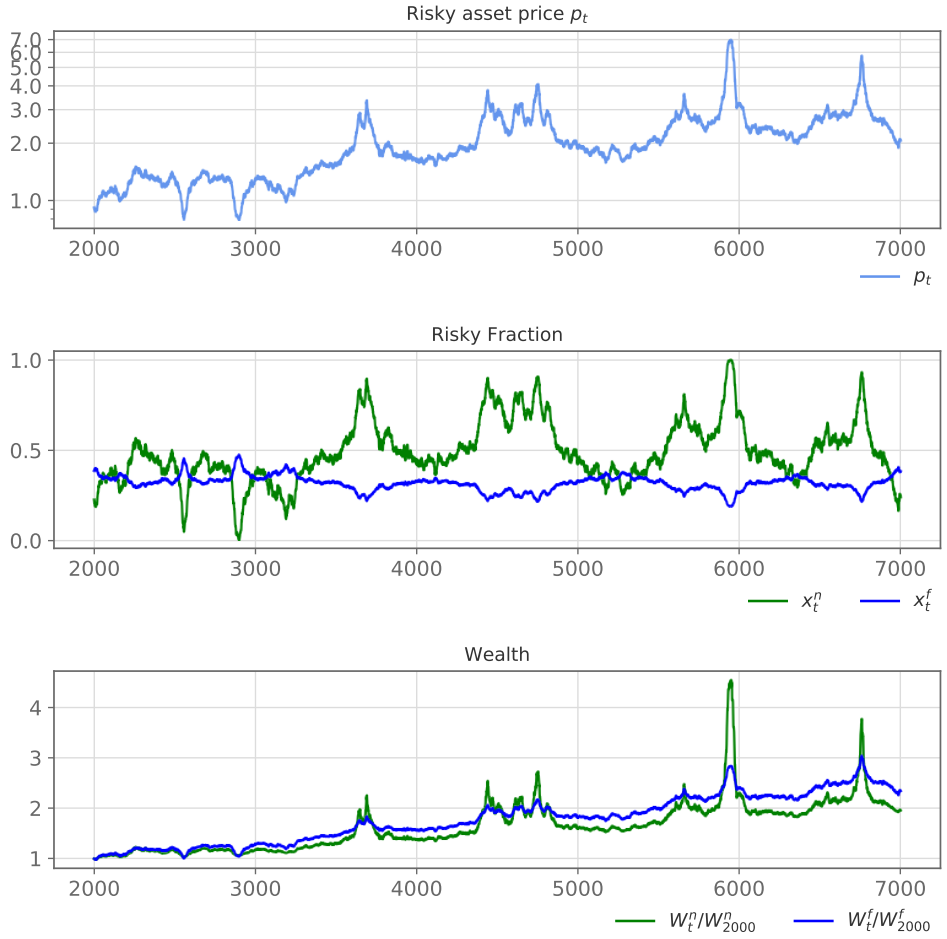

Figure 1.1: One realization of the two types of agent, two-asset market, shown over 5000 time steps (20 years) after 2000 initial time steps, using parameters of table 1.1, for an Ornstein-Uhlenbeck time-varying κ_t. The first line shows the evolution of the price p_t of the risky asset, the second shows the risky fractions of the two types of traders, and the third shows the evolution of the wealth of both traders.

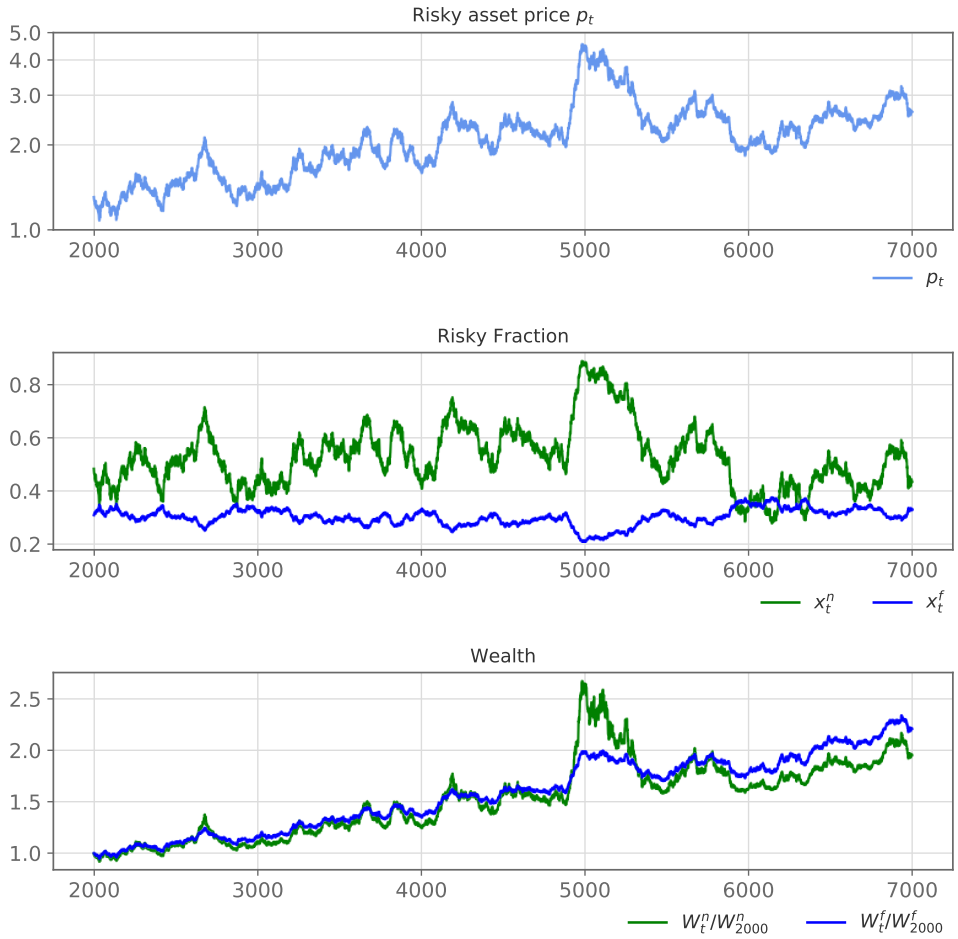

Figure 1.2: One realization of the two types of agent, two-asset market, shown over 5000 time steps (20 years) after 2000 initial time steps, using parameters of table 1.1, for a constant $\kappa = 0.98 \times 0.199375$, that is 98% of the critical value. The first line shows the evolution of the price p_t of the risky asset, the second shows the risky fractions of the two types of traders, and the third shows the evolution of the wealth of both traders.

While these two realizations were chosen representative of the general model behavior, these considerations are not systematic. In an effort to represent the performance of the two types of traders, we show in figure 1.3 the distribution of the Sharpe ratios of both traders' strategies, accompanied with the distribution of Sharpe ratio of the risky asset for reference.

Sharpe ratio is calculated as:

$$\frac{\mu - r_f}{\sqrt{Var(\mu - r_f)}} \tag{1.39}$$

where r_f is the rate of return of the risk-free rate, constant here, $\mu = \frac{1}{T-1}\sum_{t=1}^{T}\frac{W_t - W_{t-1}}{W_{t-1}}$ is the empirical average over time of the returns of the trader strategy, and $Var(\mu - r_f)$ is the empirical standard deviation of the returns of the trader strategy. This measure can be used to compare strategies for experiments with the same time frame T.
In figure 1.3, the risky asset Sharpe ratio is calculated using the price returns instead of the wealth returns, thus corresponding to a buy-and-hold strategy where a trader invests all their wealth in the risky asset and never move it. The calculation gives an empirical daily Sharpe ratio, with data from times $t = 5000$ and $t = 17500$ to make sure not to incorporate transient effects at the initialization. The distribution is taken over 1000 different random seeds.
As expected, the strategy of the fundamentalists have a higher performance

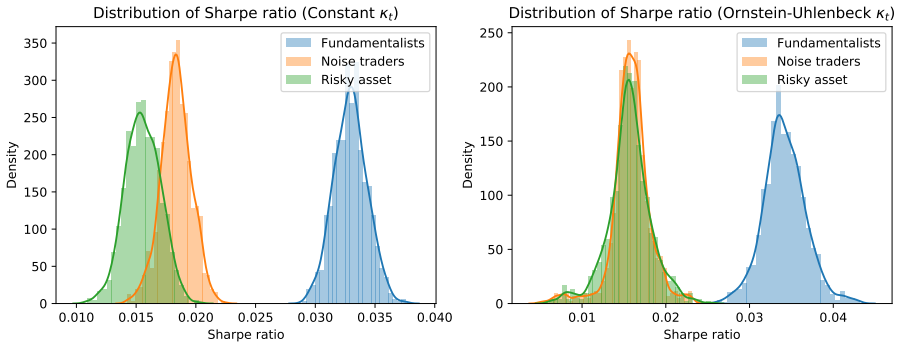

Figure 1.3: Sharpe ratio distribution, over 1000 experiments, of the strategies of the fundamentalists, noise traders and a simple buy-and-hold strategy (labeled risky asset). Sharpe ratios are calculated from daily returns, from time steps 5000 to 17500. Results are shown for constant κ_t on the left, and Ornstein-Uhlenbeck κ_t on the right.

in terms of the risk-adjusted return than the strategy of the noise traders. Note that in the constant κ case, the noise traders slightly outperform the pure risky asset strategy, but in the Ornstein-Uhlenbeck case, the two Sharpe ratios become almost identical, with only a slightly better performance for the noise traders. This can be explained by the occurrence of phases of collective organization of the noise traders, where the dominant opinion is in favor of the risky asset, thus copying the pure risky asset strategy. The

annualized average Sharpe ratios in the constant κ case are 1.56×10^{-2} for the risky asset, 1.83×10^{-2} for the noise traders, and 3.29×10^{-2} for the fundamentalists. In the case of an Ornstein-Uhlenbeck κ_t, they are 1.57×10^{-2} for the risky asset, 1.58×10^{-2}, for the noise traders, and 3.42×10^{-2} for the fundamentalists.

Chapter 2

A three-agent market model: the Dragon Rider trader

In this chapter, we will study the effect of the introduction of a variation of the fundamentalist, a trader using a bubble detecting tool in order to increase its profit. This will help us to understand the influence of our regulatory policies in the following two chapters, to get an understanding of what is attainable. It is a study in the same line as the work of Westphal and Sornette (2020b). In their paper, the two-agent market model we presented in the previous chapter is also extended with a trader similar to the fundamentalist trader – in particular, they both maximize their expected utility with the same CRRA utility function – but they have the advantage of a prediction tool determining if the market is in a bubble. The diagnosis is based on log-periodic power law detection (LPPLs), which has been studied in Sornette et al. (2009), Jiang et al. (2010) and Yan et al. (2012). They observe a reduction in the size and amplitude of the bubbles, with a destabilizing effect when the proportion of Dragon Riders become too great. The trader is purely motivated by wealth gain, and as such, its parameters may be optimized with performance metrics such as Sharpe ratio of the agent's strategy. This is why the new agent is labeled as a "Dragon Rider": compared to the fundamentalist trader, it is using its additional knowledge to profit off expected price soars by investing in them and divesting early when it expects a crash. It has been observed that such agents also have the positive externality of reducing size and number of bubbles in the model, but it is not the original goal, as it would be in the case of open market operations from a central bank. In that sense, this finding is in line with the "invisible hand" argument of Smith (1761), where personal interest of the individual agents, combined with better information, lead to more stable prices and better returns for all.

The approach in this chapter is based on a different detection tool from Westphal and Sornette (2020b), established in Westphal and Sornette (2020a). It

is simpler than the LPPLs detection scheme, and relies on a measure of the excess return to detect the stage of the bubble, and dictate the agent behavior accordingly. Note that in this chapter, we will see that it is also possible to optimize the parameters of this agent with the objective of reducing asset price bubbles. In that case, the agent will using the bubble detection tool as a Dragon Slayer. Dragon Slayers represent agents who use their knowledge about price bubble detection to reduce their amount and size, not arbitrage the information. The Dragon Slayer characteristics of the Dragon Riders described in Westphal and Sornette (2020a), that is their positive externality of reducing the amount and size of bubbles is also what justifies using the same detection tool to create Dragon Slayers in this chapter, and further in the following chapters.
In an attempt to assess the effect of this agent on the market, and use this as a benchmark for the rest of the study, we first will perform a grid exploration to evaluate the impact of various parameter settings. In the Dragon Rider logic, we will try to find optimal parameters in terms of the Sharpe ratio of the trader. Even though these optimal parameters reach unrealistic behavior, they help us get an understanding of the model. Furthermore, we also evaluate what happens if we optimize the parameters in terms of number and size of bubbles, that is, in a Dragon Slayer perspective, in order to understand if regulatory policies or incentives could potentially surpass the Dragon Slayer characteristics of self-interested Dragon Riders, or if instead, the more effective the strategy is for the Dragon Rider, the better it is for reducing bubble occurrence.
We will then use a benchmark parameter set to evaluate the effect to analyze a more sensitive variation of the Dragon Rider formula.

2.1 Description of the three-agent model

2.1.1 Assets

As in the two-agent model, we keep one risky asset, and one risk-free asset. The dynamics stay the same, i.e. the risk-free asset has a fixed return r_f, and the risky asset yields a dividend d_t at each time step, such that $d_t = d_{t-1}(1 + r_t^d)$ with $r_t^d = r_d + \sigma_d u_t$, where $u_t \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$. This is identical to the market setup of the two-agent model.
In similar fashion to the two-agent model, the price dynamics of the risky asset are determined by the equilibrium of supply and demand, however this will include the third trader, as explained in section 2.2. This model features the same fundamentalists and noise traders, as described in chapter 1.

2.1.2 The Dragon Rider/Slayer strategy

Let us now present the new addition to the two-agent, two-asset model from chapter 1. This agent uses a bubble detection tool to calibrate its strategy, whether it is directed towards Dragon Slayer characteristics, or Dragon Rider characteristics. We will first present the features of this tool.

Bubble detection tool

The bubble detection tool that we use is the same as in the internal report of Westphal and Sornette (2020a) . It is based on the exponential moving average y_t of the excess return, that we call overpricing momentum for the rest of the study:

$$y_t = a y_{t-1} + (1-a)(r_t - \bar{r}), \tag{2.1}$$

where $r_t = \dfrac{p_t}{p_{t-1}} - 1$ is the risky asset return, $\bar{r}$ is the long-term return and a is the memory parameter. As explained in the previous chapter, we note that $\dfrac{1}{1-a}$ is the number of time steps characteristic of the frame considered by the exponential moving average.
The idea is that this overpricing momentum y_t informs on the current trend of the market, more specifically on whether the returns are higher or lower than the long-term average, while smoothing the result over a certain window through the filter of averaging. In that way, day-to-day abrupt shocks that are not representative of trends will be ignored by this measure. Note that $y_t = H_t - \bar{r}$.
Now, in order to detect abnormal growth, we set a threshold $y_2 > 0$ for what is considered to be an overly high absolute discrepancy between the return and the long-term average. That is, we use an indicator that will react to the spread $|y_t| - y_2$. The chosen indicator is a logistic function of the following form:

$$\lambda(y) = \frac{1}{1 + e^{-\frac{|y|-y_2}{s}}}, \tag{2.2}$$

where $s > 0$ is the sloppiness of the indicator.
It is a mapping from $\mathbb{R}$ to $\left[\dfrac{1}{1+e^{\frac{y_2}{s}}}, 1\right)$ that measures the likelihood of being in a bubble. When $|y_t|$ is very high compared to y_2, i.e. when the returns have reached levels considered too far above or too far below the acceptable variation y_2, the indicator goes to 1 as the denominator goes down to 1. Conversely, when the exponential moving average of the returns is perfectly equal to the long-term average, the logistic function reaches its minimum $\dfrac{1}{1+e^{\frac{y_2}{s}}}$. Depending on the value of $\dfrac{y_2}{s}$, this value can be closer to 0 (in the case of $y_2 \gg s$) or closer to $\dfrac{1}{2}$ (in the case of $y_2 \ll s$).

While we can make the probability of a crash arbitrarily small in this setup, it is important to note that this probability is never zero. In that sense, the traders who use this detection tool constantly have a slight expectation of a crash. In extreme conditions, i.e. $y_2 \ll s$ or very high sloppiness, the probability of a crash never goes lower than $1/2$. Figure 2.1 displays the effect of sloppiness on the logistic function. Note that we will discuss the precise values of each parameter when we look at the Dragon Rider results in section 2.3. For now, the values used are set for general understanding purposes, and can be considered arbitrary. We can confirm here that this

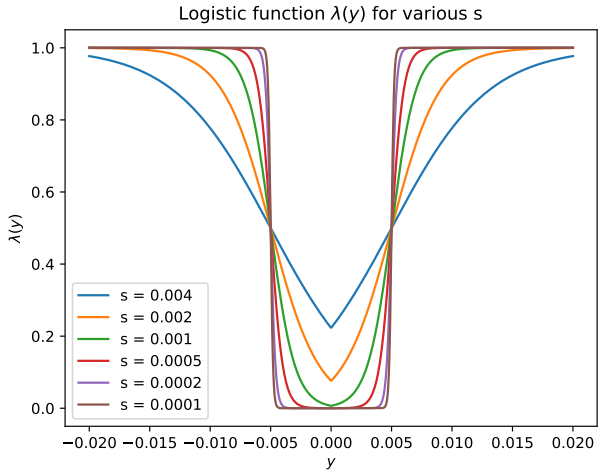

Figure 2.1: Logistic function as a function of the overpricing momentum y, for different values of the sloppiness s. In this figure, $y_2 = 0.005$

version of the logistic function has a minimum $\dfrac{1}{1+e^{\frac{y_2}{s}}}$ that can be closer to 0 if $y_2 \gg s$, or closer to $\dfrac{1}{2}$ if s is higher. The function measures abnormality in the overpricing momentum, the norm being the value $y = 0$.
Figure 2.2 showcases what this bubble detection tool looks like on a price path exhibiting a bubble. The price path corresponds to a realization of the market for the two-trader, two-asset market. Alongside are shown the behavior of the overpricing momentum y_t and the logistic function $\lambda(y_t)$. As we can see, there are two main price peaks at times 600 days and 900 days approximately. The beginning of those peaks are characterised by abnormal increase in the overpricing momentum, that is increasing deviation of the exponential moving average of the returns, from the long-term average return. In turn, $\lambda(y_t)$ goes to 1 whenever the absolute value of y_t exceeds the threshold y_2.

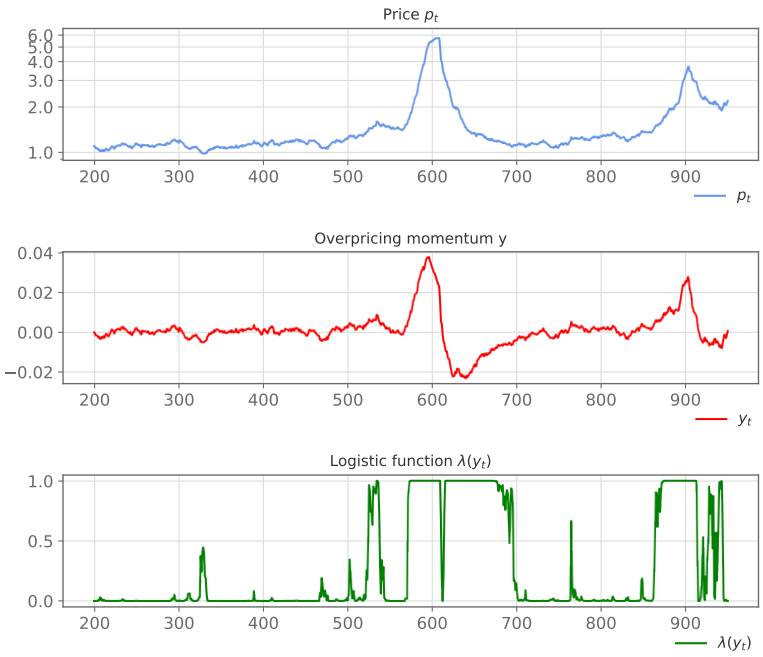

Figure 2.2: Example of overpricing momentum y_t and logistic function $\lambda(y_t)$ for a price path of the two-trader, two-asset model. In this figure, $a = 0.95$, corresponding to a 20 day window, $s = 0.0005$, and $y_2 = 0.005$.

Design of the trader

The Dragon Riders are designed as a variation of the fundamentalist traders. Just as the fundamentalists, their goal is to maximize their individual expected utility at each time step. The Dragon Riders also have the same CRRA utility. Since the characteristics are identical among Dragon Riders, the collective behavior of this group of traders can be summarized with a representative agent for the aggregated wealth of all the traders. This representative agent is set to have the same coefficient of relative risk aversion $\gamma = -x\dfrac{u''(x)}{u(x)}$ as the representative agent of the fundamentalist traders. The reason behind all those similarities is the fact that the Dragon Riders are meant to be regular investors like the fundamentalists, only with the difference of added knowledge about bubble detection. They have the freedom to invest a fraction of its allocated wealth in the risky asset, and the remainder is invested in the risk-free asset. The risky fraction is denoted x_t^{DR}. Again, as in chapter 1, the risky fraction is set to remain in $[0, 1]$, that is, no borrowing or short-selling is allowed. Their wealth structure also remains as follows:

$$W_t^{DR} = W_{t-1}^{DR}\left[1 + r_f + x_{t-1}^{DR}\left(r_t + \frac{d_t}{p_{t-1}} - r_f\right)\right]. \tag{2.3}$$

However, the Dragon Rider differs from the fundamentalist in the sense that they show an improved estimation of the expected return at time t. In the case of the fundamentalist, the expected value of the return at time t, used at $t-1$ by the trader in their optimization, i.e. $E_{r_t} := E_{t-1}[r_t]$, is taken to be equal to the long-term average return $\bar{r}$. This is a default best guess for the fundamentalist.
In the case of the Dragon Rider, the estimation of E_{r_t} is improved using the bubble diagnostic $\lambda(y_t)$ mentioned in 2.1.2, in the following way:

$$E^{DR}_{r_t} = w(\underbrace{(1-\lambda(y_t))\cdot y_t}_{I} - \underbrace{\lambda(y_t)\cdot \text{sign}(y_t) y_2}_{II}) + \bar{r}. \tag{2.4}$$

And then the risky fraction of the Dragon Rider is chosen in an analogous manner to that of the fundamentalist, through utility maximization using this estimate:

$$x^f_{t-1} \approx \frac{E^{DR}_{r_t} + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_f}{\gamma\sigma^2}. \tag{2.5}$$

First, note that when the exponential moving average of the excess return goes closer to zero, i.e. the returns are close to $\bar{r}$, so is the expectation of the return. This is necessary and also showcases the fact that this agent is an extension of the fundamentalist trader, who always uses $\bar{r}$.
Figures 2.3 and 2.4 show an examples of the behavior of the expected return formula for the Dragon Rider, on a large bubble. Note that when the threshold is lower, the Dragon Rider divests slightly too early from the risky asset, before the bubble crash. Now, as explained before, when $\lambda(y_t)$ grows close to 1, it indicates that y_t has reached abnormally large values, meaning that the exponential moving average of the returns has strayed too far from the acceptable value threshold y_2. This is where part II of the equation is important, and part I vanishes. It pushes the expected return in the opposite direction of the sign of the overpricing momentum $sgn(y_t)$, so that if the price is exhibiting a positive bubble, the Dragon Rider divests from the risky asset, and vice versa. It is a set correction factor of y_2 because this corresponds to the approximate value of the expected size of the bubble when the agent sells the asset, i.e. the expected size of the crash. Making part II equal to $-\lambda(y_t)y_t$ directly also makes sense, but in practice the variations were too volatile as a strategy.
Conversely, when $1-\lambda(y_t)$ grows close to one, it means that the rates of return are within a reasonable distance of the average return $\bar{r}$, i.e. the distance is not too far from y_2. This is where the level confidence in the existence of the bubble, controlled by s, is important, because it corresponds to situations where $|y_t| - y_2 \ll s$. In those cases, Dragon Riders have a tendency to invest in the risky asset in the same direction of the overpricing momentum. They do it proportionally to the overpricing momentum itself y_t, such that they

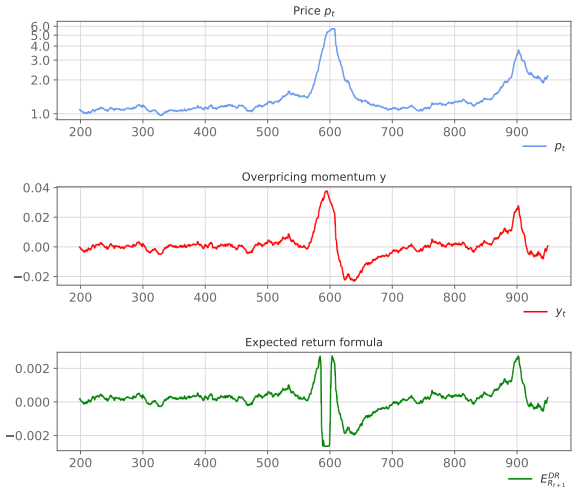

Figure 2.3: Expected return formula for the Dragon Rider, for a price path of the two-trader, two-asset model. In this figure, a=0.95, corresponding to a 20 day window, s=0.0005, and $y_2 = 0.03$

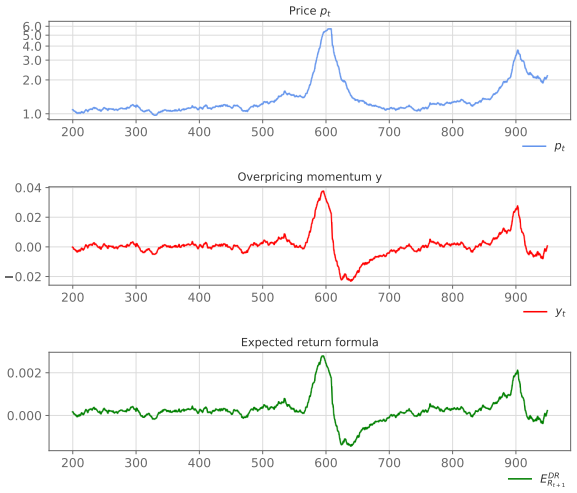

Figure 2.4: Expected return formula for the Dragon Rider, for a price pathof the two-trader, two-asset model. In this figure,a=0.95, correspondingto a 20 day window,s=0.0005, and $y_2 = 0.05$

are only lightly nudging their investment strategy either way when y_t stays close to zero.
There is, however, a critical point around y_2, where the Dragon shifts its strategy from part I to part II. For example, in the case of a positive bubble, the design of the formula will cause the Dragon Rider to go from being inclined to invest in the risky asset to the opposite, compared to a the standard set by a constant $\bar{r}$ return. The effect of the shift around the critical point is continuous due to $\lambda(y_t)$.
Finally, the amount of the corrections is proportional to a dampening factor w, used to make sure that the expected return doesn't exhibit unreasonably high corrections.

2.2 Price derivation

The equilibrium of supply and demand still holds, however it features one more agent.

$$\begin{aligned}
0 = \;& \Delta D_t^f + \Delta D_t^n + \Delta D_t^{DR} \\
= \;& W_{t-1}^f \left(x_t^f \left[1 + r_f + x_{t-1}^f \left(r_t + \frac{d_t}{p_{t-1}} - r_f \right) \right] - x_{t-1}^f \frac{p_t}{p_{t-1}} \right) \\
& + W_{t-1}^n \left(x_t^n \left[1 + r_f + x_{t-1}^n \left(r_t + \frac{d_t}{p_{t-1}} - r_f \right) \right] - x_{t-1}^n \frac{p_t}{p_{t-1}} \right) \\
& + W_{t-1}^{DR} \left(x_t^{DR} \left[1 + r_f + x_{t-1}^{DR} \left(r_t + \frac{d_t}{p_{t-1}} - r_f \right) \right] - x_{t-1}^{DR} \frac{p_t}{p_{t-1}} \right)
\end{aligned}$$

We recall that the writing of the risky fraction of the Dragon Rider are similar to that of the fundamentalist, with the exception of E_{r_t}, which is variable in the case of the Dragon Rider. It follows that similar manipulations of this equation as in 1.4.2 in chapter 1 lead to a quadratic equation of the price with changed parameters used in the simulations.

2.3 Dragon Rider grid exploration

In this section, we analyze the scope of the effect Dragon Rider model. On the one hand, as in the work of Westphal and Sornette (2020a), we explain the effect of parameters a, y_2, and s on the Dragon Rider behavior, especially the formula of the expected return. On the other hand, we perform a grid exploration in order to assess two things: the theoretical extent of the effect of Dragon Riders on the market, specifically in reducing the size and number of bubbles, and the effect of optimal Dragon Riders (in the sense of self-interested traders who optimize for wealth related measures).
This study will allow us, in chapters 3 and 4, to draw more informed comparisons between the effect of Dragon Riders and the effects of a direct strategy by a regulatory institution. Using the analysis, we will assess whether direct policy can be as efficient as the self-regulation of markets, in this simplistic setting.

Measures for the rest of the study

In the following, we will be measuring the effect of Dragon Riders on the market, but also the performance of the traders. In order to do so, we use different measures. We divide them in two categories: the measures designed to assess the effectiveness of the trader, and the measures used to evaluate the risky asset. The former are typically used in the perspective of trading strategy validity, and focus on the gains of the traders, while

the latter are related to both volatility and the emergence of bubbles in the market. Note that these measure will also be useful in chapter 3, when we study the impact of open market operations on the model.

Sharpe ratio of the trader strategies We use Sharpe ratio to assess the effectiveness of the strategies of the traders. Sharpe ratio of the Dragon Rider is a relevant measure of performance of the trader. As mentioned in chapter 1, it is calculated as:

$$\frac{\mu^{DR} - r_f}{\sigma^{DR}} \tag{2.6}$$

where $\mu^{DR} = \frac{1}{T-1}\sum_{t=1}^{T}\frac{W_t^{DR} - W_{t-1}^{DR}}{W_{t-1}^{DR}}$ is the empirical average over time of the returns of the trader strategy, and σ^{DR} is the empirical standard deviation of the returns of the trader strategy. This measure can be used to compare strategies for experiments with the same time frame T. In the context of our grid exploration, the total number of time steps was 10000, but we only started the calculations at $t = 2000$, to allow for the initial transient phase to subside.

Measures related to the price behavior We use moments of the returns distribution, and two additional measures to assess the characteristics of the market:

1. *Volatility*: a standard measure of market stability, calculated as the empirical standard deviation of the returns

$$v_p = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\frac{p_t - p_{t-1}}{p_{t-1}} - \mu_p\right)^2} \tag{2.7}$$

   where we note the empirical average of the returns $\mu_p = \frac{1}{T-1}\sum_{t=1}^{T}\frac{p_t - p_{t-1}}{p_{t-1}}$. While high volatility can indicate propensity of the market to be risky, it is not a complete assessment of the price path in the sense that does not differentiate positive or negative deviations from the norm. Sharp increases in price are not an undesired quality in a price path, but rather the sharp decreases.

2. Conversely, *skewness* is a measure that separates upward and downward motions in price. It is calculated as:

$$\frac{1}{T}\sum_{t=1}^{T}\left(\frac{\frac{p_t - p_{t-1}}{p_{t-1}} - \mu_p}{v_p}\right)^3 \tag{2.8}$$

3. The *kurtosis* is the fourth standardized moment of the returns, and essentially captures the amount of outliers on either side of the mean. The normal distribution has a kurtosis of 3, therefore it is usual to measure the excess kurtosis as the difference:

$$\frac{1}{T}\sum_{t=1}^{T}\left(\frac{\frac{p_t-p_{t-1}}{p_{t-1}}-\mu_p}{v_p}\right)^4-3 \tag{2.9}$$

4. *Value-at-Risk* 5%: it is the lower boundary of the returns in 95% of the cases. Here, the cases are the time steps.

5. *Peak number* Price peaks with time scale k are defined as in Westphal and Sornette (2020b): if p_t is the price at time t, then there is a peak at time t_0 if

$$\forall t \in [t_0-k, t_0+k], p_{t_0} \geq p_t.$$

The time scale k corresponds to the minimum distance between consecutive peaks. We also choose $k = 250$ days ($\sim$ one trading year), as in Westphal and Sornette (2020b).
The measure itself counts the number of peaks in the considered time frame T.

6. *Average drawdown* Peaks are defined as above, while *valleys* are defined as the lowest price point between two peaks.
From there, the *drawdown* associated with a peak p_{t_0}, with nearest next valley p_{t_1} is $log(\frac{p_{t_0}}{p_{t_1}})$. Figure 2.5 shows an example of a peak, valley and corresponding drawdown. The average drawdown is simply its average over all peaks in the time frame (except the last).

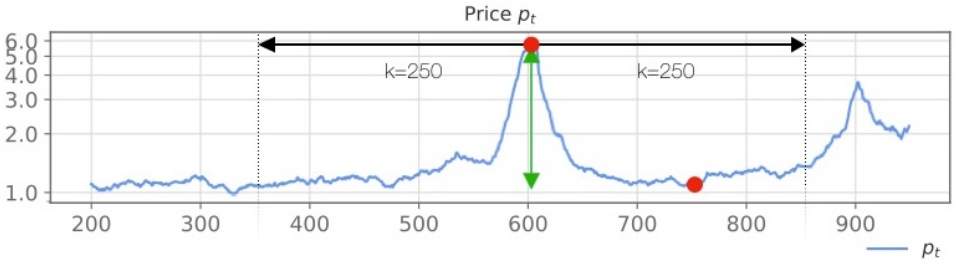

Figure 2.5: Example of a price path where there is a peak, for $k = 250$ days. The red dots are the peak and valley, and the green arrow indicates the drawdown. The price path is generated using the two-asset, two-trader model with Ornstein-Uhlenbeck κ.

Empirical behavior of the model

Let us now recall from report of Westphal and Sornette (2020a) the initial observed effects of Dragon Riders on the market.

Parameter choice Until now, we have not discussed the scale of each parameter of the Dragon Rider. It is important to recall that the setting is artificial, and these parameters are subject to a few decisions made about the market settings. However, the values of the parameters cannot be arbitrary, as they are to be effectively used as detection methods for phenomena that occur in the two-asset, two-trader model of chapter 1. We will see in part 2.3 how to optimize the parameters of the Dragon Rider, but for now, it is important to have an initial idea of what their scale should be, in order to focus the grid exploration better.
The Dragon Rider model works with some exogenous constants, but four parameters can be changed, to modify its behavior. These parameters are a, y_2, s, and w. As explained previously, a is an indicator of the memory of the Dragon Rider (the detection tool considers an exponential moving average of scope $\sim \frac{1}{1-a}$ when calculating the overpricing momentum), y_2 represents the threshold at which the Dragon Rider exits the bubble because the risk of crash is considered too high, and the sloppiness s helps calibrate the sensitivity of the response (as illustrated in Figure 2.1).
For a, we use the memory length $\dfrac{1}{1-a}$ as a reference for what makes sense as a value. It doesn't make sense to go lower than 2 days (it is the extreme boundary for defining an average), which corresponds to $a = 0.5$. More reasonably, a value of around 7 days corresponds to $a = 0.85$. On the other extreme, one trading year, i.e. 250 days, or two trading years correspond to $a = 0.996$ and $a = 0.998$ respectively. Going above makes the notion of the overpricing momentum lose meaning, because the exponential moving average will converge towards the long-term average.
The threshold parameter y_2 should be understood as a multiple of the long-term return $\bar{r}$. At the smallest, it is reasonable to tolerate excess returns of around $\bar{r}$ before exiting a bubble, while at the largest, it is likely that and excess return $50\bar{r}$ precedes a crash. Also note that the expected volatility of the daily returns is very large compared to its expected value (with parameters from chapter 1, it is 125 times higher). This also justifies such a tolerance level for the overpricing momentum. Since we use $\bar{r} = 0.00016$, this corresponds to minimum y_2 at 0.00016, and maximum around 0.008.
The parameter s can be varied at will, but it makes sense to compare it to y_2. As we explained, when s is much larger than y_2, the "probability of being in a bubble" $\lambda(y_t)$ will have a minimum of around $\dfrac{1}{2}$. This could be problematic since the Dragon Rider behavior risks becoming erratic. As per the recommendations of the report of Westphal and Sornette (2020a), we keep s

between 0.0001 and 0.002 with this range of y_2. With these ranges, the minimal value of the minimum of the logistic function is as low as 1.8×10^{-35}. Lastly, we choose not to vary w, as this one is set as in Westphal and Sornette (2020a). The idea is that the risky fraction of the Dragon Rider is, by construction of the expected return formula, symmetric around the initial risky fraction. The initial risky fraction is chosen at 0.3, like the fundamentalist traders, since they are an extension of them. The scaling w of the formula should therefore at least ensure that the risky fraction stays within 0 and 1. For that reason, the scaling factor w is chosen has the following upper boundary:

$$w = \frac{1}{y_2}\left(\bar{r} - r_f + \frac{d_0}{p_0}(1+\bar{r})\right) \tag{2.10}$$

which is chosen as a value for w. After considerations on the effect of each parameter variation, all other parameters fixed as specified in 1 the report concludes that the combination $a = 0.95$, $y_2 = 0.008$, $s = 0.0005$, $w = 0.035$ has good performance in terms of Sharpe ratio of the Dragon Rider.

Effect on the market The paper of Westphal and Sornette (2020a) shows the effect on the market of the Dragon Riders with parameters listed above. Interestingly, the more Dragon Riders are present in the market, the higher reduction there is of the average number of peaks as well as the average size of peaks. As a reminder, positive bubbles are identified as price peaks. We recall that peaks with time scale k are defined as follows: if p_t is the price at time t, then there is a peak at time t_0 if

$$\forall t \in [t_0 - k, t_0 + k], p_{t_0} \geq p_t.$$

The performance of other traders is also affected by the introduction of Dragon Riders in the market. The Sharpe ratios of noise traders, fundamentalist traders and Dragon Riders increase as the proportion of Dragon Riders increases from 0% to 50%. That is, the reduction of bubbles on the market seems to increase the performance of both other traders' strategies.

Grid exploration: extent of market impact

We now consider the extent of the market impact that the Dragon Rider can have, by performing a grid exploration on the parameters a, s and y_2. The goal is partly to obtain an increase in Sharpe ratio compared to the parameter set decided on in Westphal and Sornette (2020a), but also to observe what happens if we optimize for different measures, such as number and size of peaks, in particular if optimal market parameters coincide with optimal Dragon Rider parameters.
For the latter purposes, we draw a so-called correlation analysis, to see if optimizing the parameters with respect to Dragon Rider wealth gain yields

the same results as optimizing them with respect to Dragon Slayer characteristics. In that framework, high correlation would go in the direction of an "invisible hand" of Smith (1761) type of hypothesis, where the utility of Dragon Rider traders is aligned with the overall stability of the market, reducing the size and frequency of bubbles as they increase their wealth. Low correlation would tend to indicate that regulatory entities targeting bubbles specifically could have a stronger impact on them than the Dragon Riders.

Grid values We assess the effect of variations of parameters a, y_2, and s, on these measures. We use the following ranges:

$$\begin{aligned} a &\in [0.5, 0.85, 0.98, 0.992, 0.996, 0.998] \\ s &\in [0.0001, 0.0002, 0.0005, 0.001, 0.002] \\ y_2 &\in [0.0001, 0.002, 0.004, 0.006, 0.008] \end{aligned}$$

The search explores all 175 combinations. Other parameters are set to a the parameters from chapter 1. We perform the grid search for Ornstein-Uhlenbeck κ_t and constant κ_t. The relevant result is the former, since constant κ does not account for bubble behavior, but we will use the data of constant κ later. As mentioned above, the total number of time steps of the grid exploration experiments is 10000. For Sharpe ratio calculations, we start at $t = 5000$, to allow for the initial transient phase to subside.
A very important parameter in this search is the initial wealth of the Dragon Rider. In order to keep the scale of the total amount of initial wealth traded constant over the different experiments, including in the following experiments where the effect of the initial wealth of the Dragon Riders is studied, we maintain a constant initial wealth for the ensemble of fundamentalists and Dragon Riders, i.e. $W_0^f + W_0^{DR} = 10^9$, to match the two-asset model. For this reason, we refer to the wealth of the Dragon Rider in percentages. In our grid exploration, we use a 15% amount of Dragon Riders. According to Westphal and Sornette (2020a), this proportion of Dragon Riders does not completely suppress bubbles as in higher proportions, but starts to have an effect. This allows us to compare the extent of the effect on the bubbles.
For each combination, we average each measure over the same 100 seeds. We then look at average parameter values over the top 10%, 20%, and 50% performing combinations, for all the different parameter values. Given the results, we observe the top 10% in our findings. For instance, if we are using a measure like average drawdown, where a lower number is more desirable, we will use sort the combinations in ascending order and take the first values. Conversely, if the measure is Sharpe ratio, where a high number is desirable, we will sort the combinations in descending order and take the first values.
In the following, we will observe the behavior of Dragon Riders set up with

	Mean	Std
a	0.97	0.02
s	0.0005	0.0002
y_2	0.0001	0.00004

Sharpe ratio

Mean	Std
0.93	0.05
0.0006	0.0003
0.0001	0.00003

Peak number

Mean	Std
0.98	0.01
0.0006	0.0002
0.0001	0.00005

Avg drawdown

Mean	Std
0.97	0.03
0.0006	0.0001
0.0001	0.00002

Avg dd$\times$p nb

Table 2.1: Average and standard deviation of the parameter values of a, s, y_2 over the top 10% performing combinations, under Ornstein-Uhlenbeck. From left to right, top performing combinations with respect to Sharpe ratio of the Dragon Rider, price peak number over the 10000 time steps time frame, average drawdown on those peaks, and the product of the two (average drawdown times peak number).

optimal combinations of parameters, in the case of optimization for Dragon Rider Sharpe ratio, or average drawdown, peak number, and skewness.

Optimal combinations for Ornstein-Uhlenbeck κ_t

We obtain, the following optimal combinations of a, s and y_2, when we average the values of top performing 10% combinations in 4 parameters, and obtain a combination with associated standard deviation. We recall the combination from the report of Westphal and Sornette (2020a) for reference: $a = 0.95, s = 0.0005, y2 = 0.008$.

We showcase in 2.1 the measures linked to the performance of the Dragon Rider, as a trader, and the measures related to the amount and size of bubbles in the price path, because those are the ones we are interested in optimizing for, as mentioned above. Note that the average values taken over the top 10% of the combinations in terms of skewness is identical to that of the average drawdown. This signifies high correlation between the act of optimizing for either measure. We will quantify this notion in the following paragraph.
Also note that all of the combinations in 2.1 are rather similar, especially if we consider the ranges given by the standard deviation we observe. This tends to mean that the optimal combinations if optimizing for the Sharpe ratio of the Dragon Rider, are the same as the ones obtained optimizing for reduction of size and amount of peaks. However, this is only true for the top 10% performing ones, as far as we can observe here. The more systematic correlation study in the next paragraph will help clarify this.
Also note that the threshold parameter y_2 is reaching the lower boundary of the exploration. We recall that this lower boundary is chosen because it is assumed that fluctuations of the returns around $\bar{r}$ can reasonably be set to a minimum of $\bar{r}$, without representing a real threat of bubble. We will see

in an example that this already leads to erratic behavior on the part of the Dragon Rider, which would probably be tempered in real life by transaction costs, and as such, it is not realistic to extend the exploration further. However, these very small values of y_2 lead to a better control of bubbles and better Sharpe ratio overall. Note that some good control of bubbles can be obtained with high y_2 as well, but not as good as very low y_2 in this study.
In figure 2.6 is an example of the behavior of the model using 15% of Dragon Riders with parameters a, y_2, and s optimized for Sharpe ratio of the Dragon Rider, taken on one representative seed. In figure 2.7, we show behavior of the model using 15% of Dragon Riders with the initial combination of Westphal and Sornette (2020a), on the same seed and range of time steps.
As we can see, compared to the behavior of the model in figure 2.7, the optimized model in figure 2.6 has a very smooth price path and shows a dramatic and constant wealth increase for the Dragon Rider. This is done through very sensitive and frequent switching of the amount of wealth invested in either asset. The risky fraction of the model using the parameters of Westphal and Sornette (2020a) is less polarized, never reaching 0 or 1, while the optimized model switches back and forth from those values quite often. This is explained by the more dramatic variations in the expected return values used by the Dragon Rider to calibrate their wealth, as shown in the bottom row. Indeed, the small y_2 value leads the Dragon Rider to interpret small deviations from the average return as a bubble. While in the model using the parameters of Westphal and Sornette (2020a), the value of E_{r_t} has a maximum around 0.0005, the boundary of E_{r_t} in the optimized case is above 0.01.

In figure we also observe the Sharpe Ratio distribution of the three traders, over 1000 seeds, for the Sharpe ratio of Dragon Riders optimized parameter setting of 2.1. The methodology is the same as in chapter 1, to allow for comparison. The data is taken from simulations longer than those of the grid exploration, and also taken from times $t = 5000$ and $t = 17500$ to make sure not to incorporate transient effects at the initialization, as in 1. Compared to the standard parameters of $a = 0.95, s = 0.0005, y2 = 0.008$, the Sharpe ratios are much higher than with no intervention. Note that the Dragon Riders have higher Sharpe ratios than the fundamentalists and noise traders, which indicates that they benefit from the bubble detection tool, and do in fact arbitrage this information to their advantage. Overall, the results indicate that the Sharpe ratio of the Dragon Rider can be dramatically increased through parameter tweaking. Similar manipulations can also improve the other measures significantly. However, they lead to unrealistic or unreasonable Dragon Rider behavior and therefore cannot be used as a benchmark for later experiments. It is however interesting to see that the standard parameters still provide significant bubble reduction, as emphasized in Westphal

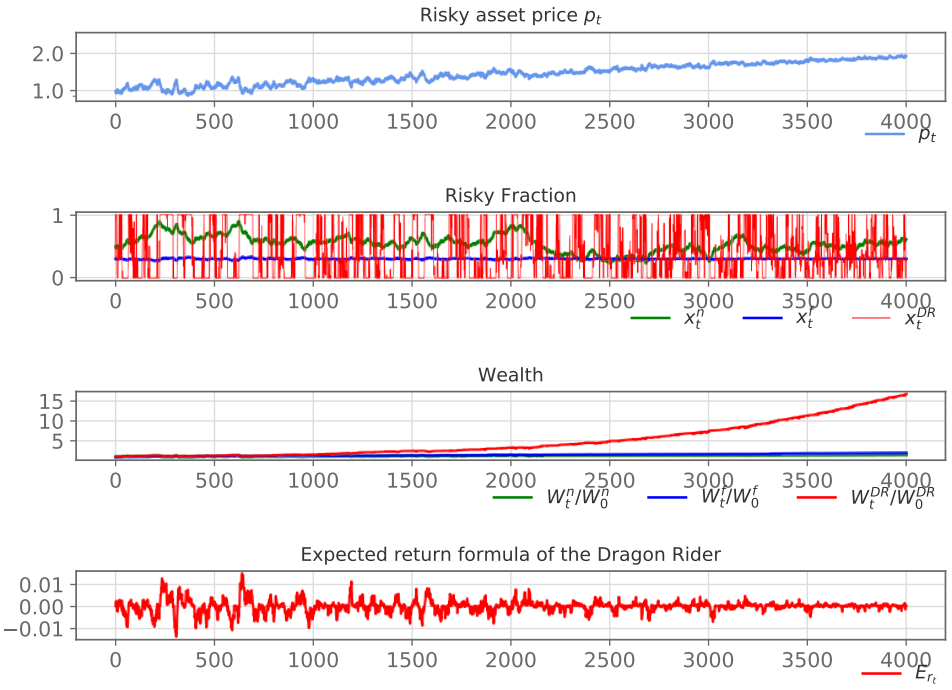

Figure 2.6: Behavior of the three-agent market model with 15% of Dragon Riders, and parameters set in chapter 1, with $a = 0.97, s = 0.0005, y2 = 0.0001$, the optimal combination obtained with respect to the Dragon Rider Sharpe ratio. The top panel shows the price path, the second panel shows the risky fractions of the three traders over time, the third panel shows the normalized wealth (with respect to initial wealth) of the three types of traders, and the last panel shows the evolution of the formula of the expected return E_{r_t} used by the Dragon Rider in the wealth allocation decision.

and Sornette (2020a).
An interesting observation we can make is that the optimization of the Sharpe ratio of the Dragon Rider leads to similar results as the optimization of the bubble size and amount related parameters. However, this convergence is only observed on the few best performing combinations, that also lead to unrealistic behavior of the trader. In the next analysis, we attempt to draw correlation results from a wider scope of combinations.

Correlation analysis

An interrogation we aim to answer is whether Dragon Rider optimizing the parameter values for their own utility maximization, for example getting the

Figure 2.7: Behavior of the three-agent market model with 15% of Dragon Riders, and parameters set in chapter 1, with $a = 0.95, s = 0.0005, y2 = 0.008$, as in Westphal and Sornette (2020a). The top panel shows the price path, the second panel shows the risky fractions of the three traders over time, the third panel shows the normalized wealth (with respect to initial wealth) of the three types of traders, and the last panel shows the evolution of the formula of the expected return E_{r_t} used by the Dragon Rider in the wealth allocation decision.

best Sharpe ratio, leads to a stabilization of the market in general.
To this end, we observe, for each combination (of the 175 that the grid search covers), average performance values over 100 seeds. This means we have arrays of the variation of the empirical expected value over 100 tries of each performance measure along the 175 combinations. We measure correlation between the 175 length arrays of different performances, using the Pearson coefficient. If $X = (x_1, ..., x_N)$ and $Y = (y_1, ...y_N)$ are two arrays of performance variation along the same set of N combinations, we calculate

$$\rho(X, Y) = \frac{Cov(X, Y)}{\sigma_X \sigma_Y},$$

where σ_X and σ_Y are the sample standard deviations of X and Y and $Cov(X, Y) = \frac{1}{N-1} \sum_{i=1}^{N} (x_i - \bar{X})(y_i - \bar{Y})$ is the sample covariance of X and Y.

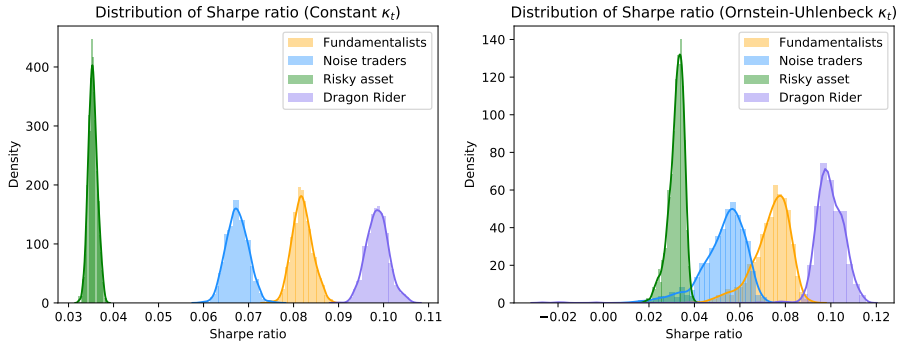

Figure 2.8: Distribution over 1000 seeds of the Sharpe ratios of the Fundamentalists, Noise traders, Dragon Riders and pure risky asset strategy, for a three-agent model with parameters from chapter 1, $a = 95$, $s = 0.0005$ and $y_2 = 0.008$ and a wealth level of 15%. Sharpe ratios are calculated between $t = 5000$ and $t = 17500$.

An absolute value of $\rho(X, Y)$ close to one indicates that optimization for X and Y are highly correlated. The sign indicates the orientation of the correlation. We highlight some important takeaways form the study in table 2.2. This table indicates that for Ornstein-Uhlenbeck κ, the optimization with respect to Sharpe ratio of the Dragon Rider is not very correlated with price related measures. The two exceptions are the peak number, which is significantly negatively correlated with $\rho = -0.47$, and the excess kurtosis of the returns which is positively correlated with $\rho = 0.62$. While this seems slightly counter intuitive, the peak number and excess kurtosis are rather negatively correlated with $\rho = -0.51$. It seems that optimizing for Sharpe ratio of the Dragon Rider might be accompanied with the effect of also lowering the amount of peaks, to some level of probability (the correlation factor is lower than 1), while increasing the excess kurtosis. There might also be, but with lower probability, an increase in Value-at-Risk at the 5% level, and with similar likelihood, a reduction in the compound measure of average drawdown times peak number.
As for the correlations between different market related measures, we first highlight that the product of average drawdown and peak number is more correlated with average drawdown than peak number. This is mostly due to the fact that there is less variability of the number of peaks, than that of the average drawdown. Note that the definition of a peak does not specify any height, which might explain why there might technically remain peaks in relatively smooth price paths. Also note that the peak number is negatively correlated with Value-at-Risk is negatively correlated with peak number, at a -0.58 level, and even more so with average drawdown, at a -0.95 level. This is explained by the fact that a higher amount of peaks leads a more

X measure	Y measure	$\rho(X, Y)$
Sharpe ratio DR	Peak number	$-\mathbf{0.47}$
Sharpe ratio DR	Average drawdown	-0.06
Sharpe ratio DR	Avg drawdown$\times$ Peak nb	-0.17
Sharpe ratio DR	Excess kurtosis	**0.62**
Sharpe ratio DR	Skewness	0.05
Sharpe ratio DR	Volatility	-0.12
Sharpe ratio DR	VaR 5%	0.18
Avg drawdown	Avg drawdown$\times$ Peak nb	**0.98**
Peak nb	Avg drawdown$\times$ Peak nb	**0.66**
Peak number	Excess kurtosis	-0.51
Skewness	Excess kurtosis	-0.32
Volatility	Excess kurtosis	0.33
Peak number	Volatility	0.54
Peak number	VaR 5%	-0.58
Avg drawdown	VaR 5%	$-\mathbf{0.95}$

Table 2.2: Pearson correlation coefficients $\rho(X, Y)$ between size 175 arrays X and Y, corresponding to the empirical expected value calculated over 100 seeds, for each parameter combination of 175 combinations. This was calculated using Ornstein-Uhlenbeck κ. The table is separated in two parts by a double line. The top part of the table shows correlations between the Sharpe ratio of the Dragon Rider, and various price-related measures. In it, the numbers in bold are close or higher than 0.5. The bottom part features correlations between said price-related measures for clarification. In that section, the numbers in bold are close to 1.

frequent occurrence of very negative returns, thus also to a worse Value-at-Risk, i.e. a lower one. Of course, one might ask if these very negative returns caused by bubble crashes are offset by the sharp growth at the beginning of the bubble, but there tends to be asymmetry between the left side and the right side of the peak, in the sense that crashes are faster than growth in at least some events. This is also nicely represented by the fact that there is even further negative correlation when it comes to average drawdown, compared to peak number alone.
While Sharpe ratio optimization of the Dragon Riders might not be completely uncorrelated with reduction of bubbles, we can observe that in this framework, the market does regulate itself in an "invisible hand" fashion. A Dragon Rider who does not care about the market still improves it to some extent. However, the correlation shows that this stabilization is imperfect, and a different set of parameters than the Dragon Rider's can improve the market even more.

The same analysis is performed on constant κ, in table 2.3. The results are noticeably different. In this context, is seems that Sharpe ratio optimization for the Dragon Rider is rather equivalent to aiming directly at reducing peak number and average drawdown, or volatility, or increasing Value-at-Risk at the 5% level. This goes in favor of self-regulation in the absence of regimes of spontaneous alignment of noise traders, in Ising fashion, as mentioned above, which are exclusive to Ornstein-Uhlenbeck κ_t in our model.
An interesting point that bears noticing is the fact that the correlations between skewness and excess kurtosis, and volatility and excess kurtosis, switch signs in the constant κ case compared to the Ornstein-Uhlenbeck case. They are also much closer to 1 in absolute value than in the Ornstein-Uhlenbeck case, where the correlation approached 0.3 in absolute value. This is due to the structural differences between the price patterns of the two cases. A caveat to this analysis is that the impact of the Dragon Rider for constant κ is much smaller than for Ornstein-Uhlenbeck κ. The absence of bubbles makes the effects of the model much less significant. In that sense, the correlation analysis is mostly confirming that fact.

X measure	Y measure	$\rho(X, Y)$
Sharpe ratio DR	Peak number	-0.94
Sharpe ratio DR	Average drawdown	-0.81
Sharpe ratio DR	Avg drawdown$\times$ Peak nb	-0.94
Sharpe ratio DR	Excess kurtosis	0.95
Sharpe ratio DR	Skewness	0.78
Sharpe ratio DR	Volatility	-0.96
Sharpe ratio DR	VaR 5%	0.96
Avg drawdown	Avg drawdown$\times$ Peak nb	0.94
Peak number	Excess kurtosis	-0.94
Skewness	Excess kurtosis	0.71
Volatility	Excess kurtosis	-0.91
Peak number	Volatility	0.88
Peak number	VaR 5%	-0.87
Avg drawdown	VaR 5%	-0.90

Table 2.3: Pearson correlation coefficients $\rho(X, Y)$ between size 175 arrays X and Y, corresponding to the empirical expected value calculated over 100 seeds, for each parameter combination of 175 combinations. Table was calculated using constant κ. The top part of the table shows correlations between the Sharpe ratio of the Dragon Rider, and various price-related measures. The bottom part features correlations between said price-related measures for clarification.

Analysis of the effects of a time-varying memory parameter

In this section, we explore an extension of the Dragon Rider behavior. Until now, the memory parameter a was constant. We recall that a is related to the memory of the Dragon Rider. More precisely, $\frac{1}{1-a}$ represents the window, in number of days, the Dragon Rider uses to compare price. In an effort to improve bubble detection, this memory is made variable in the following model generalization:

$$a(\lambda) = a_0 \times (1-\lambda)^{\beta}$$

a is made dependent on λ such that during bubbles, when λ is close to 1, the memory parameter gets shorter. That means that the model is more sensitive to price movements, when it is more likely that the risky asset is in a price bubble. Intuitively, a strictly positive β leads to a reduction the memory parameter as the Dragon Rider detects bubble regime. The constant case corresponds to $\beta = 0$. $\beta < 1$ makes a a concave function of $1-\lambda$, while $\beta > 1$ makes a a convex function of $1-\lambda$. This means that β controls the sensitivity to λ. Note that since $\lambda \in [0,1]$, this means that for cases of $\beta > 1$, $a(\lambda)$ is lower than a_0, while when $\beta < 1$, $a(\lambda)$ is higher than a_0.
We conduct an analysis of the behavior of the Sharpe ratio of the Dragon Rider trader, using the standard combination of parameters from Westphal and Sornette (2020a): $a = 0.95, s = 0.0005, y2 = 0.008$. In figure 2.9, we show plots of the evolution of several measure's empirical expected value, calculated over 100 seeds, for $T = 10000$, for the case of 15% wealth of Dragon Riders compared to the total of fundamentalists and Riders. The measures considered are average drawdown, peak number, Value-at-Risk at the 5% level, Sharpe ratio of the fundamentalist, the noise trader and the Dragon Rider. As shown in those plots, the increase of β seems to have little effect on the price related measures, and no positive effect. The increase of β also decreases the Sharpe ratios of all three types of traders slightly, but significantly.
In figure 2.10, we show analogous plots for a constant κ. The results are inconclusive in this case, with no strong effects on either measure.
We also attempt the same experiments with higher proportions of Dragon Riders up to 50% of wealth, and obtain similar results with stronger decreases in the Sharpe ratios in the case of Ornstein-Uhlenbeck κ_t. We conclude that this type of detection tool does not help the Dragon Rider increase its accuracy or efficiency.

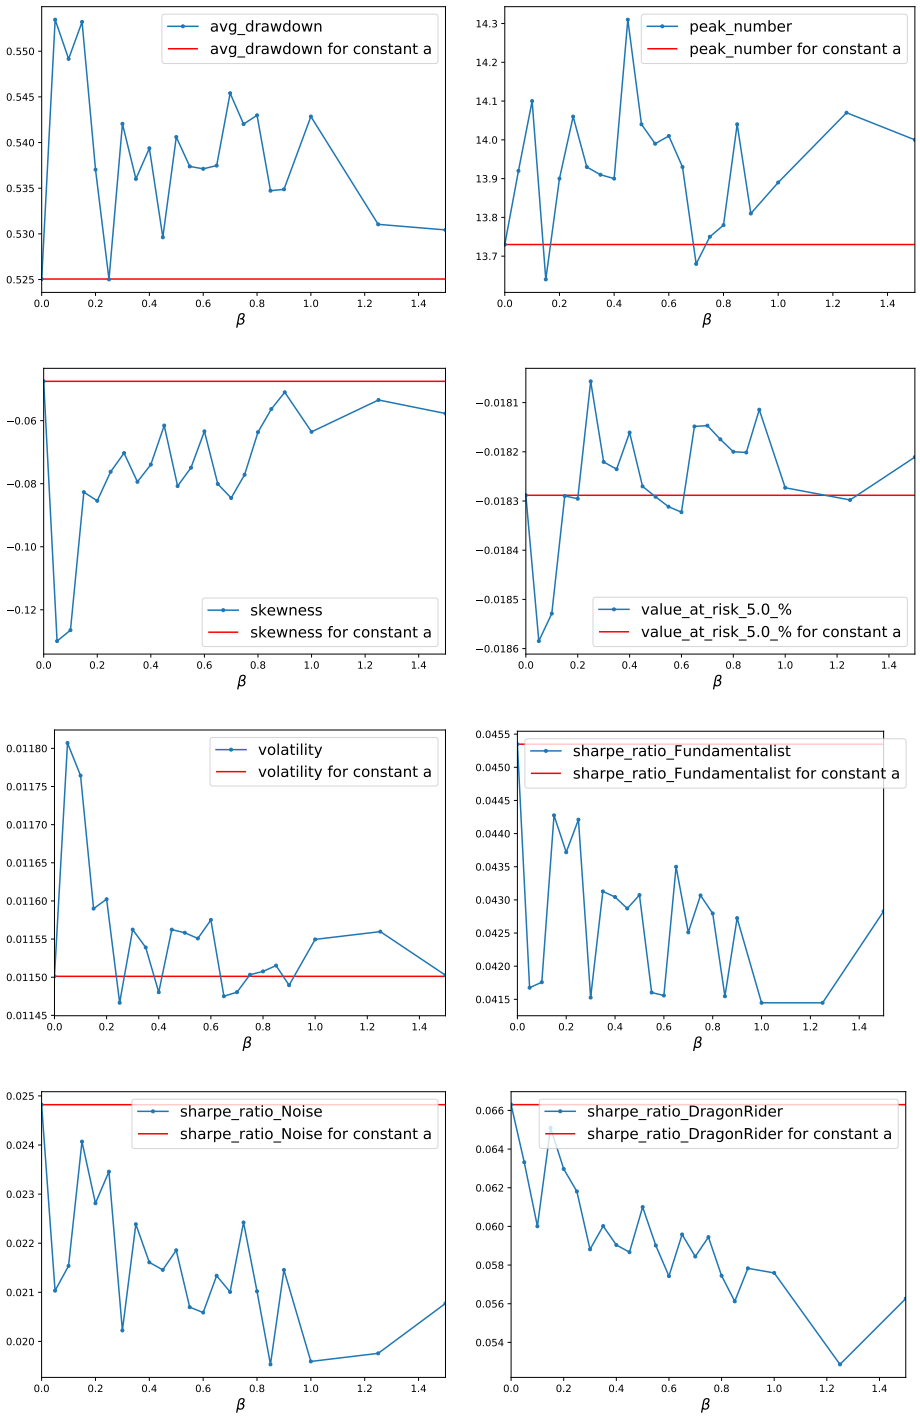

Figure 2.9: Plots of the variation of each measure average over 100 seeds, for different values of β applied to the three-agent model with the standard parameters for the Dragon Rider $a = 0.95, s = 0.0005, y2 = 0.008$. Ornstein-Uhlenbeck κ_t is used. In blue, the variation of each parameter is showed, while in red is the benchmark of constant a behavior. From left to right, top to bottom, the measures considered are average drawdown, peak number, Value-at-Risk at the 5% level, sharpe ratio of the fundamentalist, the noise trader and the Dragon Rider.

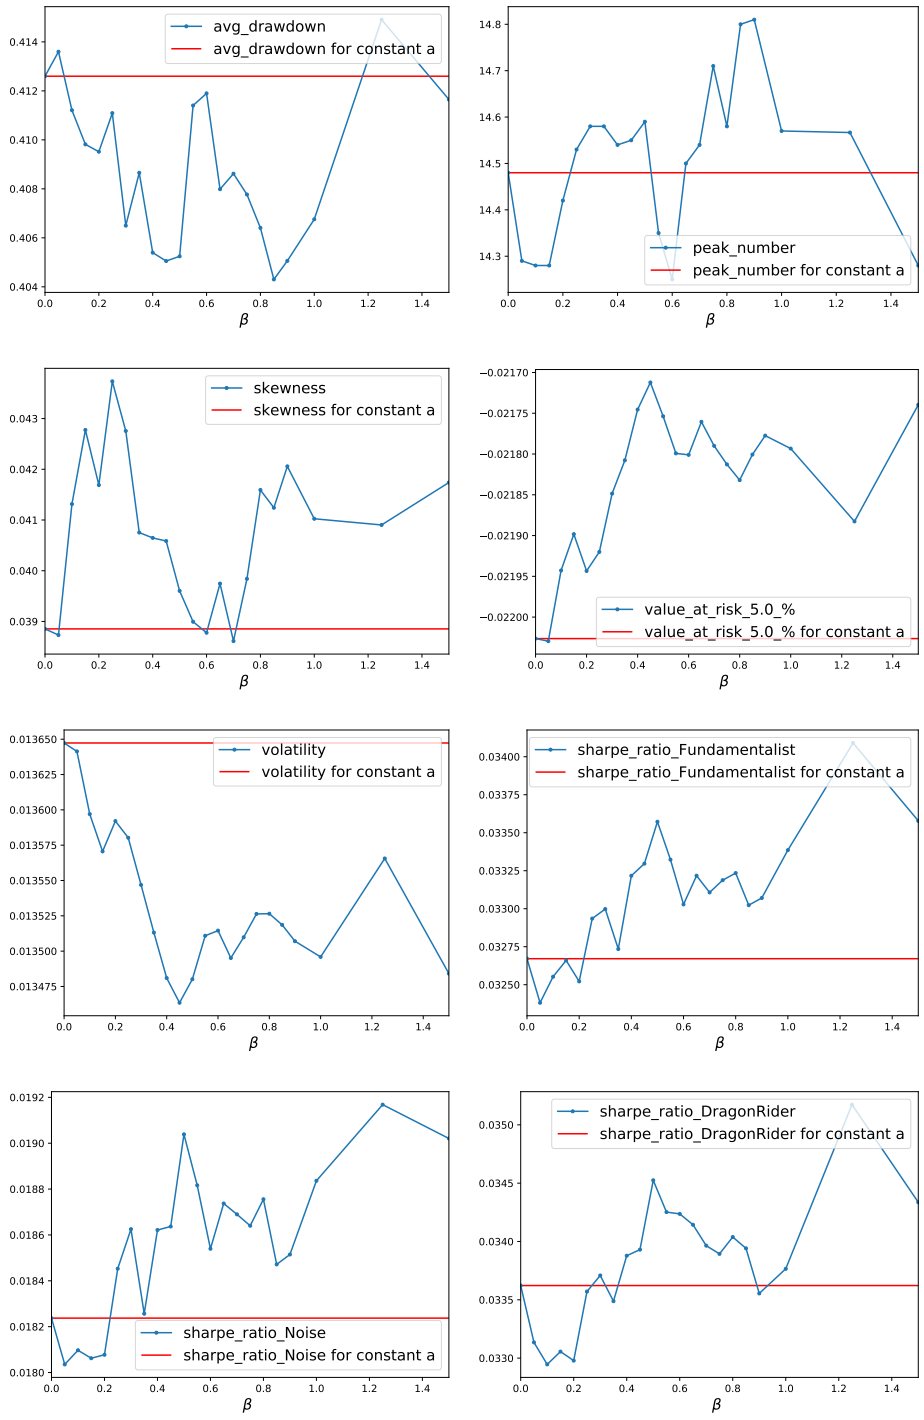

Figure 2.10: Plots of the variation of each measure average over 100 seeds, for different values of β applied to the three-agent model with the standard parameters for the Dragon Rider $a = 0.95, s = 0.0005, y2 = 0.008$. Constant κ_t is used. In blue is the variation of each parameter. In red is the benchmark of constant a behavior. From left to right, top to bottom, the measures considered are average drawdown, peak number, Value-at-Risk at the 5% level, sharpe ratio of the fundamentalist, the noise trader and the Dragon Rider.

Chapter 3

A second three-agent market model: a Dragon Slayer trader

As explained in our introduction, the goal of our study is to observe methods and consequences of market regulation. The first method is a simulation of open market operations through an agent who dedicates their wealth to making transactions to stabilize the market at the appropriate time when they anticipate the rise and crash of a bubble.
In this chapter, we present the first type of agent designed with Dragon Slayer objectives exclusively, that we will study. Recall that unlike the Dragon Rider, Dragon Slayers represent agents who use their knowledge about price bubble detection to reduce their amount and size, not arbitrage the information. In total, we study two – in the next chapter, we will study a Dragon Slayer with the ability to change the interest rate.
This is similar to what central banks would do in a context of asset price targeting. We recall that asset price targeting is not the standard targeting method for the central banks. As explained in our introduction, there is debate as to whether it is appropriate, especially in the context of reducing positive bubbles before the crash. However, we align our study with the arguments in favor of asset price targeting. One remark is that our model cannot account for inflation targeting, because of its simplicity. Central banks access to inflation data through the tracking of the price of baskets of goods. However, our market does not account for goods other than financial assets. In that context, the risk-free rate can be interpreted as a nominal interest rate, with unknown inflation and real interest rate. This does not perturb the study, since we are interested in analyzing how to prevent bubbles and crashes, making the model well adapted to our interrogations.
We also recall that central banks use several different instruments of monetary policy. The instrument we are modeling in this chapter is open market operations. We create a trader who invests its wealth on either the risky or the risk-free asset, only in order to reduce the amount and size of bubbles

in the market. In the context of conventional policy, open market operations are the purchase or sale of government bonds only, but since we apply this policy to bubbles and crashes, it is reasonable to assume that the policy enters the realm of unconventional monetary instruments. In which case, the purchase or sale of assets is also used. However, this is not unreasonable or uncommon in real life behavior. Examples ranging from the Swiss National Bank to analogous institutions in Hong Kong or Japan have shown that central banks buy large amounts of shares. This is why it makes sense to study this type of Dragon Slayer. A supporting point in this choice that we also mentioned in our introduction, is that central banks have become more and more inclined to own financial assets since the financial crisis of 2008, even when the threat of collapse is not imminent. As a last note, we also know that since 2008, the amounts of money involved have largely increased, which is why we will consider various initial wealth levels for this trader.
In the previous chapter, we performed a grid exploration to assess the impact of Dragon Riders in various parameter settings. We concluded that optimization was possible, while still using the standard parameters as a benchmark for more realistic behavior. We can now compare their performance to what we obtain here through an institution using the same tools of detection, and open market operations. We aim to reproduce the positive externality generated by the Dragon Riders.
The good performance observed in the previous chapter justifies the use of the same detection tool in our following strategy.

3.1 Description of the three-agent model

3.1.1 Market setup

The model we use is again an extension of the two-trader, two-asset model of chapter 1. The model only differs from that of chapter 2 in the third trader's characteristics, because we replace the Dragon Rider with a Dragon Slayer. As in the two-agent model, we keep one risky asset, and one risk-free asset. The risk-free asset has a constant return r_f. In addition, at each time-step the risky asset provides a dividend $d_t = d_{t-1}(1 + r_t^d)$ with $r_t^d = r_d + \sigma_d u_t$, where $u_t \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$. We keep the fundamentalist traders and noise traders identical to the previous two chapters, and introduce a Dragon Slayer type of trader.

3.1.2 Dragon Slayer strategy

Similarly to the Dragon Rider, the idea for the Dragon Slayer, is to create an extension of the fundamentalist trader, since we can dictate its behavior

through the control of the formula of estimation of the expected return of the risky asset, noted $E_{r_t}^{DS}$. It also allows us to revert to fundamentalist behavior $E_{r_t}^{DS} = r_d$ when intervention is not needed, thus replicating the intervention with minimal disturbance to the market. This is due to the fact that fundamentalists themselves are a stabilizing element of the market, since their investment is proportional to the dividend-price ratio.
Thus, we set up this trader with the following behavior characteristics. Conceptually, they behave in order to maximize their individual expected utility at each time step with the same CRRA utility as fundamentalists. The group of Dragon Slayers are summarized with a representative agent for the aggregated wealth of all the traders, just as the other types of traders. They have the freedom to invest a fraction of their allocated wealth in the risky asset, and the remainder is invested in the risk-free asset. The risky fraction is denoted x_t^{DS}. Their wealth structure is as follows:

$$W_t^{DS} = W_{t-1}^{DS}\left[1 + r_f + x_{t-1}^{DS}\left(r_t + \frac{d_t}{p_{t-1}} - r_f\right)\right]. \tag{3.1}$$

and the utility maximization implies:

$$x_{t-1}^{f} \approx \frac{E_{r_t}^{DS} + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_f}{\gamma\sigma^2}. \tag{3.2}$$

It is important to understand that while the dynamics are analogous to fundamentalists and Dragon Riders, we will set up the behavior of $E_{r_t}^{DS}$ such that this is indeed an agent of regulatory behavior. That is, their goal is not wealth maximization, but reduction of the price bubbles size and amount.
As a remark, this setup implies that the Dragon Slayer starts with an initial amount of money and invests it throughout the course of the experiment. One might ask why the Dragon Slayer is trading even when the risky asset is not showing signs of a bubble, when it is supposed to model unconventional monetary policy. We argue that in this simplified, closed model, external and discontinuous introduction of wealth and retraction of wealth in the system is not realistic. It is easier to study the effects in the context of having part of the fundamentalists turn into regulatory entities in times of bubble detection, then revert back to the fundamentalist behavior. Another argument would be that financial institutions build up their inventory of stocks during normal times, so that they can react to bubbles and sell the stock when necessary, which makes the behavior of our model analogous to real market behavior. This is enabled by having a formula for $E_{r_t}^{DS}$ that converges to r_d, the long-term average, when no bubble is detected.
The foundation for the Dragon Slayer strategy is a lean against the wind policy. The basic concept is that when a positive bubble starts to happen, the regulatory entity tries to change the trend by making the traders divest from the risky asset. That is, when Dragon Slayers anticipate a positive bubble,

they divest from the risk-free rate in order to drive the price down (partly through influence of the price momentum for noise traders, and partly by effects on the equation of supply and demand), and when it starts crashing, they slow down the fall by investing again. Dragon Riders and Dragon Slayers share the same basis for bubble detection, in the form of the overpricing

$$y_t = a y_{t-1} + (1-a)(r_t - \bar{r}), \tag{3.3}$$

where $r_t = \dfrac{p_t}{p_{t-1}} - 1$ is the risky asset return, $\bar{r}$ is the long-term return and a is the memory parameter. While the Dragon Slayer uses the same detection tool and parameters as the Dragon Rider, because of their efficiency being possibly correlated with their accuracy, it is important to note that there is a difference in strategy with the Dragon Rider situation. The Dragon Rider benefits from investing when a bubble is forming, which is why the overpricing detection tool is important, as with the Dragon Slayer. However, another key part of the strategy of the Dragon Rider is to know when to exit the risky asset position. The exit must be made before the crash happens, but possibly at the last minute. This justifies the existence of a threshold y_2 in the following equation (2.2). However, a threshold in the Dragon Slayer perspective might set the timing for entering the bubble instead of exiting it, which means that the threshold would play a very different role. In this approach, we find another way to detect the start of the crash, especially in order to control the reaction of the model right after the crash. Contrary to the Dragon Rider in chapter 2, where the crash is equivalent, in the formula, with the beginning of a negative bubble, we want to implement a policy that reacts almost immediately to the crash to counter it. There might be delays to soften the effect of the policy, but in the detection, there must be a way to pinpoint the exact moment of the crash, unlike for the Dragon Rider, when overpricing absolute value is enough to determine the strategy. Indeed, what we aim for is a lean against the wind policy where the Dragon Slayer goes from selling the asset (as the price grows too fast), to buying the asset right after the beginning of the crash.

In this perspective, we try to adapt the size of the intervention to the bubble stage. We decide to rely on the empirical derivative of the overpricing, that is the step by step acceleration of the overpricing $y_t - y_{t-1}$, because the switch in the sign of the overpricing momentum is made evident in the crossing of the value zero by the derivative. It is an approach that is in agreement with the findings of super exponential growth in the empirical price bubbles, seen in Sornette et al. (2009), Yan et al. (2012) and Jiang et al. (2010). In order not to capture too much noise, we use a smoothed derivative D_t^m calculated over m time steps in the following way:

$$D_t^m = \frac{1}{m}\sum_{k=1}^{m}(y_{t-k} - y_{t-k-1}) = \frac{1}{m}(y_{t-1} - y_{t-m-1}) \tag{3.4}$$

Figure 4.1 shows an example of the behavior of the overpricing, and the smoothed derivative of the overpricing, when confronted with a typical price path from the model of chapter 1. We use $m = 15$ in this figure. The price path is generated using the parameters of chapter 1, with Ornstein-Uhlenbeck κ_t, in order to show the behavior around price peaks. We select a zone where there are two price peaks, one at time 2500 and another shortly after, at time 2650 approximately. As we can see, while the overpricing indicates the price peaks, the value of the derivative of the overpricing is particularly apt to serve as amplitude. It is very clear for the peak at time 2500 that the sign of the derivative changes signs and crosses zero at the moment of the crash, which is helpful for a formula of the amplitude around r_d that we want to introduce. Note that the smoothing with $m = 15$ induces a delay of a few days in the response of the derivative. This causes a delay in the response of the Dragon Slayer. However, unnecessary noise caused by too low m can also cause malfunction of this regulatory force. Therefore, we will calibrate this parameter. We can also notice that the amplitude of the derivative is not symmetrical for this peak, as the growth of the returns grows slower than it falls after the crash. This leads the correction before the crash to be softer than the counteraction right after it. In this case, the peak is not too sharp since the returns actually slow down for a while before the crash. This highlights a response of the derivative that is adaptable to the characteristics of the peak.

Note that in addition to this derivative, we could use the same indicator reacting to the spread $|y_t| - y_2$ as for the Dragon Rider.

Indeed, we could consider intervening only after the overpricing has reached a certain threshold and not before, that is, only changing anything in the case of a bubble of a certain size. However, it is not the same situation. t is important to note that the threshold y_2 plays a different role in the Dragon Rider's strategy. Indeed, intuitively, the Dragon Rider benefits from only exiting the bubble at the last possible moment before the crash. However, the role of the threshold y_2 for the Dragon Slayer is to indicate the beginning of a bubble, for the Dragon Slayer to sell the risky asset. Conceptually, if the amplitude is based on the derivative of the overpricing, there is also no need to separate the intensity of response after reaching some arbitrary threshold y_2. Indeed, the derivative of the overpricing increases more and more rapidly before the crash. This is due to the super exponential nature of the growth. These observations lead us to disregard the idea of differentiating the response using some threshold, and to focus on a design that has the same response to the derivative regardless of the status of the overpricing itself, i.e. having the same response before or after the threshold is reached, with an amplitude based on the derivative of the logistic function.
Thus, we design a response such that regardless of the bubble diagnostic,

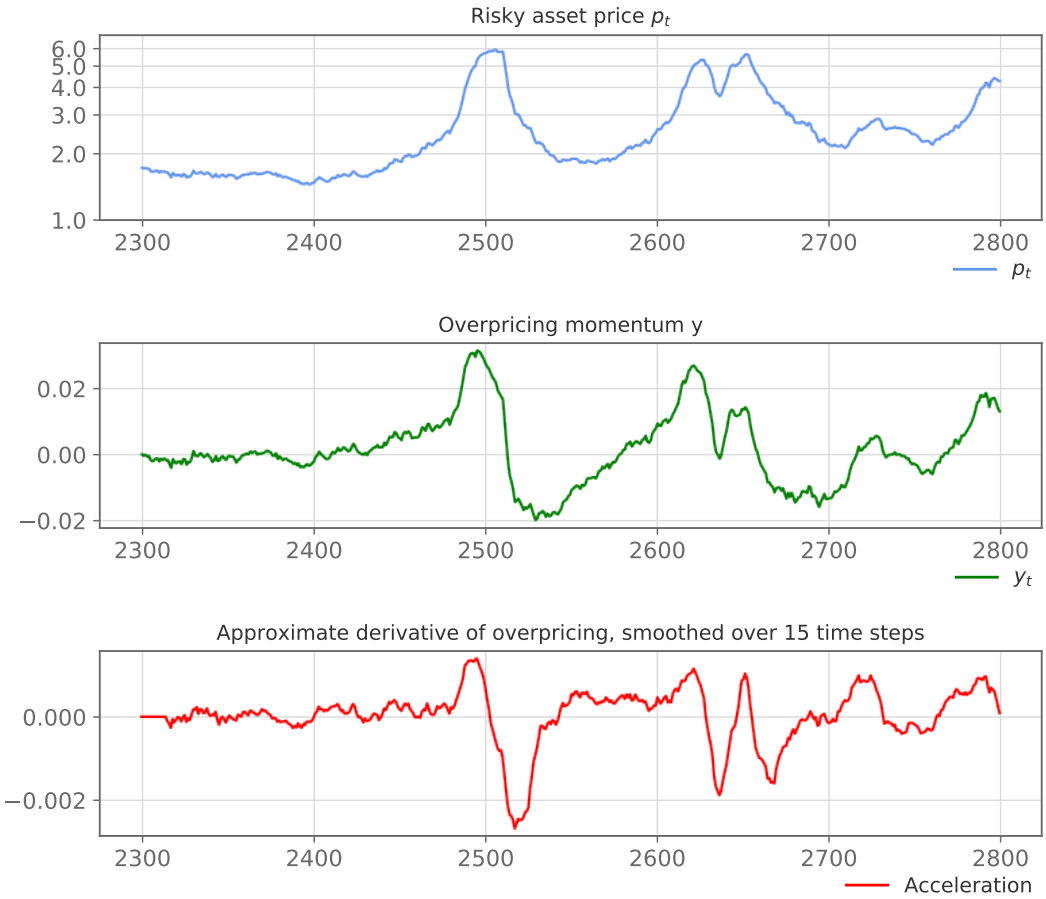

Figure 3.1: Example plots, over 500 time steps, of a price path accompanied with the corresponding overpricing momentum and overpricing acceleration (top, middle and bottom panels).

the formula of the expected return is modified with the same amplitude a_t, in which case the logistic function becomes redundant, as in equation (3.5)

$$E^{DS}_{r_{t+1}} = \bar{r} - a_t, \tag{3.5}$$

Considering the behavior of the smoothed derivative of the overpricing D^m_t, we decide to set $a_t = wD^m_t$, where w is a weight that we will calibrate considering the typical range of D^m_t, and the response of the model using the formula without threshold first, as in (3.5). We will also consider the effect of the window of smoothing m on the results, again calibrating on the model using the formula without threshold.
So, to summarize, we will study the following formula for the expected

return used by the Dragon Slayer to choose its risky fraction:

$$E^{DS}_{r_{t+1}} = \bar{r} - wD^m_t, \tag{3.6}$$

3.1.3 Price derivation

Again, the market price follows the equilibrium of the supply and demand. Similar manipulations of this equation as in 1.4.2 in chapter 1 lead to an analogous quadratic equation of the price.

$$\begin{aligned}
0 = \;& \Delta D^f_t + \Delta D^n_t + \Delta D^{DS}_t \\
= \;& W^f_{t-1}\left(x^f_t\left[1 + r_f + x^f_{t-1}\left(r_t + \frac{d_t}{p_{t-1}} - r_f\right)\right] - x^f_{t-1}\frac{p_t}{p_{t-1}}\right) \\
& + W^n_{t-1}\left(x^n_t\left[1 + r_f + x^n_{t-1}\left(r_t + \frac{d_t}{p_{t-1}} - r_f\right)\right] - x^n_{t-1}\frac{p_t}{p_{t-1}}\right) \\
& + W^{DS}_{t-1}\left(x^{DS}_t\left[1 + r_f + x^{DS}_{t-1}\left(r_t + \frac{d_t}{p_{t-1}} - r_f\right)\right] - x^{DS}_{t-1}\frac{p_t}{p_{t-1}}\right)
\end{aligned}$$

3.2 Dragon Slayer results

3.2.1 Choosing the weight w

In this section, we will choose the weight w associated with the smoothed derivative of the overpricing in the perturbation factor of the formula (3.7)

$$E^{DS}_{r_{t+1}} = \bar{r} - wD^m_t, \tag{3.7}$$

We base the level of the amplitude on the typical range of D^m_t. Typical values of D^m_t in case of a price peak have a maximum absolute value around 0.0025, as in figure 4.1, as evidenced by experiments on a 100 seeds. We recall that the baseline value of $E^{DS}_{r_{t+1}}$ is $\bar{r} = r_d = 0.00016$, and that Dragon Rider variations of $E^{DR}_{r_{t+1}}$ typically had an absolute value of a maximum 0.0005, to keep the risky fraction in an appropriate range $[0, 1]$ without too much concentration at the extremes, or in the middle. Based on this, we choose a baseline of w around 0.02 to keep this amplitude. Sensitivity analysis shows it to be an appropriate choice, when testing later (including at different wealth levels), in the wider range of $w \in \{0.0001, 0.0005, 0.01, 0.02, 0.05, 0.1\}$. The value 0.02 was found to have the best behavior for the model in terms of peak number times average drawdown.

3.2.2 Wealth variation

Given this weight value, we study the effect of the variation of wealth on the model. We find that the reduction of size and amount of bubbles gets better as the wealth ratio increases. Recall that the initial wealth ratio is

always referring to the ratio of Dragon Riders or Slayers wealth, to the sum of Dragon Riders or Slayers and fundamentalists wealth. As discussed in 2 in the case of the Dragon Rider, we keep the total wealth of fundamentalists and Dragon Slayers combined equal to 10^9, the same wealth as the noise traders, to allow for comparison. As we can see in figure 3.2, the results on peak number are a bit inconclusive because the number of peaks is relatively similar, between 13.5 and 14.2. However, driven by the significant decline in the average drawdown from 0.7 to 0.45, the product of the two steadily decreases as the wealth ratio of Slayers increases.

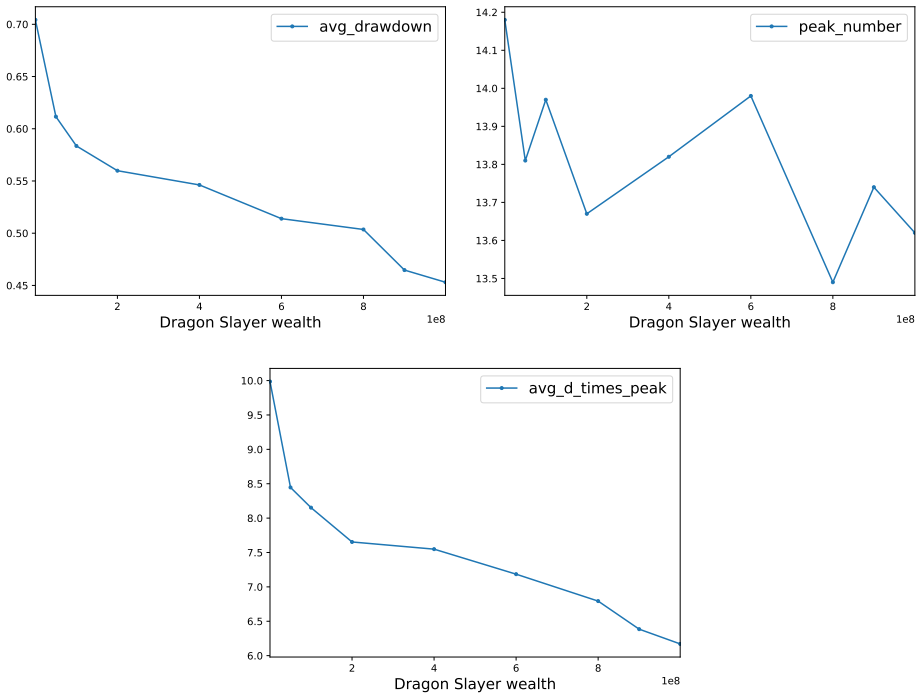

Figure 3.2: Plots of the variation of each measure average over 100 seeds, for different values of the initial wealth of the Dragon Slayer applied to the three-agent model with memory parameter for the Dragon Slayer $a = 0.95$. Total initial wealth of fundamentalists and Dragon Slayers is maintained at 10^9. Ornstein-Uhlenbeck κ_t is used. From left to right, top to bottom, the measures considered are average drawdown, peak number, and product of average drawdown and peak number

This is confirmed in the behavior of the price paths in the representative seeds, for time steps $t = 2000$ to 4000, using the parameters of chapter 1 as a foundation, $w = 0.02$, and $a = 0.95$ for the Dragon Slayer when there is one.

In figure 3.3, we show a benchmark for the price path in the case of only two traders (no Dragon Slayer). In figure 3.4, we show the same section in the case of 20% wealth ratio for the Dragon Slayer, and in 3.5 and 3.6, we show the same section for initial wealth percentages 40% and 90% respectively.
The reduction in size of the bubble at around $t = 3250$ is clear (note that the scale changes at every graph). We also notice that with increased amounts of Dragon Slayers in the 90% frame, the bubble becomes a negative bubble. The response becomes too strong, in the case of this price bubble. However overall, 90% Dragon Slayer wealth performs better than lower proportions in reducing bubbles.

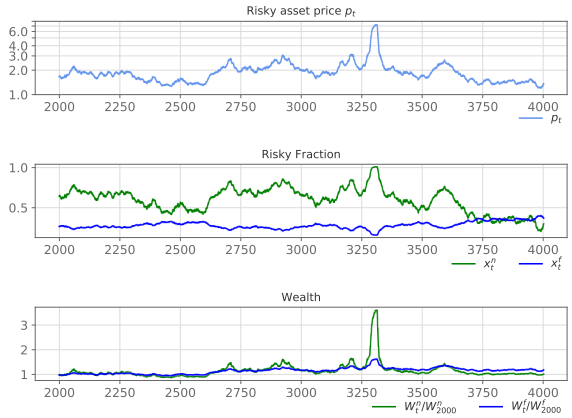

Figure 3.3: Plot of one realization (seed 126963) of the two-agent model, using parameters from chapter 1. Top panel shows the price path, second panel shows the risky fractions of the traders, lower panel shows the wealth of the traders. The section shown is from $t = 2000$ to $t = 4000$.

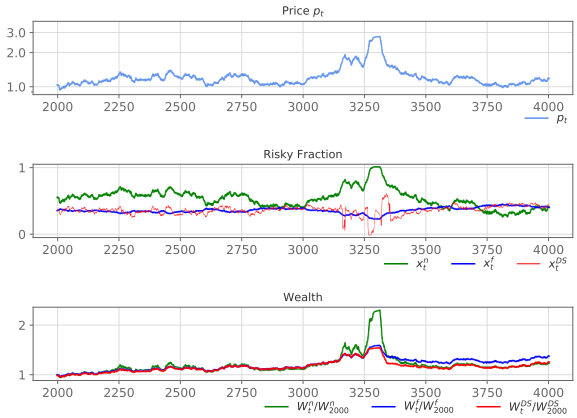

Figure 3.4: Plot of one realization (seed 126963) of the three-agent model with Dragon Slayer, using parameters from chapter 1, $w = 0.02$ and initial wealth of the Dragon Slayer at 20% of the total of Slayers and fundamentalists. Same time window and panel organization as on the left.

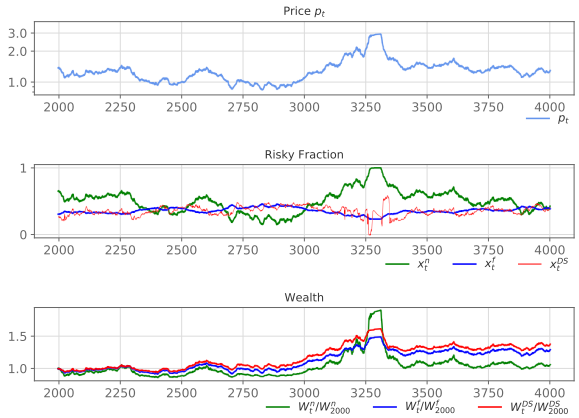

Figure 3.5: Plot of one realization (seed 126963) of the three-agent model with Dragon Slayer, using parameters from chapter 1, $w = 0.02$ and initial wealth of the Dragon Slayer at 40% of the total of Slayers and fundamentalists. Same time window and panel organization as above.

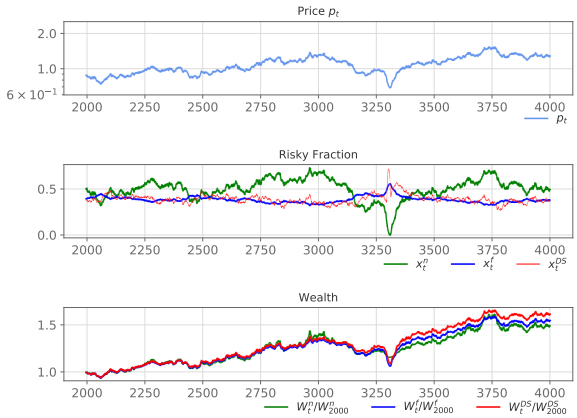

Figure 3.6: Plot of one realization (seed 126963) of the three-agent model with Dragon Slayer, using parameters from chapter 1, $w = 0.02$ and initial wealth of the Dragon Slayer at 90% of the total of Slayers and fundamentalists.Same time window and panel organization as above.

Despite the fact that the optimal wealth allocation is to replace all fundamentalists with Dragon Slayers, this is not realistic to the composition of a real market. We will keep a 50% allocation, which already allows for reduction of bubbles, in the following examples.

3.2.3 Optimization of smoothing window m

We now derive a sensitivity analysis to get the best values for m in terms of bubble size and amount reduction. Intuitively, very low m corresponds to no smoothing of the derivative, which could induce noise, while high m induces a delay in the reaction to the peak.
In figure 3.7, we show the effect of the smoothing window size on the peak number, average drawdown and product of the two. The results are a bit unclear but it seems like values between 10 and 25 provide the lowest values in size times amount of peaks. In particular, there seems to be a low amount of peaks with higher amplitude at around 15 time steps of smoothing. These results should be taken with a grain of salt since the amplitude of the changes is never very high here. The peak number varies from 13.6 to 14.2, while the average drawdown takes values from 0.053 to 0.055. We set $m = 10$ for the rest of the experiments.

3.2.4 Sharpe ratio distribution

As we can see in figure 3.8, the Sharpe ratio distribution is similar to that of the two-agent model in figure 1.3 for the fundamentalist and noise traders, and the risky asset. The same remarks apply for these three strategies with respect to the change of behavior between constant κ and Ornstein-Uhlenbeck κ_t. This means that the bubble reduction is not generally improving the efficiency of the strategies.
An interesting point is the distribution of the Dragon Slayer Sharpe ratios. It is noteworthy that in constant κ, with no bubbles, the Dragon Slayer behavior collapses with the behavior of the fundamentalist. This is the behavior that we wanted for this trader. However, in the Ornstein-Uhlenbeck case, the Dragon Slayer has a lower average Sharpe ratio than the fundamentalist does. This is to be expected, since the strategy of the Dragon Slayer is to go against the market during a bubble.

3.2.5 Comparison with Dragon Rider

In table 3.1, we use the combination of parameters $a = 0.95, s = 0.0005, y_2 = 0.008$ for the Dragon Rider to compare the effects of Dragon Riders with those of Dragon Slayers with $a = 0.95$ and $w = 0.02$. As is evidenced by the amounts, the Dragon Rider is much more effective at reducing size and especially amounts of bubbles than the Dragon Slayer. However, compared

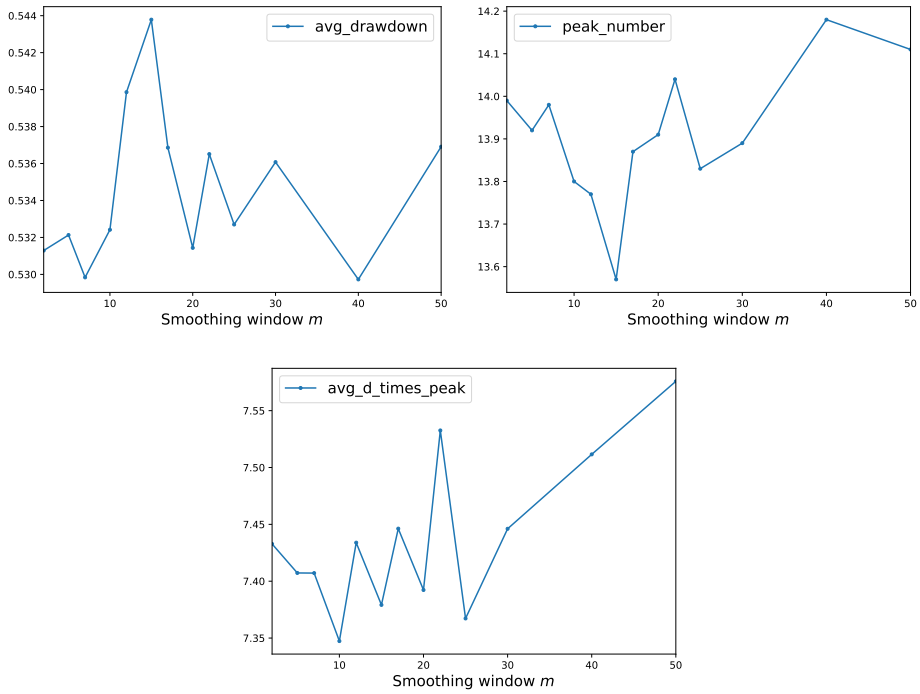

Figure 3.7: Plots of the variation of each measure average over 100 seeds, for different values of the smoothing window of the derivative m applied to the three-agent model with memory parameter for the Dragon Slayer $a = 0.95$. Ornstein-Uhlenbeck κ_t is used. From left to right, top to bottom, the measures considered are average drawdown, peak number, and product of average drawdown and peak number

to no intervention at all, the Dragon Slayer manages to reduce the number of bubbles by almost two thirds.

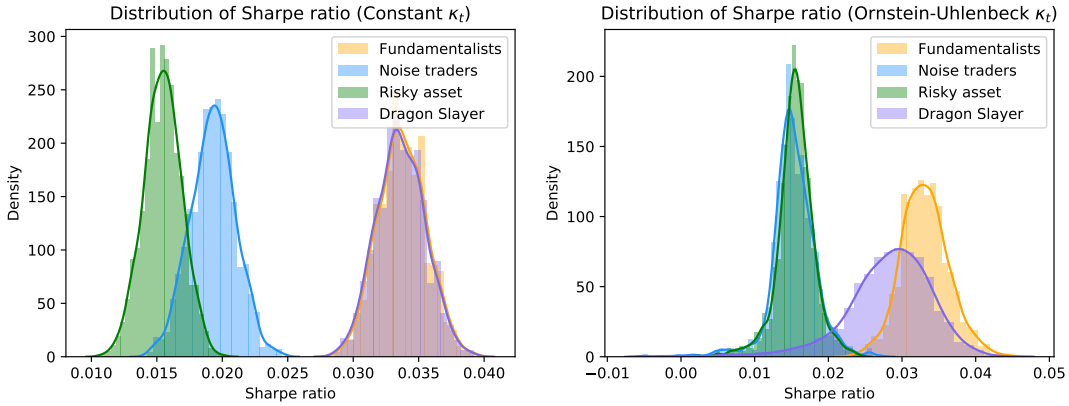

Figure 3.8: Distribution over 1000 seeds of the Sharpe ratios of the Fundamentalists, Noise traders, Dragon Slayers and pure risky asset strategy, for a three-agent model with parameters from chapter 1, $a = 95$, $w = 0.02$ and a wealth level of 90%. Sharpe ratios are calculated between $t = 5000$ and $t = 17500$.

%	Model	Peak nb	Avg dd	Avg dd $\times$ p nb
0	2 agent	19.9	0.58	11.5
20	DS	13.7	0.56	7.7
20	DR	9.47	0.28	2.65
100 (opt)	DS	13.6	0.45	6.12
60 (opt)	DR	1.26	0.46	0.58

Table 3.1: Comparison between Dragon Rider and Dragon Slayer of average value over 100 seeds of average drawdown, peak number, and the product, for 20% of initial wealth, and the respective optimum initial wealth. The measures of the two-agent model are at the top row as a benchmark. Measures are taken over a time period $T = 10000$.

Chapter 4

Regulation through interest rate modification

In this chapter, we extend the model from chapter 1 in a different direction from chapters 2 and 3. Instead of considering traders and their positive externality on the market, or open market operations as a means for regulatory entities to reduce the size and amount of bubbles directly, we design regulatory strategies through interest rate alteration. As explained in our introduction, the regulation of markets can take several aspects. One of them, used by central banks, is the modification of the interest rate, through various mechanisms such as setting the overnight rates or influencing the exchange rates. Our model is very simple, so we will summarize these interventions through a modification of the nominal interest rate directly, i.e. modification of the interest rate r_f previously considered constant. The concept of the regulatory agent is still the "Dragon Slayer" type. To reduce confusion in the naming, we will refer to this agent as the "interest rate Dragon Slayer", and the agent from chapter 3 as the "Dragon Slayer trader" when the context is not clear. In the absence of this clarification, by default, "Dragon Slayer" refers to the interest rate Dragon Slayer throughout the chapter.
Note that in essence, we will also adopt lean against the wind policy when it comes to interest rate modification. This is similar to the approach in chapter 3. However, the influence on the interest rate also influences the wealth of the traders

4.1 Market setup

We again use the market setup of the two-agent model. The difference here is that we introduce a new agent, a Dragon Slayer with the ability to modify the interest rate.

4.1.1 Basic market setup

Assets As in the two-agent model, we keep the structure of one risky asset, and one so-called risk-free asset. However, the latter is no longer fully risk-free, since now, its return r_t^f, can be modified by the Dragon Slayer at each step. The asset still represents a bank account or government-backed bond. This represents the nominal interest rate. The price dynamics of the risky asset are determined by the equilibrium of supply and demand and recalled in section 4.1.2. The risky asset yields the same dividend d_t at each time step, as in the previous chapters, such that $d_t = d_{t-1}(1 + r_t^d)$ with $r_t^d = r_d + \sigma_d u_t$, where $u_t \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$.

Fundamentalist traders and noise traders We keep a representative fundamentalist trader with the following wealth allocation principle:

$$x_{t-1}^f \approx \frac{E_{r_t} + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_t^f}{\gamma\sigma^2}. \tag{4.1}$$

And the representative noise trader keeps the same wealth allocation scheme as well:

$$x_t^n = \frac{1}{N_t^+ + N_t^-} \sum_{i=1}^{N_{t-1}^+} (1 - \xi_i(p_{t-1}^+)) + \frac{1}{N_t^+ + N_t^-} \sum_{j=1}^{N_{t-1}^-} (1 - \xi_j(p_{t-1}^-)), \tag{4.2}$$

where $\xi(p)$ are Bernoulli random numbers with parameter p.
Wealth dynamics for noise traders and fundamentalist traders remain:

$$W_t = W_{t-1}\left[1 + r_t^f + x_{t-1}\left(r_t + \frac{d_t}{p_{t-1}} - r_t^f\right)\right] \tag{4.3}$$

Bubble detection tool We keep the overpricing definition:

$$y_t = a y_{t-1} + (1-a)(r_t - \bar{R}), \tag{4.4}$$

where $r_t = \frac{p_t}{p_{t-1}} - 1$ is the risky asset return, $\bar{R}$ is the long-term return and a is the memory parameter.

4.1.2 Dragon Slayer scheme

Basic concept

The initial concept that we use is to find a way to modify the interest rate around a constant long-term average r_f in such a way that market crashes are dampened. This involves slowing down price growth when it reaches abnormally high levels, and slowing the falls. This corresponds to increasing

the interest rate when the price is about to crash (in order to encourage the investors to converge towards the risk-free asset), and decreasing it when the price is currently crashing in order to make the risky asset comparatively more attractive. The average interest rate is maintained at a constant level r_f. We use the evolution of the overpricing (4.4) as a guide for the variation around r_f. When a positive bubble is expected, the Dragon Slayer should be attempting to slow down the growth of the bubble, by increasing the interest rate, therefore $r_t^f > r_f$. When the crash starts, the Dragon Slayer should be trying to slow down the fall of the price by making the risk-free asset less attractive, i.e. decreasing the interest rate, therefore $r_t^f < r_f$. Since the overpricing quantifies the growth of the returns with the long-term average as a reference point, we use its evolution as in chapter 3 to calibrate the amplitude a_t in the following:

$$r_t^f = r_f + a_t, \tag{4.5}$$

where a_t is the amplitude of the correction prescribed by the Dragon Slayer. Note that in terms of notation and timing, the Dragon Slayer chooses the interest rate r_t^f, that will apply to the risk-free asset from t to $t+1$, at $t-1$ and given all available information at $t-1$. We can also write this formula in the following form (where a_t is now a multiplying factor for r_f):

$$r_t^f = r_f(1 + a_t). \tag{4.6}$$

with a_t is the amplitude of the relative correction prescribed by the Dragon Slayer. For the same reasons as above, a_t should be positive during the rise of the bubble, and negative during the crash. In practice, we use second formula (4.6) for better understanding of the chosen amplitude.

Amplitude of the correction

We attempted to assess the effect of interest rate r_t^f modifications on the price at t. The larger ambition here was to obtain a formula for a_t from 4.6 as a function of the state of the market at $t-1$ by examining the variation of price at t with respect to the interest rate modifications at $t-1$. We used Taylor expansion on the price equation to achieve this goal in first order approximation. As explained above, the interest rate r_t^f, applying to the risk-free asset from t to $t+1$, is chosen at $t-1$.
In practice, we calculated the impact of a variation dr_t^f on the price p_t. Recalling the price equations in the two-agent model from chapter 1, we performed a Taylor expansion which was inconclusive at this level. We refer to the appendix for the details of the Taylor expansion.
In the end, we decided to drop this approach to the determination of the amplitude. Indeed, there are two main caveats to this line of reasoning. Firstly,

a lot of the state at $t-1$ cannot be perfectly known by the Dragon Slayer. It is not reasonable to assume that the Dragon Slayer knows the risky fractions of the other traders, for example, because in reality the risk aversion of the population is not clear. Whether the Dragon Slayer could access the dynamics and proportion of noise traders and fundamentalists is also debatable, so we can only settle on averages for most values. Secondly, the approximation we make in the appendix does not account for the full effect of the noise traders on the price. We treat the risky fraction of noise traders at time t as a constant independent of the price, to simplify the calculations, but in reality the herding effect amplifies the effect of the variations of the interest rate on the price. Therefore, the estimation can only be taken as a lower boundary of the real price impact of a change in interest rate.
And finally, this approach to the determination of the amplitude is very specific to our model. We aim for a general formula that would possibly apply to extensions of this model. This is why we opted for a formulation based on empirical behavior rather than theoretical manipulation of the price equation.

Formula for the interest rate

Following the discussion in section 4.1.2, we find decide on a policy independently from the price equation. Conceptually, we again need a lean against the wind policy for the interest rate Dragon Slayer. We decide, based on the behavior of the Dragon Slayer trader model in chapter 3, to use the same overprice derivative based amplitude function. In order not to capture too much noise, we use a smoothed derivative D_t^m calculated over m time steps in the following way:

$$D_t^m = \frac{1}{m}\sum_{k=1}^{m}(y_{t-k} - y_{t-k-1}) \tag{4.7}$$

We recall in figure 4.1 a comparative view of the overpricing and smoothed derivative of the overpricing, on the price path of the model of chapter 1 containing a peak. We use $m = 15$ in this figure. The price path is generated using the parameters of chapter 1, with Ornstein-Uhlenbeck κ_t, in order to show the behavior around price peaks. The sign of the derivative changes signs and crosses zero $\frac{m}{2}$ steps after the maximum value of y_t, which is close to the behavior we want for a_t We choose to use the following formulation for the equation:

$$r_t^f = r_f(1 + wD_t^m). \tag{4.8}$$

Note that this is very close to the model in chapter 3. In particular, under certain weight w values, the formulas of the risky fraction of the fundamentalists in this model could be affected in a similar way as the formulas of the risky fraction of the Dragon Slayer traders in chapter 3. Indeed, the risky

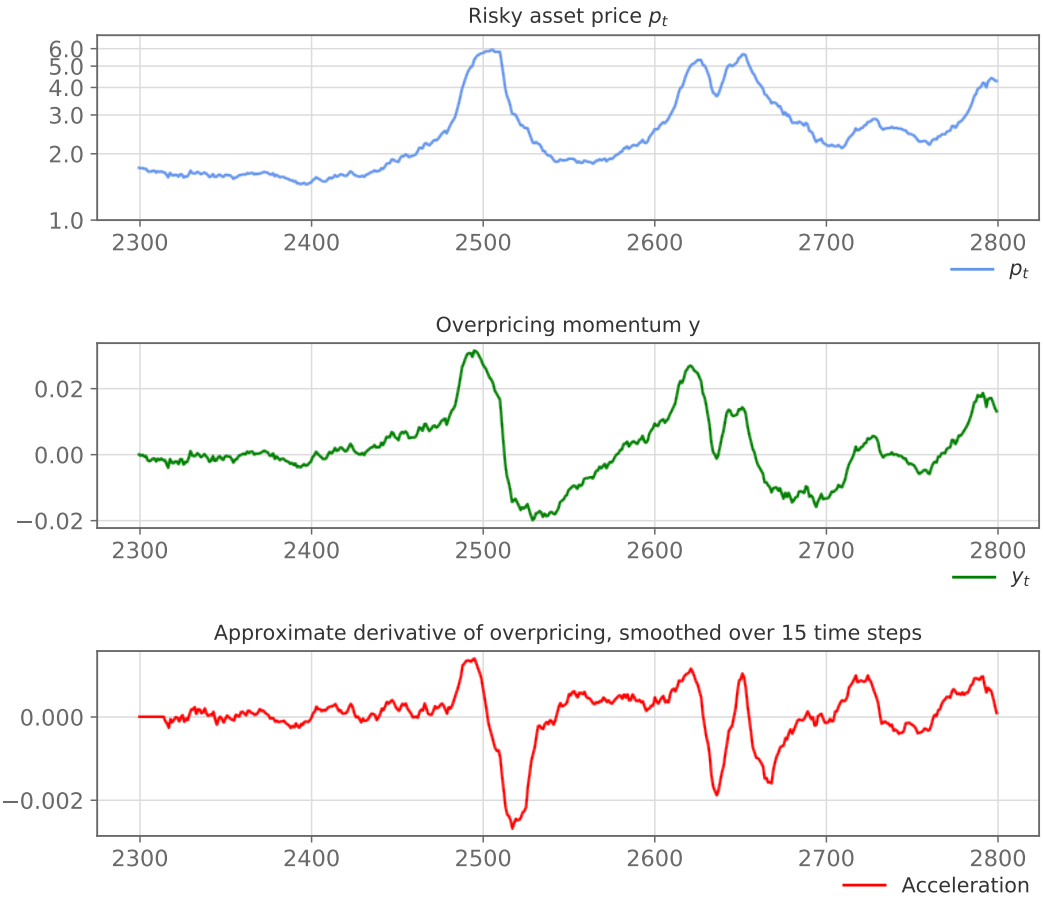

Figure 4.1: Behavior of the derivative of the overpricing, plotted over 500 trading days, with the parameters of chapter 1. We use Ornstein-Uhlenbeck κ. Top panel is the price, middle panel is the overpricing, and the third panel is the empirical derivative, smoothed over 15 time steps.

fraction of the fundamentalists here is defined as :

$$\begin{aligned} x^f_{t-1} &\approx \frac{E_{r_t} + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r^f_t}{\gamma\sigma^2} \\ &= \frac{r_d + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_f(1+wD^m_t)}{\gamma\sigma^2} \end{aligned} \tag{4.9}$$

while the risky fraction of the Dragon Slayer traders from chapter 3 is defined as:

$$x_{t-1}^f \approx \frac{E_{r_t}^{DS} + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_f}{\gamma\sigma^2} = \frac{r_d - wD_t^m + \frac{d_{t-1}(1+r_d)}{p_{t-1}} - r_f}{\gamma\sigma^2} \tag{4.10}$$

which corresponds to the same effect, with a different weight setting. However, the difference lies in the fact that the modification of r_f^t also induces a difference in wealth for the traders, thus changing the effect that they have in the market. For instance, increasing r_f^t for a couple of time steps not only makes fundamentalists switch part of their investments to the so-called risk-free asset, but they also increase their wealth in doing so, as opposed to the Dragon Slayer traders who were losing wealth to influence the market, as we saw looking at their Sharpe ratios. This means that the approach of the fundamentalists in this chapter is still that of wealth maximizers. This possibly gives modifications in the derivative of the overpricing momentum a bigger effect on the price, even with equal contribution in the risky fraction. Indeed, as the fundamentalists gain wealth, their influence on the price is heightened.
Keeping these differences in mind, we now study the effect of parameters a and w on the performance of the intervention of the Dragon Slayer.

4.2 Dragon Slayer results

We initially performed a grid search over the values of the weight w and the memory parameter a in order to get an idea of the best parameters for reducing the amount and size of bubbles. The effect is studied with Ornstein-Uhlenbeck κ, to test the influence on the size and amount of bubbles. The smoothing of the derivative is kept at 15. We study combinations of parameters in the following ranges:

$$a \in \{0.7, 0.85, 0.95, 0.98, 0.992, 0.996, 0.998\}$$
$$w \in \{1, 5, 10, 100, 200, 500, 1000\}$$

Note that the range for w is purposefully large, because the effect of the interest rate on the behavior of the price is not well quantified initially. However, note that in order to keep the same order of magnitude as the Dragon Slayer traders of chapter 3, we need a weight level w of the same order as $0.02/r_f$, i.e. of the order of magnitude of 500. Note that the weight could possibly need to be lower, because of the additional effects linked to the wealth modification of the fundamentalist and noise traders, that are not present in

chapter 3. The range for a corresponds to a memory going from 3 days to 500.
We find initially that the best combination, in terms of peak number and average drawdown as well as qualitative behavior for individual seeds, is around $a = 0.996$, $w = 200$. Note that the weight is around what we expect, albeit smaller, which could be explained by the additional effects of the interest rate modification on the wealth of the traders. We showcase the behavior of this combination, and refine the study through sensitivity analysis in the following.

4.2.1 Effects of varying a

We compare the behavior of the model for $w = 200$ and values of a in $[0.992, 0, 998]$. We note an optimal value at around 0.996 for average drawdown times peak number. Note that this optimum is mostly driven by the variations in the average drawdown, since the peak number seems to be slightly higher in the neighborhood of 0.996. However, as we already remarked, the peak number could go slightly up as bubbles become smaller, since peaks are not defined by an absolute size but rather comparatively to the rest of the price path. Which is why we consider the product of the two measures to be the most relevant information here. Another remark is that the performance seems to have rather high sensitivity around 0.996, even if the general trend indicates an optimum.

4.2.2 Effects of varying w

Now, we evaluate the behavior of the model for w in $[100, 400]$, for $a = 0.996$. In figure 4.3, we show the effect of the weight on the measures of average drawdown, peak number and the product of the two. We conclude that 200 is an optimal value for w. The optimum is quite clear here, and mostly driven by average drawdown too.

In figures 4.4 and 4.5, we showcase the behavior on an example experiment, for Ornstein-Uhlenbeck κ, with no intervention and with intervention on the interest rate with parameters $a = 0.996$ and $w = 0.1$. As we can see, the positive bubbles are hindered by the response of the Dragon Slayer. Note that some other bubbles are not always hindered, and there is sometimes overreaction on negative bubbles, which explains the remaining bubbles even with Dragon Slayer intervention.

4.2.3 Effects of varying m

Finally, we vary the smoothing window to assess the sensitivity of the model to this parameter. We vary the smoothing parameter within $\{2, 5, 7, 10, 12, 15, 20, 25, 30\}$. The values either 7 or 15 seem optimal. The value 7 is optimal in terms of

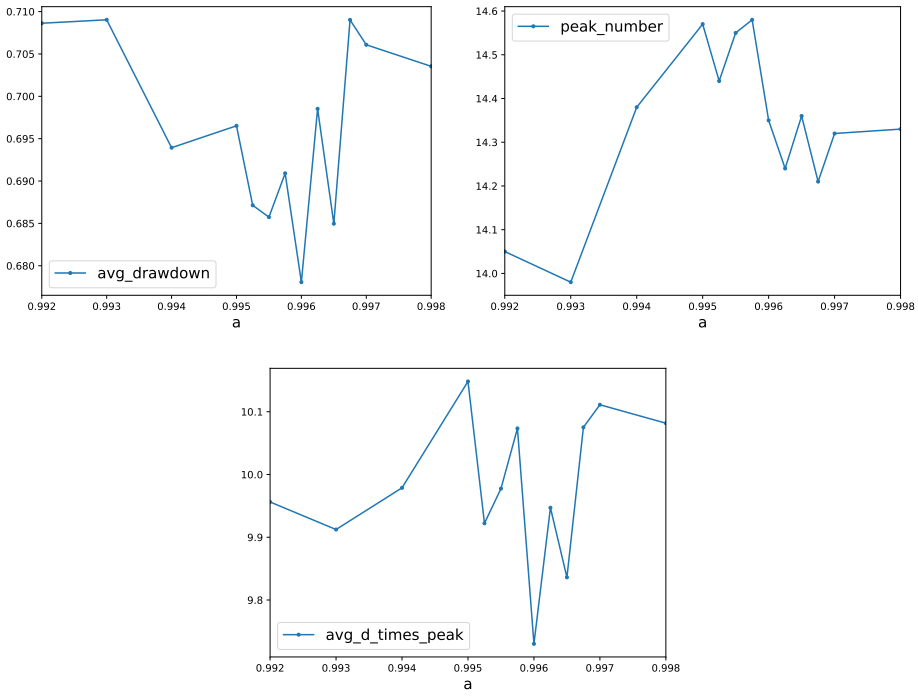

Figure 4.2: Plots of the variation of each measure average over 100 seeds, for different values of the memory parameter a, applied to the three-agent model with weight $w = 200$. Ornstein-Uhlenbeck κ_t is used. From left to right, top to bottom, the measures considered are average drawdown, peak number, and product of average drawdown and peak number

peak number, while the value 15 is optimal in terms of average drawdown. We keep 15 as a smoothing parameter for the summary performance, because it is deemed preferable to have a larger amount of smaller bubbles. Furthermore, the product of the average drawdown and peak number is also slightly smaller for $m = 15$.

4.2.4 Summary results

We compare the results of the different models in table 4.1. We give a reference performance of the model with no intervention, as well as the results of the model with Dragon Riders, Dragon Slayer traders, and interest rate Dragon Slayers. Note that the model from this chapter is not the worst performing of the three, but it has worse average drawdown than the model without intervention. Comparatively with the Dragon Slayer trader model, it seems that the effect on bubbles is less strong. We recall that the interest rate Dragon Slayers have potentially the same effects on the risky fractions of the fundamentalists as the modification applied on the fraction of Dragon

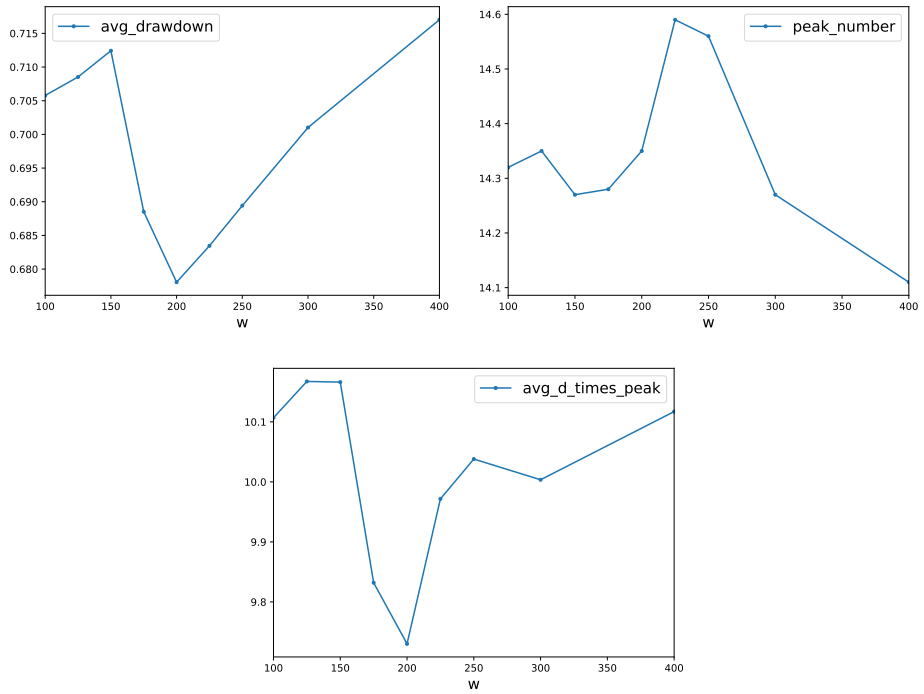

Figure 4.3: Plots of the variation of each measure average over 100 seeds, for different values of the weight w, applied to the three-agent model with memory paramete $a = 0.996$. Ornstein-Uhlenbeck κ_t is used. From left to right, top to bottom, the measures considered are average drawdown, peak number, and product of average drawdown and peak number

Slayer traders in chapter 3. However, the interest rate modifications in the present chapter concern the entirety of the fundamentalists, when Dragon Slayer traders only ever represent a fraction of the total. Furthermore, there is also an effect on the wealth of the noise traders and the fundamentalists, when the interest rate is modified, that is not present in the case of the Dragon Slayer traders. Dragon Slayer traders indeed lose wealth, as mentioned in chapter 3, when they counter the bubbles in lean against the wind policy and divest from a sharply rising in price risky asset. Conversely, the fundamentalists of the present chapter are fueled by utility maximization when they choose to divest from the risky asset and move their wealth to the so-called risk-free asset. Such effects can explain the discrepancy between the two strategies results. Note that the performance of the model of the present chapter is also thwarted by the overreaction to some negative price bubbles.

Figure 4.4: Plot of one realization (seed 314382) of the model without Dragon Slayer intervention, using parameters from chapter 1. Top panel shows the price path, second panel shows the risky fractions of the traders, third panel shows the wealth of the traders, bottom panel shows the constant risk-free rate . The section shown is from $t = 2000$ to $t = 4000$.

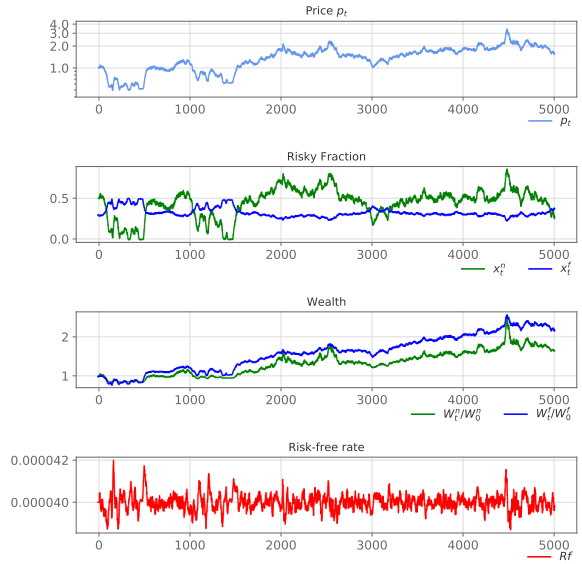

Figure 4.5: Plot of one realization (seed 314382) of the three-agent model with Dragon Slayer, using parameters from chapter 1, $w = 0.1$ and $a = 0.996$. Top panel shows the price path, second panel shows the risky fractions of the traders, third panel shows the wealth of the traders, bottom panel shows the variations of the risk-free rate.

% wealth	Model	Peak nb	Avg d	Avg dd × p nb
0	2 agent	19.9	0.58	11.54
100 (opt)	DS chapter 3	13.6	0.45	6.12
60 (opt)	DR	1.26	0.46	0.58
	DS chapter 4	14.35	0.67	9.6

Table 4.1: Comparison between optimal Dragon Rider and optimal Dragon Slayers of chapters 3 and 4 of average value over 100 seeds of average drawdown, peak number, and the product. The measures of the two-agent model are at the top row as a benchmark. Measures are taken over a time period $T = 10000$.

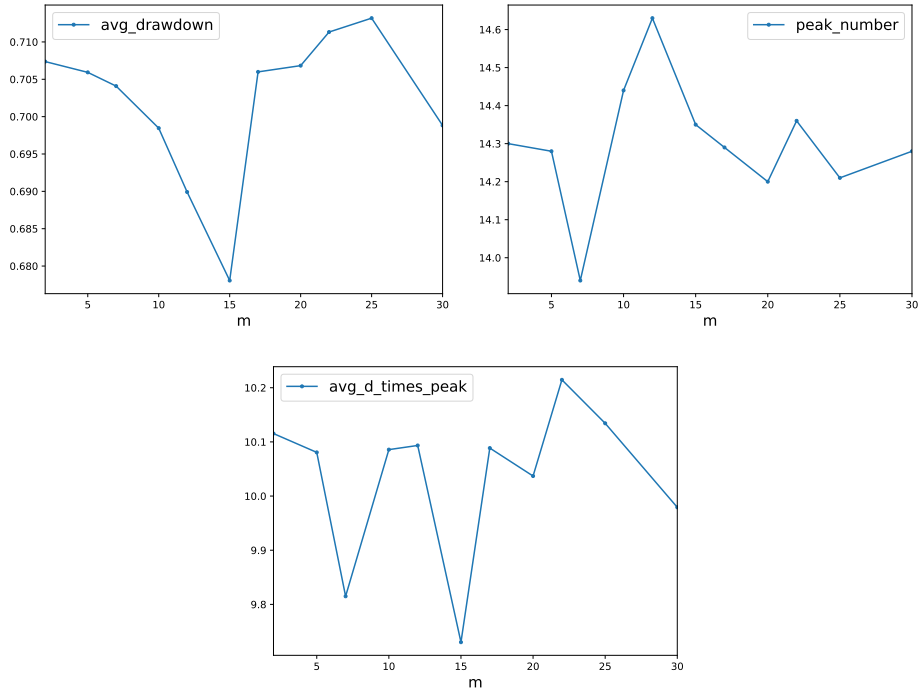

Figure 4.6: Plots of the variation of each measure average over 100 seeds, for different values of the smoothing window m of the derivative, applied to the three-agent model with memory parameter $a = 0.996$ and $w = 200$. Ornstein-Uhlenbeck κ_t is used. From left to right, top to bottom, the measures considered are average drawdown, peak number, and product of average drawdown and peak number

Conclusion

The main objective of this thesis was to design and study the effect of monetary policies aimed at reducing the frequency and size of price bubbles on the market. In order to do so, we extended an agent-based model developed in Kaizoji et al. (2015), Westphal and Sornette (2020b) and Westphal and Sornette (2020a) among others. This agent-based model, which we presented in chapter 1, was shown to exhibit market bubbles, in the form of super-exponential growth of the returns of the price of the risky asset, in Kaizoji et al. (2015). We, in turn, studied three main strategies to reduce these price bubbles. The first one was a derivation of the work of Westphal and Sornette (2020a), with Dragon Riders optimized to enhance their Dragon Slaying properties, in chapter 2. We recall that Dragon Riders are agents who use their ability to predict extreme events to profit from this information, while Dragon Slayers are agents who use the same ability in order to prevent such extreme events. The denomination stems from the work of Sornette (2009); Sornette and Ouillon (2012), who proposed the name "Dragon-King" to describe statistical outliers of strong significance, such as asset price bubbles and their subsequent crashes. We also modeled a purely Dragon Slayer oriented agent implementing lean against the wind policy using open market operations in chapter 3. Lastly, we applied lean against the wind policy through interest rate modification, designing another Dragon Slayer for this purpose, in chapter 4.

As we found, regulatory policies can reduce certain bubbles significantly, especially when those are detected early, in our models. This goes in favor of the arguments of Roubini (2006) and Filardo (2005), who support early intervention for price-targeting monetary policy when price bubbles are forming. It is notable that Dragon Riders generally performed better in reducing the size and frequency of bubbles, than lean against the wind policies as we designed them, even when the Dragon Rider policy was optimized with respect to the Sharpe ratio of the trader. We recall in further detail that with a parameter setting designed for a better Dragon Rider performance, turning

60% of the fundamentalists into Dragon Riders could bring the number of peaks from 19.9 to 1.26, and bring their average size from 0.58 to 0.46 (recall that the initial price is set to 1). This is to be put in perspective with the fact that replacing all fundamentalists with the Dragon Slayer traders of chapter 3 brought the number of peaks from 19.9 to 13.6 only, while their average size was brought to 0.45. And finally, the interest rate modifying Dragon Slayers of chapter 4 brought the number of peaks to 14.35 and actually raised their average size to 0.67. This could mean that in this setting, self-regulation has better effects on the market than external intervention by a regulatory entity. However, this discrepancy could also be due to the differences in the design of the policies. Note that the interest rate policy was less effective than the open market operations, which could be dues to issues of assessing of the effect of such a transversal policy as interest rate control. We also cannot exclude the hypothesis that differently shaped lean against the wind policies could have more significant impact on the market. Finally, also note that we found ways to optimize the Dragon Riders such that exceptionally favorable behavior of the market could be induced. One might ask if these strategies would work in real market scenarios, or if they are a result of overfitting to our simple model. Similarly, one might ask if early intervention in asset price targeting is only worth pursuing in our simple model. Such a question could be answered through extension of the model, robustness testing with different parameters of the model, or real data applications of the model.

To go further into detail, among others, this work could be extended in the following ways. It would be pertinent to observe the applications of these regulatory policies to extensions of the initial model of Kaizoji et al. (2015). For instance, in the work of Chiarella et al. (2009), a higher dimension of the risky asset price is introduced. Assessing the impact of the regulatory policies of our model to this type of more complex market could uncover valuable insights. Different currencies could also be introduced, as in Westerhoff (2008), to account for exchange rate intervention by the central banks. Furthermore, quantifying interaction between the different regulatory instruments we studied would help understand the optimal behavior of central banks better. One might also model fiscal policies instead and study their interaction with our model. And, of course, the application of all these strategies to real market data would definitely yield stronger confidence in the use of monetary instruments that we prone in this report.

Appendix A

Appendix

In this appendix, we detail the calculations for a study on what amplitude should be chosen for an interest rate prescription formula, in order to modify the price path of a two type of agent model introduced in chapter 1. We attempted to assess the effect of interest rate r_t^f modifications on the price at t. More precisely, we attempted to obtain an estimate of $\frac{\partial p_t}{\partial r_t^f}$. The larger ambition here was to obtain a formula for a_t from 4.6 as a function of the state of the market at $t-1$ by examining the variation of price at t with respect to the interest rate modifications at $t-1$. We used Taylor expansion on the price equation to achieve this goal in first order approximation. Note that the interest rate r_t^f, applying to the risk-free asset from t to $t+1$, is chosen at $t-1$. In practice, we calculated the impact of a variation dr_t^f on the price p_t. However this was inconclusive in providing good amplitude prescription. There is a discussion to be had about whether the Dragon Slayer would realistically be informed about the risky fractions of the representatives of other traders, but for the purposes of this study, we assume it does.

Price equation with time-dependent interest rate Let us recall the implications of the equilibrium of supply and demand. As in chapter 1, this translates into the following equation for the excess demand of the fundamentalists and noise traders:

$$\Delta D_t^f + \Delta D_t^n = 0 \tag{A.1}$$

The difference is that the interest rate, featured in both those excess demands, is now dependent on time. Because of the time-dependent interest rate in the wealth dynamics, as featured in equation (4.3), the excess demand formula

for both traders now becomes for $i \in \{n, f\}$:

$$\begin{aligned}\Delta D_t^i &= x_t^i W_t^i - x_{t-1}^i W_{t-1}^i \frac{p_t}{p_{t-1}} \\ &= W_{t-1}^i \left(x_t^i \left[1 + r_t^f + x_{t-1}^i (r_t + \frac{d_t}{p_{t-1}} - r_t^f) \right] - x_{t-1}^i \frac{p_t}{p_{t-1}} \right) \qquad \text{(A.2)} \\ & \qquad \text{(A.3)}\end{aligned}$$

Therefore, we can use the following equation:

$$\begin{aligned}0 = \;& \Delta D_t^f + \Delta D_t^n \\ = \;& W_{t-1}^f \left(x_t^f \left[1 + r_t^f + x_{t-1}^f (r_t + \frac{d_t}{p_{t-1}} - r_t^f) \right] - x_{t-1}^f \frac{p_t}{p_{t-1}} \right) \\ & + W_{t-1}^n \left(x_t^n \left[1 + r_t^f + x_{t-1}^n (r_t + \frac{d_t}{p_{t-1}} - r_t^f) \right] - x_{t-1}^n \frac{p_t}{p_{t-1}} \right), \qquad \text{(A.4)}\end{aligned}$$

in order to derive the price as a fuction of the interest rate change Δr_t^f. The same derivations lead to a quadratic equation of the form

$$a_t p_t^2 + b_t p_t + c_t = 0, \qquad \text{(A.5)}$$

with

$$\begin{aligned}a_t = \;& \frac{1}{p_{t-1}} \left(W_{t-1}^f x_{t-1}^f (\frac{E_{r_{t+1}} - r_t^f}{\gamma \sigma^2} - 1) + W_{t-1}^n x_{t-1}^n (x_t^n - 1) \right) \qquad \text{(A.6)} \\ b_t = \;& W_{t-1}^f \left(x_{t-1}^f \frac{1}{p_{t-1}} \frac{d_t (1 + r_d)}{\gamma \sigma^2} + \frac{E_{r_{t+1}} - r_t^f}{\gamma \sigma^2} \left(x_{t-1}^f (\frac{d_t}{p_{t-1}} - r_t^f - 1) + r_t^f + 1 \right) \right) \\ & + W_{t-1}^n x_t^n \left(x_{t-1}^n (\frac{d_t}{p_{t-1}} - r_t^f - 1) + r_t^f + 1 \right) \qquad \text{(A.7)} \\ c_t = \;& W_{t-1}^f \frac{d_t (1 + r_d)}{\gamma \sigma^2} \left(x_{t-1}^f (\frac{d_t}{p_{t-1}} - r_t^f - 1) + r_t^f + 1 \right) \qquad \text{(A.8)}\end{aligned}$$

Which obtains a positive root

$$p_t = \frac{-b_t - \sqrt{b_t^2 - 4 a_t c_t}}{2 a_t} \qquad \text{(A.9)}$$

as explained in chapter 1. Now, we will draw a few observations from this equation before obtaining an estimate of $\frac{\partial p_t}{\partial r_t^f}$.

Observations on the relationship between r_t^f and p_t From a high-level viewpoint, we can already infer the nature of the relationship between p_t and r_t^f, all other things being equal.
Note that for the purpose of the approximations, a lot of parameters can be taken around their expected value, independently of t. However, the price at $t-1$ is not one of them, and neither is the dividend d_{t-1} (depending on the growth rate r_d, it can heavily depend on t.
As a first understanding of the price equation, let us note that all other things being equal, the parameter a_t from (A.5) is an affine function of r_t^f, b_t is a degree 2 polynomial in r_t^f, and c_t is affine in r_t^f. This makes the numerator of the function approximately a polynomial of degree 2 in r_t^f (approximately because there is a square root of a power of 4 involved), and the denominator an affine function of r_t^f. p_t is an algebraic fraction in r_t^f. Asymptotically, when r_t^f becomes very large, p_t is linear in r_t^f. Note that p_t is well defined when $r_t^f = 0$, i.e. it has a value in $\mathbb{R}$.

Taylor expansion Let us now derive a Taylor expansion to understand the effect on the price of a slight modification of the interest rate dr_t^f such that $dr_t^f \ll r_t^f$. We are interested in a first degree approximation, in order to get an estimate of what the amplitude of corrections should look like. To that end, we rewrite the factors a_t, b_t, c_t as functions of r_t^f:

$$a_t = \alpha_a r_t^f + \beta_a, \tag{A.10}$$

$$b_t = \alpha_b (r_t^f)^2 + \beta_b r_t^f + \gamma_b, \tag{A.11}$$

$$c_t = \alpha_c r_t^f + \beta_c, \tag{A.12}$$

with

$$\alpha_a = -\frac{W_{t-1}^f}{p_{t-1}}\frac{x_{t-1}^f}{\gamma\sigma^2}, \tag{A.13}$$

$$\beta_a = \frac{W_t^n x_{t-1}^n (x_t^n - 1)}{p_{t-1}} + \frac{W_{t-1}^f x_{t-1}^f \left(E_{r_{t+1}} - \gamma\sigma^2\right)}{p_{t-1}\gamma\sigma^2}, \tag{A.14}$$

$$\alpha_b = -W_{t-1}^f \frac{1 - x_{t-1}^f}{\gamma\sigma^2}, \tag{A.15}$$

$$\beta_b = \frac{W_{t-1}}{\gamma\sigma^2}\left(E_{r_{t+1}}(1 - x_{t-1}^f) - 1 - x_{t-1}^f\left(\frac{d_t}{p_{t-1}} - 1\right)\right) + W_{t-1}^n x_t^n (1 - x_{t-1}^n), \tag{A.16}$$

$$\gamma_b = \frac{W_{t-1}^f}{\gamma\sigma^2}\left[\frac{x_{t-1}^f d_t(1 + r_d)}{p_{t-1}} + E_{r_{t+1}}\left(x_{t-1}^f\left(\frac{d_t}{p_{t-1}} - 1\right) + 1\right)\right] + W_{t-1}^n x_t^n \left(x_{t-1}^n\left(\frac{d_t}{p_{t-1}} - 1\right) + 1\right), \tag{A.17}$$

$$\alpha_c = W_{t-1}^f (1 - x_{t-1}^f)\frac{d_t(1 + r_d)}{\gamma\sigma^2}, \tag{A.18}$$

$$\beta_c = W_{t-1}^f \frac{d_t(1 + r_d)}{\gamma\sigma^2}\left[1 + x_{t-1}^f\left(\frac{d_t}{p_{t-1}} - 1\right)\right] \tag{A.19}$$

Then, using these more compact notations, we can infer a formula for the first degree expansion of the price with respect to r_t^f. First, we calculate the Taylor expansion of $\frac{1}{2a_t}$ for small variations of r_t^f:

$$\begin{aligned}\frac{1}{2a_t}\bigg|_{r_t^f + dr_t^f} &\approx \frac{1}{2a_t}\bigg|_{r_t^f} - \frac{1}{2}\frac{\partial_R a_t}{a_t^2}\bigg|_{r_t^f} dr_t^f \\ &=: A_1 dr_t^f + B_1,\end{aligned} \tag{A.20}$$

with

$$\begin{aligned}A_1 &= \frac{-\alpha_a}{2(\alpha_a r_t^f + \beta_a)^2}, \\ B_1 &= \frac{1}{2a_t}\bigg|_{r_t^f}.\end{aligned} \tag{A.21}$$

Now, we calculate the first degree Taylor expansion of $-b_t - \sqrt{b_t^2 - 4a_t c_t}$:

$$\begin{aligned} -b_t - \sqrt{b_t^2 - 4a_t c_t}\Big|_{r_t^f + dr_t^f} \approx & \left[-b_t - \sqrt{b_t^2 - 4a_t c_t}\right]_{r_t^f} \\ & - \left[\partial_R b_t \frac{2b_t \partial_R b_t - 4a_t \partial_R c_t - 4c_t \partial_R a_t}{2\sqrt{b_t^2 - 4a_t c_t}}\right]_{r_t^f} dr_t^f \\ = & \; A_2 dr_t^f + B_2, \end{aligned} \tag{A.22}$$

with

$$\begin{aligned} A_2 &= -\left(2\alpha_b r_t^f + \beta_b\right) \left[\frac{b_t(2\alpha_b r_t^f + \beta_b) - 2(\alpha_a c_t + \alpha_c a_t)}{\sqrt{b_t^2 - 4a_t c_t}}\right]_{r_t^f}, \\ B_2 &= \left[-b_t - \sqrt{b_t^2 - 4a_t c_t}\right]_{r_t^f}. \end{aligned} \tag{A.23}$$

With this compact notation, we simply have:

$$\frac{dp_{t+1}}{dr_t^f} \approx A_1 B_2 + A_2 B_1 \tag{A.24}$$

Calibrating the corrections using the approximation Based on the approximate value of $\frac{dp_{t+1}}{dr_t^f}$, we can decide how to calibrate the amplitude of corrections on the interest rate that we perform using the Dragon Slayer. Ideally, if we have the exact value of $\frac{dp_{t+1}}{dr_t^f}$, we can deduce a linear approximation of the modification we need to perform in order to approach the desired value. Indeed, if:

$$\frac{dp_{t-1}}{dr_t^f} \approx A_1 B_2 + A_2 B_1$$

then the consequence to a perturbation dr_t^f will be

$$dr_t^f \approx \frac{dp_{t-1}}{A_1 B_2 + A_2 B_1} \tag{A.25}$$

and if we want to maintain a constant return $r_t = \bar{r}$, for example, we can deduce the following approximations:

$$\begin{aligned} & r_t && = \bar{r} \\ \iff & \frac{p_t + dp_t - p_{t-1}}{p_{t-1}} && = \bar{r} \\ \iff & y_t + \bar{r} + \frac{dp_t}{p_{t-1}} && = \bar{r} \\ \iff & dp_t && = -y_t p_{t-1} \end{aligned}$$

This gives us the target dp_t we aim for. We use the approximation $y_t \approx y_{t-1}$ to obtain an approximate prescription of dr_t^f:

$$dr_t^f \approx \frac{-y_t p_{t-1}}{A_1 B_2 + A_2 B_1} \tag{A.26}$$

The above result (A.26) could theoretically be used the foundation for the amplitude in the formula for r_t^f. However, this was inconclusive in our study, because of two main reasons. Firstly, it is doubtful that the Dragon Slayer could be expected to know the state of the market so completely at $t-1$. Applying averages to this equation gives unsatisfying results, probably because of some high parameter sensitivity that is overlooked in such an attempt. Secondly, the concept of the approach is nearing on overfitting to this specific model. We want to be able to generalize our results to other extensions of the model.

Bibliography

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, 503–546.

Ang, A., & Timmermann, A. (2012). Regime changes and financial markets. *Annu. Rev. Financ. Econ.*, *4*(1), 313–337.

Arrow, K. J. (1965). *Aspects of the theory of risk-bearing*. Yrjö Jahnssonin Säätiö.

Axelrod, R. (2006). Agent-based modeling as a bridge between disciplines. *Handbook of computational economics*, *2*, 1565–1584.

Baumol, W. J. (1957). Speculation, profitability, and stability. *The Review of Economics and Statistics*, 263–271.

Bernanke, B., & Gertler, M. (2000). *Monetary policy and asset price volatility* (Tech. Rep.). National bureau of economic research.

Bernanke, B. S., & Gertler, M. (2001). Should central banks respond to movements in asset prices? *american economic review*, *91*(2), 253–257.

Board of Governors of the Federal Reserve System. (2020a). *Factors affecting reserve balance.* Retrieved 15.08.2020, from https://www.federalreserve.gov/releases/h41/

Board of Governors of the Federal Reserve System. (2020b). *Reserve requirements.* Retrieved 4.08.2020, from https://www.federalreserve.gov/monetarypolicy/reservereq.htm

Chan, N. T., LeBaron, B., Lo, A. W., & Poggio, T. (1999). Agent-based models of financial markets: A comparison with experimental markets. *Unpublished Working Paper, MIT Artificial Markets Project, MIT, MA*.

Chiarella, C., Dieci, R., & He, X.-Z. (2009). Heterogeneity, market mechanisms, and asset price dynamics. In *Handbook of financial markets: Dynamics and evolution* (pp. 277–344). Elsevier.

Cœuré, B. (2019). A tale of two money markets: fragmentation or concentration. In *speech at the ecb workshop on money markets, monetary policy implementation and central bank balance sheets* (Vol. 12).

Committee, B. M. (2019). Monetary policy frameworks and central bank market operations. *BIS Markets Committee Papers, October*, 95.

Eser, F., Carmona Amaro, M., Iacobelli, S., & Rubens, M. (2012). The use of the eurosystem's monetary policy instruments and operational framework since 2009. *ECB occasional paper*(135).

European Central Bank. (2020). *The eurosystem's instruments.* Retrieved 29.07.2020, from https://www.ecb.europa.eu/mopo/implement/html/index.en.html

Filardo, A. (2003). Should monetary authorities prick asset price bubbles. *Preliminary draft, Bank for International Settlements, Basel*.

Filardo, A. J. (2001). Should monetary policy respond to asset price bubbles? some experimental results. *Some Experimental Results (July 2001). FRB of Kansas City Working Paper*(01-04).

Filardo, A. J., et al. (2000). Monetary policy and asset prices. *Economic Review-Federal Reserve Bank of Kansas City*, *85*(3), 11–38.

Fisher, I. (1930). *Theory of interest: as determined by impatience to spend income and opportunity to invest it*. Augustusm Kelly Publishers, Clifton.

Friedman, M., & Friedman, M. (1953). *Essays in positive economics*. University of Chicago press.

Gatley, T. (2020). How to burst a bubble. *Gavekal Dragonomics*, 4.

Harras, G., Tessone, C. J., & Sornette, D. (2012). Noise-induced volatility of collective dynamics. *Physical Review E*, *85*(1), 011150.

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

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 & organization*, *74*(3), 149–162.

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 & Organization*, *112*, 289–310.

Khort, R. (2016). *The market impact of exploiting financial bubbles* (Unpublished doctoral dissertation). Master's thesis, ETH Zurich, Switzerland.

Kim, G.-r., & Markowitz, H. M. (1989). Investment rules, margin, and market volatility. *Journal of Portfolio Management*, *16*(1), 45.

Lera, S. C., & Sornette, D. (2017). Evidence of a bimodal us gdp growth rate distribution: A wavelet approach. *Quantitative Finance and Economics*, *1*(1), 26–43.

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.

Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of monetary Economics*, *15*(2), 145–161.

Mizuta, T. (2016). A brief review of recent artificial market simulation (agent-based model) studies for financial market regulations and/or rules. *Available at SSRN 2710495*.

Modigliani, F., & Miller, M. H. (1963). Corporate income taxes and the cost of capital: a correction. *The American economic review*, *53*(3), 433–443.

Modigliani, F., & Miller, M. H. (1965). The cost of capital, corporation finance, and the theory of investment: reply. *The American Economic Review*, *55*(3), 524–527.

Morgenstern, O., & Von Neumann, J. (1953). *Theory of games and economic behavior*. Princeton university press.

Roubini, N. (2006). Why central banks should burst bubbles. *International Finance*, *9*(1), 87–107.

Samanidou, E., Zschischang, E., Stauffer, D., & Lux, T. (2007). Agent-based models of financial markets. *Reports on Progress in Physics*, *70*(3), 409.

Shiller, R. J. (2000). *Irrational exuberance*. Princeton university press.

Smith, A. (1761). *The theory of moral sentiments*. Penguin.

Sornette, D. (2009). Dragon-kings, black swans and the prediction of crises. *arXiv preprint arXiv:0907.4290*.

Sornette, D., & Ouillon, G. (2012). Dragon-kings: mechanisms, statistical methods and empirical evidence. *The European Physical Journal Special Topics*, *205*(1), 1–26.

Sornette, D., Woodard, R., & Zhou, W.-X. (2009). The 2006–2008 oil bubble: Evidence of speculation, and prediction. *Physica A: Statistical Mechanics and its Applications*, *388*(8), 1571–1576.

Walras, L. (1954). *Elements of pure economics*. Routledge.

Westerhoff, F. H. (2008). The use of agent-based financial market models to test the effectiveness of regulatory policies. *Jahrbucher Fur Nationalokonomie Und Statistik*, *228*(2), 195.

Westphal, R., & Sornette, D. (2020a). How market intervention can prevent bubbles and crashes. , 38. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3683858

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

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).

Yan, W., Rebib, R., Woodard, R., & Sornette, D. (2012). Detection of crashes and rebounds in major equity markets. *International Journal of Portfolio Analysis and Management*, *1*(1), 59–79.

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

Authored by (in block letters):
For papers written by groups the names of all authors are required.

Name(s):	**First name(s):**

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the 'Citation etiquette' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date	**Signature(s)**

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.