INSTITUTE
OF ECONOMICS

Scuola Superiore
Sant'Anna

International Doctoral Programme in Economics

Essays on the emergence of endogenous financial fluctuations

A Thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy in Economics

Ph.D. Candidate:

Jacopo Staccioli

Supervisor:

Prof. Giovanni Dosi

Academic Year 2017–2018

This page intentionally left blank.

Abstract

The present Thesis investigates the root causes and the complex mechanisms underlying the emergence of fluctuations in stock markets' dynamics, with a special focus on those that originate from within the market, i.e. endogenously, as opposed to those arising as (possibly exaggerated) response to external, freshly available fundamental news. It consists of three independent and self-contained essays, in the sense that they tackle distinct research questions and none of them builds upon the others' results, each employing a very different scientific methodology but all attributable to the literature of complexity economics in its financial flavour. The first essay proposes a heterogeneous agent model of a stock market in which a risky asset and a bond are exchanged, and demand on behalf of a group of traders is subject to random shocks. The second develops a parsimonious agent-based model of a stock market in which a large population of high-frequency traders exchange a long-lived security. Finally, the third empirically investigate large databases of intra-day volatility trajectories from two major stock markets indices, the S&P500 and the EURONEXT 100, along the lines of functional principal component analysis.

JEL classification: C58, C62, C63, D84, G11, G12, G15.

Keywords: Heterogeneous Agents, Evolutionary Finance, Stylised Facts, Intra-day Financial Dynamics, Functional Principal Component Analysis.

Acknowledgements

I acknowledge my sincere appreciation of those who have contributed to the present Thesis and supported me in any way throughout its development.

First of all, I wish to express my heartfelt gratitude towards my supervisor Prof. Giovanni Dosi for the continuous guidance during my doctoral studies, for his patience, motivation, and immense knowledge. His mentoring has proven invaluable at every step of my research.

Very special thanks go to my co-authors, Prof. Pietro Dindo, Dr. Mauro Napoletano, and Prof. Matteo Barigozzi, who supported me both technically and spiritually in pursuing and refining my research. Without their precious help, the present Thesis would not have seen the light.

My earnest thanks also go to Prof. Giulio Bottazzi, Prof. Mikhail Anufriev, and Prof. Alessio Emanuele Biondo, whose comments have helped rounding the previously sharp corners of my research ideas.

I thank my fellow colleagues and acquaintances at the Scuola for the stimulating discussions and for all the regenerating carefree moments we have spent together during the last four years.

I'm grateful to the OFCE – Sciences Po, Sophia Antipolis, France, and the Department of Statistics, London School of Economics and Political Science, London, UK, for granting me access to their research facilities and social networks during my visiting periods.

Last, but certainly not least, I wish to thank my family and Marica. None of them has ever doubted my ability and commitment to finish this journey, even at times in which I had doubts myself.

Contents

List of Figures

List of Tables

This page intentionally left blank.

1 Introduction

The present Thesis investigates the root causes and the complex mechanisms underlying the emergence of fluctuations in stock markets' dynamics, with a special focus on those that originate from within the market, i.e. endogenously, as opposed to those arising as (possibly exaggerated) response to external, freshly available fundamental news. The recent history of financial turmoil events, including prolonged financial crises and sudden flash crashes, demonstrates the importance, under both the positive and the normative perspectives, of a thorough understanding of financial fluctuations for the development of market regulation aimed at mitigating their negative effects. Yet, to date, there is still little agreement on behalf of policymakers on why financial markets occasionally experience sudden and unexpected turbulence, and what are the necessary actions to prevent it or alleviate its symptoms. What is clear however, more so to practitioners than to scholars, is that many episodes regularly unfolding on financial markets cannot be properly explained on the grounds of traditional neoclassical theories based on the so-called efficient market hypothesis (EMH). Loosely speaking, the latter claims that markets are efficient, in the sense that

> *"security prices at any time 'fully reflect' all available information".*
>
> (Fama, 1970, p. 383)

The logically consequent corollary is that, especially for short-term horizons in which risk-premia can be assumed to be constant, stock market prices evolve as a random walk and future price changes are completely unpredictable (see Bachelier, 1900; Samuelson, 1965). In other words, no amount of computation over what is known today would in any way improve a forecast of the price change tomorrow. As Malkiel puts it ,

> *"Taken to its logical extreme,* [the random walk hypothesis] *means that a blindfolded monkey throwing darts at a newspaper's financial pages could*

> *select a portfolio that would do just as well as one carefully selected by the experts."* (Malkiel, 2016, p. 24)

This stands clearly at odds with empirical evidence (and popular wisdom) suggesting the existence of sophisticated techniques, e.g. fundamental and technical analysis, whose skills have been for decades highly demanded and rewarded by financial institutions, which are able to 'beat the market' in a way that is far too systematic than what can be expected from for a purely random draw (see Lo and MacKinley, 2002 and references therein; for a critique see Malkiel, 2003; Timmermann and Granger, 2004). American business magnate Warren E. Buffett, himself a critic of the EMH (see Buffett, 1984), is but an example. A few studies identify a significant relation between price-earning ratios and subsequent returns performance (Basu, 1977; Bhargava, 2014; Breen and Savage, 1968; Breen, 1968; Dreman and Berry, 1995; Jensen, 1978; McWilliams, 1966; Nicholson, 1968). Others find consistent seasonality effects which are unrelated to fundamental news, e.g. around the turn of the week, of the month, of the year, and around holidays (Lakonishok and Smidt, 1988; Thaler, 1987), and correlations between market performance and bizarre variables such as weather (Hirshleifer and Shumway, 2003), traders' morning testosterone levels (Coates and Herbert, 2008) and their physical manifestations of stress (Lo et al., 2005).

Motivated by the overwhelming empirical evidence (of which the aforementioned is but an instance) and by a different introspection, based on psychological and cognitive biases affecting economic decision-making on the one hand, and on the evolutionary process of biological and ecological systems on the other, a growing population of theorists from different branches of the economic discipline have recently questioned previously incontrovertible notions such as market efficiency, equilibrium and perfect rationality (see e.g. Barberis, 2013; Benartzi and Thaler, 1995; Bouchaud, 2008; Campbell, 2000; Farmer, 2002; Farmer and Lo, 1999; Hirshleifer, 2001; Hirshleifer and Luo, 2001; Hirshleifer, 1977; Kahneman, 2003; Kahneman and Tversky, 1979; Kirman, 1992; Lo, 2004; Nelson and Winter, 1982). Within the financial literature this revolution, largely still under way, has spurred a number of mostly interrelated fields. A few labels include behavioural finance (Barberis and Thaler, 2003) and evolutionary finance (Hens and Schenk-Hoppé, 2005), whose boundaries are, to a good extent, a matter of personal taste. These fields are united in relaxing some of the assumptions of rigorous neoclassical economics and in embracing the possibility of framed and boundedly-rational decision making on behalf of heterogeneous agents based upon simple heuristics, thereby encompassing the possibility of

prolonged mispricing and of inefficient market selection mechanisms. The economy is regarded as a complex adaptive system, namely a system with multiple elements adapting or reacting to the pattern that these elements themselves create, and in which a perfect understanding of the individual parts does not automatically convey a similar understanding of the whole system's behaviour (see Arthur, 1999, 2014; Kirman, 2011; Miller and Page, 2007; Rosser, 1999).

> *"Such systems arise naturally in the economy. Economic agents, be they banks, consumers, firms, or investors, continually adjust their market moves, buying decisions, prices, and forecasts to the situation these moves or decisions or prices or forecasts together create. But unlike ions in a spin glass, which always react in a simple way to their local magnetic field, economic elements (human agents) react with strategy and foresight by considering outcomes that might result as a consequence of behaviour they might undertake. This adds a layer of complication to economics that is not experienced in the natural sciences. Conventional economic theory chooses not to study the unfolding of the patterns its agents create but rather to simplify its questions in order to seek analytical solutions."*
>
> (Arthur, 1999, p. 107)

Complexity economics, rather than being a new theory in competition with the established neoclassical paradigm, is a generalisation of the latter that accounts for out-of-equilibrium behaviour, as opposed to restricting uniquely to strategies that consistently induce no further reaction from the system's ecology.

The concept of bounded rationality captures both the cognitive limitations of economic actors and the limited time available to them when making decisions. This is especially true in a hectic environment such as that of a stock exchange trading floor, in which split-second decisions have far-reaching implications in terms of gains and losses. As Herbert A. Simon recognises,

> *"it appears probable that, however adaptive the behavio[u]r of organisms in learning and choice situations, this adaptiveness falls far short of the ideal of 'maximi[s]ing' postulated in economic theory. Evidently, organisms adapt well enough to 'satisfice'; they do not, in general, 'optimi[s]e'."*
>
> (Simon, 1956, p. 129)

> *"Because administrators satisfice rather than maximi[s]e, they can choose without first examining all possible behavio[u]r alternatives and without ascertaining that these are in fact all the alternatives. Because they treat the world as rather empty and ignore the interrelatedness of all things [...],*

> *they can make their decisions with relatively simple rules of thumb that do not make impossible demands upon their capacity for thought. Simplification may lead to error, but there is no realistic alternative in the face of the limits on human knowledge and reasoning."* (Simon, 1997, p. 119)

Bounded rationality should not be seen as a limitation, but rather as an enhancing mechanism by means of which an economic agent can make a (possibly sub-optimal) decision when the postulates required by rational maximising behaviour are not met (see also Conlisk, 1996; Simon, 1979). In this framework, the concept of heuristic is a leading instrument within the so-called 'adaptive toolbox' (Goldstein and Gigerenzer, 2002).

> *"A heuristic is a strategy that ignores part of the information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods."* (Gigerenzer and Gaissmaier, 2011, p. 454)

Gigerenzer (2008) argues that the intuition of heuristics producing second-best results is a misconception. He recognises that, in many situations, carrying out a proper optimisation procedure is hindered by computational intractability, while in others the latter may turn out to be less accurate because of lack of robustness and estimation errors. Interestingly, DeMiguel et al. (2009) show that the out-of-sample performance of the naïve diversification strategy $1/N$ can exceed that of a number of supposedly more 'optimal' strategies, including the Markowitz (1952) mean-variance portfolio and others based on sophisticated Bayesian techniques.[1] The intrinsic difficulty attached to the process of gathering what Fama (1970) terms *"all available information"* had been well recognised by John M. Keynes long before the EMH was even conceived. He argues that investors' sentiment and market psychology (in his own words, *"animal spirits"*) play an important role in everyday financial markets (see also Shiller, 2016).

[1] Anecdotal evidence reports that Markowitz himself did not use his award-winning optimization technique for his own retirement investments, relying instead on the naïve $1/N$ heuristic (Gigerenzer, 2008). Similarly, in spite of his sophisticated stock-picking value-investing techniques based on careful fundamental analysis, Warren E. Buffett's will consists of extraordinarily simple rules:

> *"My advice to the trustee could not be more simple: Put 10% of the cash in short-term government bonds and 90% in a very low-cost S&P500 index fund. [...] I believe the trust's long-term results from this policy will be superior to those attained by most investors – whether pension funds, institutions or individuals – who employ high-fee managers."* (2013 Shareholder Letter, Berkshire Hathaway Inc., p. 20)

> *"Investment based on genuine long-term expectation is so difficult as to be scarcely practicable. He who attempts it must surely lead much more laborious days and run greater risks than he who tries to guess better than the crowd how the crowd will behave; and, given equal intelligence, he may make more disastrous mistakes."* (Keynes, 1936, Chapter 12)

This excerpt also highlights the role played by the so-called 'beauty contest' mechanism, according to which instead of picking a stock (a person's face in the original Keynesian metaphor) that she genuinely considers undervalued (attractive), e.g. because the issuing firm displays sound fundamentals, a short-term trader may well be better off guessing what would the majority of other traders' choose, and then make a selection based on some inference from her knowledge of public perceptions. Even better, a trader could infer what other traders infer about the general perception, and so on and so forth.

> *"It is not a case of choosing those which, to the best of one's judg[e]ment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees."* (Keynes, 1936, Chapter 12)

This type of behaviour, which serves as a microfoundation for many principles of Dow theory and technical analysis, can generate self-sustaining departure of market prices from the consensus of underlying fundamental values, and it can ultimately fuel so-called rational bubbles (Allen et al., 2006).

The modelling attempts that embrace the aforementioned complex systems philosophy have proven able to reproduce some of the empirical properties, the so-called 'stylised facts', that conventional economic theory struggles to reconcile (milestones in this literature include Arthur et al., 1997; Blume and Easley, 1992, 2006; Brock and Hommes, 1998; Levy et al., 1994; Lux, 1995, 1998; Lux and Marchesi, 2000; Sandroni, 2000, and more recent follow-ups). Due to the intrinsic non-linear nature of boundedly-rational behaviour and the high dimensionality associated with agents' heterogeneity, these models raise important methodological challenges. In particular, there exists a trade-off between deductive methods that necessitate unrealistic simplifying assumptions (albeit more realistic than traditional neoclassical ones) and inductive methods which are however more difficult to falsify and generalise. Among the deductive methods, heterogeneous agents models (Hommes, 2006) build upon the mathematical theory of dynamical systems and propose analytically tractable equations

that describe the laws of motion of the economy. Since the complexity of these models grows dramatically with the underlying dimensionality, in order to obtain a close form solution, it is customary to focus on a very limited number of agent types (usually no more than two or three), and to assume simple pricing schemes (e.g. Walrasian) that have no actual counterpart in real market settings. Among the inductive methods, agent-based models (LeBaron, 2006) rely upon extensive numerical simulations and are therefore not constrained by analytical tractability issues. These models are usually very high-dimensional (it is not uncommon to include hundreds or thousands of different agents) and attempt at thoroughly replicating realistic economic scenarios. On the one hand, they provide an extremely flexible way of testing very complex behavioural and microstructural assumptions at virtually no cost (besides the time complexity of the underlying algorithms). On the other hand, the proper robustness of their result can be difficult to assess, and the extent of their generalisability is still an open question.[2]

The present Thesis is intended as a contribution to the literature of complexity economics in its financial flavour. It consists of three independent and self-contained essays, in the sense that they tackle distinct research questions and none of them builds upon the others' results, each employing a very different scientific methodology. The Thesis is organised as follows.

Chapter 2, titled *"Asset prices and wealth dynamics in a financial market with random demand shocks"*, proposes a heterogeneous agents model of a stock market in which a long-lived risky asset and a risk-free bond are exchanged, and demand on behalf of a group of traders is subject to random shocks. The impact of demand shocks upon market clearing prices depends on the relative wealth dynamics of the different groups of traders involved. By studying the local stability of deterministic and random fixed points of the random dynamical system whose state variables are the agents' relative wealth, the risky asset's return, and the dividend yield, we provide conditions on agents' portfolios under which such price impact is either maximal, when the traders subject to demand shocks dominate, minimal, when the traders subject to demand shocks vanish, or endogenously determined, when all traders survive and their relative wealth dynamics is a mean reverting process. Overall, the price impact of demand shocks brings an intrinsic penalty to the group of traders subject to them. Contrary to previous contributions in which portfolios that are more skewed towards the

[2]It is worth noting that controlled laboratory experiments also belong to the inductive class of methods (Bloomfield and Anderson, 2010).

risky security provide better survival chances, we find that a riskier portfolio may systematically fail to outperform a relatively safer one, should the former be subject to demand shocks.

Chapter 3, titled *"An agent-based model of intra-day financial markets dynamics"*, develops a parsimonious agent-based model of a stock market in which a risky security is exchanged at high frequency among a large population of heterogeneous traders. We design our simulations to closely replicate the timing structure of an existing stock market, namely the EURONEXT, in order to map the iterations of our algorithm with proper calendar time. We show that simple behavioural assumptions, such as chartist vs. fundamentalist behaviour and an endogenous participation scheme based on past profitability signals, are sufficient to simultaneously reproduce many of the stylised facts that financial markets exhibit at the high-frequency level of time granularity. These include properties related to returns (leptokurtosis, absence of linear autocorrelation, volatility clustering), trading volumes (volume clustering, correlation between volume and volatility), and timing of trades (number of price changes, autocorrelation of durations between subsequent trades, heavy tail in their distribution, order-side clustering).

Chapter 4, titled *"A 2-step functional principal component analysis of intra-day volatility trajectories"*, has a markedly statistical nature and does not directly implement the founding elements of complexity economics. However, it is meant to empirically support the modelling attempts therein by providing empirical evidence about some of the stylised facts that these models aim at reproducing. Moreover, it shares with the complex systems mindset an intrinsically high-dimensional character, ascribable to the realm of big-data analytics. We empirically investigate intra-day volatility trajectories from two major stock markets indices, the S&P500 and the EURONEXT 100, along the lines of recent developments in functional data analysis (FDA)· We propose a novel 2-step procedure based on functional principal component analysis for reducing the dimension of these large databases to a small set of curves. We furthermore propose a model, based on a CAPM-inspired distinction between the market component of the constituents' volatility, namely the part that correlates with the index as a whole, from the residual idiosyncratic component, that scores a good performance in predicting the original volatility trajectories. The analysis highlights the importance of the so-called common idiosyncratic volatility, i.e. the presence of a correlation structure across volatilities of distinct assets that persists after the underlying CAPM 'betas' have been filtered out, in the aforementioned predic-

tion exercise. Moreover, a visual inspection of the various functional principal components involved in our procedure and their corresponding loading coefficients gives some insight about a few stylised facts that financial markets' data typically exhibit (e.g. U-shaped intra-day activity and volatility clustering), and the patterns of international substitution and complementarity in place between the American and European stock markets under study.

To the best of our knowledge, all the aforementioned research questions have never been addressed by the incumbent literature.

This page intentionally left blank.

2

Asset prices and wealth dynamics in a financial market with random demand shocks[1]

2.1. Introduction

In this chapter, we investigate the impact of random demand shocks on the price dynamics of a risky asset. When a group of investors sell (buy) in a coordinated fashion, the impact on prices depends on the size of the group. If the group holds most of the supply of the asset, the latter sells at a large discount (premium). Conversely, if the group with demand shocks is small, the market is able to absorb their supply (demand) without a significant price adjustment. Therefore, the price impact of demand shocks upon market clearing prices depends on the relative wealth dynamics of different groups of traders. The purpose of this chapter is to investigate the long-run equilibrium interplay between wealth-driven market selection and the price impact of random demand shocks on behalf of a group of traders.

In order to address this question, we develop a heterogeneous agent model (HAM) mimicking a stylised financial economy in which a long-lived risky security and a risk-free bond are traded by a population of heterogeneous agents. The HAMs literature has emerged in the last couple of decades with the aim at modelling possible sources of endogenous risk, stemming from the mutual interaction of heterogeneous and boundedly rational trading strategies. We refer

[1]This chapter is a joint work with Pietro Dindo (Department of Economics, Università Ca' Foscari Venezia, Italy). The manuscript is published as Dindo, Pietro and Jacopo Staccioli (2018). 'Asset prices and wealth dynamics in a financial market with random demand shocks'. *Journal of Economic Dynamics and Control* 95, pp. 187–210. DOI: 10.1016/j.jedc.2018.08.009. Pietro Dindo acknowledges support from the Marie Curie International Outgoing Fellowship PIOF-GA-2011-300637 within the 7th European Community Framework Programme.

to Hommes (2006) for a survey. Contributions close to ours in spirit, namely that investigate the joint asset prices and wealth dynamics, include Chiarella and He (2001), Chiarella et al. (2006), Anufriev et al. (2006), Anufriev and Bottazzi (2010), Anufriev and Dindo (2010), Evstigneev et al. (2011), Anufriev et al. (2012), and Palczewski et al. (2016). To the best of our knowledge, no contribution within the HAMs literature has provided sufficient conditions for the survival of traders exposed to demand shocks and their impact on the asset price dynamics.

Within our exchange economy, investors' demand for the bond absorbs only a fraction of its total supply so that its price can be taken as fixed. Investors' demand for the risky asset instead absorbs all of its supply. As a consequence, traders are exposed, other than to the exogenous risk of its dividend payout, also to the endogenous risk of possible variations in its price. If, for example, a sizeable fraction of traders, as measured by aggregate wealth, decide to sell the asset, they will be able to do so only at a lower price. Conversely, when a small fraction of agents decides to sell the risky security, all of their demand is absorbed by the market without a significant price variation. A symmetric reasoning applies to the case in which a group of traders simultaneously decide to buy more of the risky asset. We assume that two groups of traders are active in the market. The first group is subject to joint demand shocks, which we model as an exogenous Markov process that moves their demand between two levels, either a high fraction or a low fraction of wealth to be invested in the risky asset. The second group is not vulnerable to demand shocks and maintains a constant position in the risky asset.

Assuming joint demand shocks that are modelled as an exogenous stochastic process, our contribution is also linked to the concept of *noise traders* and to the impact of noise unrelated to the exogenous dividend process, as put forth by Black (1986). This idea has been incorporated by different streams of literature to investigate the effect on market dynamics of portfolio decisions that are not strictly based on rational expectations (cf. e.g. DeLong et al., 1990, 1991). However, the HAMs literature has traditionally modelled noise traders with strategies driven by deterministic feedback mechanisms from realised market outcomes, mainly in terms of chartist (i.e. trend extrapolating) versus fundamentalist (i.e. mean reverting) rules (see e.g. Chiarella, 1992, Brock and Hommes, 1998). Little has been done so far to study the interaction between strategies incorporating a *truly* random component that continually perturbs the system away from equilibrium. A notable exception is Chiarella et al. (2011),

investigating a traditional fundamentalists vs. chartists asset pricing model in which the fundamental price follows a random walk process. Along the same lines, we extend the incumbent HAMs literature by analysing the equilibrium effect of a stochastic portfolio.

There might be multiple reasons for these shocks to occur synchronously within a group: all traders' demand is updated by looking at the same, possibly noisy, signals, e.g. an exogenous source of sentiment; alternatively, all these investors have the same institutional or financial constraints, e.g. they are all hedge funds who are overly exposed to the risky asset and might need to reduce their position due to new regulations or an industry shock.

As we shall discover, the most interesting scenario is the one in which the group that is vulnerable to demand shocks also acts more aggressively on the market, in the sense of committing to hold, on average, a larger share of wealth in the risky security, relative to the other group. On the one hand, taking more fundamental risk brings a positive reward since riskier securities carry higher expected excess-returns (when looking only at the exogenous component of risk). On the other hand, a portfolio which is exceedingly skewed towards the risky security also turns out to be vulnerable to losses due to a buy-high and sell-low phenomenon. In particular, when these aggressive traders sell and their market impact is high, they have to endure a substantial haircut. We show that the gain that follows a joint purchase is not large enough to completely offset the loss following a joint sale. We are able to analytically investigate this trade-off and provide sufficient conditions for survival of all traders and the ensuing endogenous asset prices fluctuations in equilibrium. In particular, we define the transmission of demand shocks onto market prices as the *pass-through*. This term, borrowed from signal processing, is already established in other fields of the economic discipline to denote the extent to which a certain signal (in our case, the stochastic process driving the demand shocks) is incorporated into another (in our case, the market clearing price of the asset).[2]

From a mathematical point of view, the model outlined above can be described by a random dynamical system. The state variables are the agents' relative wealth, the risky asset's return, and the dividend yield. The stochastic component plays the role of the joint demand shocks. Our system cannot be properly studied by perturbing its deterministic skeleton; since demand shocks

[2]Examples include how much of the foreign exchange-rate dynamics is reflected onto imported goods prices (exchange-rate pass-through), or how much of a tax levied on, say, producers is shifted to consumers (tax incidence pass-through).

constitute the main driver of the dynamics, their effect cannot be switched off without loss of generality. Technically, we study the local stability of the underlying deterministic *and* random fixed points.

Our results are as follows. From the aforementioned stability analysis, it is possible to distinguish three possible long-run scenarios. In particular, we shall provide sufficient conditions under which either the traders with a constant exposure dominate, i.e. there is no pass-through, or the traders subject to demand shocks dominate, i.e. there is maximal pass-through in equilibrium. A first finding is that the constant portfolio can systematically earn a higher return and eventually dominate the economy, even if it consists of a position which is on average strictly less than that of the stochastic one. The most interesting outcome from an economic standpoint arises when both the aforementioned polar equilibria are simultaneously unstable and neither group dominate. When this is the case, the extent of the pass-through becomes endogenous and is coupled with the agents' relative wealth dynamics. We show that the emergence of endogenous pass-through, due to long-run heterogeneity of traders either subject or immune to demand shocks, arises as a generic outcome, i.e. it occurs for a non-degenerate region of the parameters space.

Overall, the pass-through of demand shocks into market prices brings an intrinsic penalty for stochastic traders due to the aforementioned effect of buying high and selling low. In other words, we find that a riskier portfolio may systematically fail to outperform a relatively safer one, should the former be subject to demand shocks. The extent of such failure is tightly linked to the rate of growth of the dividend of the risky asset. Intuitively, this parameter balances the relative importance, for the stochastic traders' accumulation of wealth, of jumps in the risky asset's price when switching between high and low investment, and of dividend payments. Depending on its value, we can distinguish two regimes of long-run heterogeneity, and thus of endogenous pass-through. When the dividend grows quickly, *cœteris paribus*, the buy-high and sell-low negative effect is quickly overcome by the stream of dividend payments favouring those traders that are the most exposed in the risky asset. Conversely, when the growth rate is smaller than a certain threshold, both trader types may co-exist even when constant traders always invests a strictly lower fraction of wealth in the risky asset, compared to stochastic ones. As we shall see, both analytically and numerically, losses on behalf of stochastic traders can be, on average, so severe that both groups of agents survive.

The chapter is organised as follows. The financial model is outlined in the

next section, where the underlying random dynamical system governing the economy is derived. Then, Section 2.3 analyses the representative agent restriction, namely the no pass-through case and the maximal pass-through case. These two extreme cases are useful for the analysis of the full model, presented in Section 2.4. Here, sufficient conditions for all possible long-run outcomes and the related emergence of endogenous pass-through are obtained. Section 2.5 provides numerical simulation evidence of some interesting scenarios generated by our model and the sensitivity analysis for its relevant parameters. Finally, Section 2.6 concludes and proposes conceivable extensions to our framework.

2.2. The model

Our framework is common to a number of contributions in the HAMs literature, in particular Anufriev et al. (2006) and Anufriev and Dindo (2010). We first lay down the general market model and then introduce our specific assumptions regarding demand shocks in Section 2.2.1.

Consider a financial economy where a long-lived risky stock and a risk-free bond are traded in discrete time $t \in \mathbb{N} \cup \{0\}$. The amount of circulating shares of the risky security is constant, while the supply of the bond is perfectly elastic. The market is populated by an arbitrarily large number of traders divided in two groups indexed by $i = 1,2$. The detailed behaviour of each group is described shortly afterwards.

Before trade at time t starts, each trader in group i chooses the fraction $x_{i,t}$ of her current wealth $W_{i,t}$ to be invested in the risky asset. The latter pays a dividend D_t at the beginning of each period. The residual fraction of wealth, $1 - x_{i,t}$, is allocated to bond purchase, yielding a constant risk-free rate of return $r_f > 0$. The current level of individual wealth for group i depends on the past trading decision at time $t-1$, i.e. on the relative allocation of wealth between the risky and the risk-free investment, on the amount of dividend paid D_t, and on security prices. The price of the risky asset P_t is determined in equilibrium by equating aggregate demand and aggregate supply. Both the price and the dividend are expressed in terms of the bond's price, serving as the *numéraire*, conventionally normalised to 1 in every period. The evolution of wealth for the agents within group i reads

$$W_{i,t} = (1 - x_{i,t-1})W_{i,t-1}(1 + r_f) + x_{i,t-1}W_{i,t-1}\left(\frac{P_t + D_t}{P_{t-1}}\right). \tag{2.2.1}$$

The economy runs through a series of temporary equilibria in which the market

clearing condition is satisfied. By normalising, without loss of generality, the supply of the risky asset to 1, market clearing condition amounts to

$$P_t = \sum_{i=1,2} x_{i,t} W_{i,t}. \tag{2.2.2}$$

The system of eqs. (2.2.1) and (2.2.2) describes a growing economy at a rate that is in part deterministic, corresponding to the risk-free return r_f, and in part dependent on the realised price and dividend. It is easy to check this by summing the individual wealth over all the agents

$$W_t = \sum_{i=1,2} W_{i,t} = W_{t-1}(1+r_f) + \left[P_t + D_t - P_{t-1}\left(1+r_f\right)\right]. \tag{2.2.3}$$

It is therefore useful to get rid of the constant growth component r_f and define the rescaled variables

$$w_{i,t} = \frac{W_{i,t}}{(1+r_f)^t}, \qquad p_t = \frac{P_t}{(1+r_f)^t}, \qquad d_t = \frac{D_t}{(1+r_f)^t}, \tag{2.2.4}$$

and the associated rescaled model

$$\begin{cases} p_t = \sum_{i=1,2} x_{i,t} w_{i,t} \\ w_{i,t} = w_{i,t-1}\left[1 + x_{n,t-1}\left(\dfrac{p_t}{p_{t-1}} - 1 + e_t\right)\right] \end{cases}, \tag{2.2.5}$$

where e_t stands for the (rescaled) dividend yield, defined as

$$e_t := \frac{D_t}{P_{t-1}(1+r_f)} = \frac{d_t}{p_{t-1}}. \tag{2.2.6}$$

Note that in the first equation of system (2.2.5) the current price level p_t appears both in the LHS and in the RHS, as determinant of the level of wealth $w_{i,t}$. Substituting the second equation into the first and solving for p_t yields the explicit price dynamics

$$p_t = p_{t-1} \frac{\sum_{i=1,2} w_{i,t-1} x_{i,t} \left[1 + x_{i,t-1}(e_t - 1)\right]}{\sum_{i=1,2} w_{i,t-1} x_{i,t-1}(1 - x_{i,t})}. \tag{2.2.7}$$

It will prove useful for the subsequent analysis to normalise each group's wealth by total wealth. Computing the wealth share $\varphi_{i,t}$ we obtain

$$\varphi_{i,t} := \frac{w_{i,t}}{\sum_{i=1,2} w_{i,t}} = \varphi_{i,t-1} \frac{1 + x_{i,t-1}(r_t + e_t)}{1 + (r_t + e_t)\sum_{i=1,2} \varphi_{i,t-1} x_{i,t-1}}, \tag{2.2.8}$$

where the last equality comes from the second equation of system (2.2.5), and r_t denotes the capital gain (in excess of the risk-free return). It is immediate

to reformulate the latter in terms of individual wealth shares from eqs. (2.2.7) and (2.2.8) as follows:

$$r_t := \frac{p_t}{p_{t-1}} - 1 = \frac{\sum_{i=1,2} \varphi_{i,t-1} \left[x_{i,t}(1 + e_t x_{i,t-1}) - x_{i,t-1} \right]}{\sum_{i=1,2} \varphi_{i,t-1} x_{i,t-1}(1 - x_{i,t})}. \tag{2.2.9}$$

The return depends on the wealth-weighted average of the relative portfolio position of each group between the current and the preceding period, and on the dividend yield process. Regarding the latter, in the past literature a few distinct practices have emerged. Anufriev and Bottazzi (2010), Anufriev et al. (2006) and Chiarella and He (2001) all assume an endogenous dividend dynamics such that the dividend yield is an *i.i.d.* process; this directly implies that any change in the price of the risky security causes an essentially instantaneous adjustment in the paid dividend of an identical magnitude. Evstigneev et al. (2011) instead anchor the dividend to the aggregate wealth in the economy. As opposed to linking the dividend to the endogenous dynamics of the economy, Chiarella et al. (2006) and Anufriev and Dindo (2010) implement an exogenously growing dividend process with *i.i.d.* rate of growth. We follow the latter approach but for the sake of simplicity, in view of the already non-trivial effect of random demand shocks, we opt for a fully deterministic process. Switching the dividend noise off allows us to focus on traders' behaviour as the only source of randomness in the model.[3] That said, we make the following assumption:

Assumption 2.1. *The paid dividend grows geometrically at a rate* $g > 0$

$$d_t = d_{t-1}(1 + g). \tag{2.2.10}$$

We deliberately restrict to a (strictly) positive rate of growth since our focus is on the selective capacity of markets when risky assets, in the absence of demand shocks, yield a higher coupon with respect to the bond (at least on average). Strictly speaking, our Assumption implies that the risky asset has a higher return whenever its price grows at a constant rate. This is the reason why the incumbent literature (see e.g. Anufriev and Dindo, 2010) finds that fittest strategies are those most skewed towards the risky asset. However, as we shall see, the presence of random demand shocks might hinder such excess returns.

In terms of the dividend yield, Assumption 2.1 translates into

$$e_t = \frac{d_t}{p_{t-1}} = e_{t-1} \frac{1 + g}{1 + r_{t-1}}. \tag{2.2.11}$$

[3] Within a conceivable extension of the present model featuring a stochastic dividend dynamics, it would indeed be interesting to study the correlation between the two processes.

It is apparent in eq. (2.2.11) that a negative feedback coming from the past realised return negatively affects the current level of the yield. Eqs. (2.2.8), (2.2.9), and (2.2.11) together describe the overall dynamics of the wealth shares, the rate of return, and the dividend yield. They can be summarised in the following 3-dimensional system:

$$\begin{cases} \varphi_{1,t} = \varphi_{1,t-1} \dfrac{1 + x_{1,t-1}(r_t + e_t)}{1 + (r_t + e_t)\sum_{i=1,2} \varphi_{1,t-1} x_{1,t-1}} \\ r_t = \dfrac{\sum_{i=1,2} \varphi_{i,t-1} \left[x_{i,t} \left(1 + e_t x_{i,t-1}\right) - x_{i,t-1} \right]}{\sum_{i=1,2} \varphi_{i,t-1} x_{i,t-1} \left(1 - x_{i,t}\right)} \\ e_t = e_{t-1} \dfrac{1+g}{1+r_{t-1}} \end{cases} . \tag{2.2.12}$$

Note that since by definition $\sum_{i=1,2} \varphi_{i,t} = 1$, the wealth share of the second group can be residually derived as $\varphi_{2,t} = 1 - \varphi_{1,t}$.

2.2.1. Trader behaviour

The specification of traders' investment choice $x_{i,t}$, $i = 1,2$, closes the model. We restrict investment rules to those of a fixed-mix type (see Mulvey and Kim, 2010). This class of strategies ensure that at any point in time the fraction of wealth invested in each of the underlying assets is kept at a constant level. As opposed to the buy-and-hold rule, when a fixed-mix strategy is adopted the trader rebalances her portfolio to keep the weight $x_{i,t}$ of the risky security unchanged over time by selling (respectively, buying) the asset if its price has increased (respectively, decreased). Fixed-mix strategies differ from those typically embedded in other HAMs in that the weight of each asset is insensitive to realised market outcomes such as past returns or dividend yields. We make this assumption for analytical tractability purposes.

In order to be economically meaningful, the price of an asset supplied in finite amount needs to be, at the very least, strictly positive. A sufficient condition thereof, stemming from eq. (2.2.7) and widely adopted in the literature, coincides with preventing each trader from short-selling or leverage-buying the risky asset, i.e. requiring $x_{i,t} \in (0,1)$, $i = 1,2$. Note that this condition is not necessary, since the restriction only needs to apply to the aggregate market-portfolio. Note also that this assumption does not stand at odds with empirical evidence: for instance, studying a sample of U.S. domestic equity funds in the 1994–2000 period, Almazan et al. (2004) show that short-sale restrictions are not exceptional, due to both regulatory and self-imposed constraints.

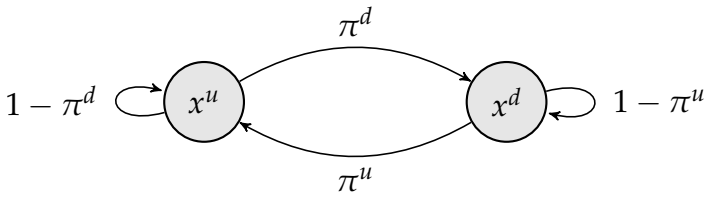

Figure 2.1.: The Markov process $\{x_t\}$.

The first group of traders is characterised by a constant fraction $x_{1,t} = \overline{x}$. We shall name this group the *constant* group and denote it by $\mathfrak{C}$. The second group of traders, instead, jump back and forth between two distinct levels of exposure, $x_{2,t} = x^u$ or $x_{2,t} = x^d$, with $x^u < x^d$, at random times. Hence, the portfolio of traders in such group consists of a stochastic variable with two admitted values. In particular, we model the transition between these two levels by an exogenous Markov process. We shall name this group the *stochastic* group and denote it by $\mathfrak{S}$.

The following Assumption fully summarises the behaviour of the two groups of agents.

Assumption 2.2. *Each trader in group $\mathfrak{C}$ adopts the portfolio rule $\overline{x} \in (0,1)$ in every period.*
Each trader in group $\mathfrak{S}$ adopts the portfolio rule $x_t \in \{x^u, x^d\}$ according to a Markov process $\{x_t,\ t \in \mathbb{N}\}$, characterised by the following transition probability matrix:

$$\mathcal{P} = \begin{array}{c} \\ x^u \\ x^d \end{array} \begin{array}{c} \begin{array}{cc} x^u & x^d \end{array} \\ \begin{pmatrix} 1-\pi^d & \pi^d \\ \pi^u & 1-\pi^u \end{pmatrix} \end{array}, \tag{2.2.13}$$

with

$$0 < x^d < x^u < 1, \qquad \pi^u > 0, \qquad \pi^d > 0. \tag{2.2.14}$$

Given transition probabilities $\mathcal{P}$, we name $\mathbb{P}$ the induced probability measure on sets of sequences $\{x_t,\ t \in \mathbb{N}\} \in \bigtimes_{\infty}\{x^u, x^d\}$.

The Markov process $\{x_t\}$, also pictured in Fig. 2.1, constitutes the driver of demand shocks. The transition probability π^d determines the probability of suffering a negative shock, conditional on being exposed to the risky asset at a *'up'* level. The transition probability π^u determines the probability of suffering a positive demand shock, conditional on being exposed to the risky asset at a

'down' level. Given transition probabilities, the average duration of a *'up'* and a *'down'* level of investment on a given realisation $\{x_t\}_{t=0}^{\infty}$ equals $(\pi^d)^{-1}$ and $(\pi^u)^{-1}$, respectively.

In terms of the two groups of traders just outlined, the 3-dimensional random dynamical system is the composition of group $\mathfrak{S}$ stochastic demand $\{x_t\}$ and of the four maps $\mathcal{F}_{x_{t-1},x_t} : \mathcal{D} \to \mathcal{D}$, each defined over the phase space $\mathcal{D} = \Delta \times (-1, +\infty) \times \mathbb{R}_{++}$ and given by

$$\mathcal{F}_{x_{t-1},x_t} : \begin{cases} \varphi_t = \varphi_{t-1} \dfrac{1 + x_{t-1}(r_t + e_t)}{1 + (r_t + e_t)\left[\varphi_{t-1} x_{t-1} + (1 - \varphi_{t-1})\overline{x}\right]} \\ r_t = \dfrac{\varphi_{t-1}\left[x_t(1 + e_t x_{t-1}) - x_{t-1}\right] + (1 - \varphi_{t-1}) e_t \overline{x}^2}{\varphi_{t-1} x_{t-1}(1 - x_t) + (1 - \varphi_{t-1})\overline{x}(1 - \overline{x})} \\ e_t = e_{t-1} \dfrac{1 + g}{1 + r_{t-1}} \end{cases}, \tag{2.2.15}$$

where we adopt the convention that φ_t denotes the time-t aggregate wealth share of traders within group $\mathfrak{S}$, i.e.

$$\varphi_t := \frac{w_{\mathfrak{S},t}}{w_{\mathfrak{C},t} + w_{\mathfrak{S},t}}, \tag{2.2.16}$$

so that $1 - \varphi_t$ residually stands for the time-t aggregate wealth share of traders within group $\mathfrak{C}$.

2.3. Representative trader economies

Let us first consider an economy populated by traders of the same type, either constant or stochastic. This case is insightful since it allows to properly disentangle the features of the long-run dynamics implied by each behavioural specification from those stemming from the market interaction induced by demand shocks. The resulting reduced system is lower dimensional with respect to (2.2.15) since the underlying wealth share φ_t is fixed to zero (respectively, one) whenever we consider an economy populated uniquely by constant (respectively, stochastic) traders.

2.3.1. The economy with a constant trader: no pass-through

When the economy is populated only by constant traders, demand shocks do not play a role and there is no pass-through. Imposing $\varphi_t = 0 \, \forall t$, the market

dynamics $\mathcal{F}$ in (2.2.15) is fully deterministic and reduces to $\tilde{\mathcal{F}}^{\mathfrak{C}} : \tilde{\mathcal{D}} \to \tilde{\mathcal{D}}$, with $\tilde{\mathcal{D}} = (-1, +\infty) \times \mathbb{R}_{++}$, given by

$$\tilde{\mathcal{F}}^{\mathfrak{C}} : \begin{cases} r_t = e_t \dfrac{\overline{x}}{1-\overline{x}} \\ \\ e_t = e_{t-1} \dfrac{1+g}{1+r_{t-1}} \end{cases} . \tag{2.3.1}$$

The following Proposition characterises the asymptotic dynamics of the economy.

Proposition 2.1. *The market dynamics* (2.3.1) *admits a unique, globally stable, fixed point*

$$\tilde{\mathfrak{C}} = \left(g \, , \, g \, \frac{1-\overline{x}}{\overline{x}} \right). \tag{2.3.2}$$

Proof. See Appendix A.1. □

At the fixed point, the price of the risky security grows, proportionally to the dividend, at a rate of g. Intuitively, since the asset is in finite constant supply, its price has to fully account for the new wealth injected by means of dividend payments. From the first equation of system (2.3.1), since the portfolio rule $\overline{x}$ is itself constant, the dividend yield level must adjust accordingly.

The existence of a globally stable value of dividend yield and return can be easily understood by looking exclusively at the dynamics of the dividend yield. In fact, by lagging and substituting the first equation into the second, it is easy to further reduce system (2.3.1) to a 1-dimensional map $\tilde{f}^{\mathfrak{C}} : \mathbb{R}_{++} \to \mathbb{R}_{++}$ solely in terms of the dividend yield:

$$e_t = \tilde{f}^{\mathfrak{C}}(e_{t-1}) = e_{t-1} \frac{1-\overline{x}}{1+\overline{x}(e_{t-1}-1)} (1+g). \tag{2.3.3}$$

Map (2.3.3) admits a unique non-trivial (i.e. strictly positive) fixed point

$$e^* = g \frac{1-\overline{x}}{\overline{x}}. \tag{2.3.4}$$

Substituting e^* into the first equation of system (2.3.1) yields the equilibrium return $r^* = g$. The fixed point is globally stable because for $g > 0$, as we assume, map (2.3.3) is increasing, concave, with derivative greater than unity in $e = 0$ and smaller than unity in e^*.[4]

[4] For the analysis of the $g \leq 0$ case we refer the reader to Anufriev and Dindo (2010, Section 3.1).

2.3.2. The economy with a stochastic trader: maximal pass-through

When the economy is populated only by stochastic traders, demand shocks do play an essential role and pass-through is maximal. Imposing $\varphi_t = 1\,\forall t$, the market dynamics $\mathcal{F}$ in (2.2.15) still depends on the last two investment levels. For each pair $(x_{t-1}, x_t) \in \bigtimes_2 \{x^u, x^d\}$, the returns and dividend yield dynamics can be written as $\tilde{\mathcal{F}}^{\mathfrak{S}}_{x_{t-1},x_t} : \tilde{\mathcal{D}} \to \tilde{\mathcal{D}}$, with $\tilde{\mathcal{D}} = (-1, +\infty) \times \mathbb{R}_{++}$, given by

$$\tilde{\mathcal{F}}^{\mathfrak{S}}_{x_{t-1},x_t} : \begin{cases} r_t = \dfrac{x_t\,(1 + e_t x_{t-1}) - x_{t-1}}{x_{t-1}\,(1 - x_t)} \\ e_t = e_{t-1} \dfrac{1+g}{1+r_{t-1}} \end{cases} . \tag{2.3.5}$$

The returns and dividend yield dynamics can be characterised by studying the random fixed point of (2.3.10), as defined in the following. As a matter of notation, we shall use x_{-1} and x_{-2} to denote the 1-period and 2-period lagged values of a generic realisation x of $\{x_t\}$.

Definition 2.1. *Call $\tilde{\mathcal{D}}^*$ the space of all random vectors $(R^*, E^*) : \bigtimes_2 \{x^u, x^d\} \to \tilde{\mathcal{D}}$. A* random *fixed point of system (2.3.5) is a random vector $(R^*, E^*) \in \tilde{\mathcal{D}}^*$ such that $(R^*, E^*)_{x_{-1},x} = \tilde{\mathcal{F}}_{x_{-1},x}(R^*, E^*)_{x_{-2},x_{-1}}$, for all (x_{-2}, x_{-1}) and (x_{-1}, x) in $\bigtimes_2 \{x^u, x^d\}$.*

The following Proposition shows that, in an economy where all traders are subject to demand shocks, there exists a globally stable random fixed point.

Proposition 2.2. *For every realisation of the Markov process $\{x_t\}$, the market dynamics (2.3.5) admits a unique (random) fixed point*

$$\tilde{\mathfrak{S}} = \left(g\,\frac{x(1-x_{-1})}{x_{-1}(1-x)} + \frac{x - x_{-1}}{x_{-1}(1-x)},\; g\,\frac{1-x_{-1}}{x_{-1}} \right), \qquad \forall (x_{-1}, x) \in \bigtimes_2 \{x^u, x^d\}. \tag{2.3.6}$$

Moreover, $\tilde{\mathfrak{S}}$ is globally stable, i.e.

$$\lim_{t\to\infty} \tilde{\mathcal{F}}^{\mathfrak{S}}_{x_{-1},x} \left(\tilde{\mathcal{F}}^{\mathfrak{S}}_{x_t,x_{-1}} \circ \tilde{\mathcal{F}}^{\mathfrak{S}}_{x_{t-1},x_t} \circ \cdots \circ \tilde{\mathcal{F}}^{\mathfrak{S}}_{x_0,x_1}(R_0, E_0) \right) = (R^*, E^*)_{x_{-1},x}, \tag{2.3.7}$$

for every sequence $\{x_t\}$, for all $(x_{-1}, x) \in \bigtimes_2 \{x^u, x^d\}$, and for every initial condition $(R_0, E_0) \in \tilde{\mathcal{D}}$.

Proof. See Appendix A.2. □

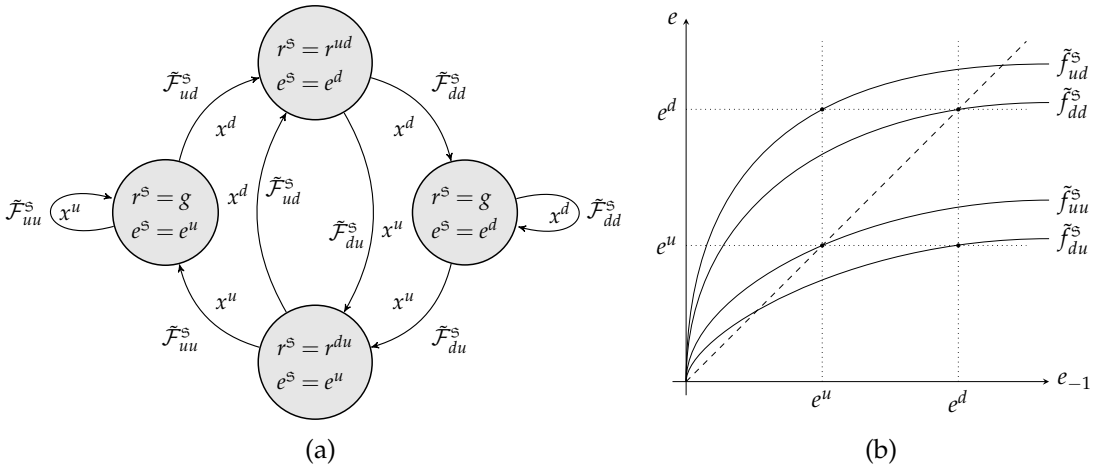

Figure 2.2.: (a): The four states associated with random fixed point $\tilde{\mathfrak{S}}$ and the transition between them depending on the realisation of the stochastic process $\{x_t\}$. (b): The four realisations of the stochastic map $\tilde{f}_{x_{t-2},x_{t-1}}$ and the bisector line (dashed).

Random fixed point $\tilde{\mathfrak{S}}$ consists of 4 states $(R^*, E^*)^{\tilde{\mathfrak{S}}}_{x_{-2},x_{-1}} \in \tilde{\mathcal{D}}$, associated to all possible couples $(x_{-2}, x_{-1}) \in \bigtimes_2 \{x^u, x^d\}$. The four maps $\tilde{\mathcal{F}}^{\tilde{\mathfrak{S}}}_{x_{-1},x}$, associated to all possible couples $(x_{-1}, x) \in \bigtimes_2 \{x^u, x^d\}$, govern the transition between them, as depicted in Fig. 2.2(a). In particular, there exist:

◇ three equilibrium values of the return

$$r^{\tilde{\mathfrak{S}}}(x_{-2}, x_{-1}) = \begin{cases} g & \text{if} \quad x_{-2} = x_{-1} \\ r^{ud} = g\dfrac{x^d(1-x^u)}{x^u(1-x^d)} - \dfrac{x^u - x^d}{x^u(1-x^d)} & \text{if} \quad (x_{-2}, x_{-1}) = (x^u, x^d) \\ r^{du} = g\dfrac{x^u(1-x^d)}{x^d(1-x^u)} + \dfrac{x^u - x^d}{x^d(1-x^u)} & \text{if} \quad (x_{-2}, x_{-1}) = (x^d, x^u) \end{cases}, \tag{2.3.8}$$

◇ two equilibrium values of the dividend yield

$$e^{\tilde{\mathfrak{S}}}(x_{-2}, x_{-1}) = \begin{cases} e^u = g\dfrac{1-x^u}{x^u} & \text{if} \quad x_{-2} = x^u \\ e^d = g\dfrac{1-x^d}{x^d} & \text{if} \quad x_{-2} = x^d \end{cases}. \tag{2.3.9}$$

In order to grasp an intuition about the existence of (and the convergence to) random fixed point $\tilde{\mathfrak{S}}$, it is instructive to rewrite the dynamics exclusively in

terms of dividend yields. By lagging and substituting the first equation into the second, it is easy to further reduce system (2.3.5) to a 1-dimensional stochastic map $\tilde{f}^{\mathfrak{S}}_{x_{t-2},x_{t-1}} : \mathbb{R}_{++} \to \mathbb{R}_{++}$ solely in terms of the dividend yield:

$$e_t = \tilde{f}^{\mathfrak{S}}_{x_{t-2},x_{t-1}}(e_{t-1}) = e_{t-1}\frac{x_{t-2}(1-x_{t-1})}{x_{t-1}\left[1+x_{t-2}(e_{t-1}-1)\right]}(1+g). \tag{2.3.10}$$

Map (2.3.10), pictured[5] in Fig. 2.2(b), depends on both past realisations x_{t-1} and x_{t-2} of the demand Markov process $\{x_t\}$. When these two realisations are equal, such as (x^u,x^u) or (x^d,x^d), map $\tilde{f}^{\mathfrak{S}}_{x_{t-2},x_{t-1}}$ in (2.3.10) corresponds to map $\tilde{f}^{\mathfrak{C}}$ in (2.3.3) when, respectively, $\overline{x} = x^u$ or $\overline{x} = x^d$. Analogously, the dividend yield of the random fixed point, e^u or e^d, corresponds to e^* in (2.3.4) when, respectively, $\overline{x} = x^u$ or $\overline{x} = x^d$. Note how a higher fraction of wealth invested in the risky asset corresponds to a higher pressure on risky asset prices and thus to a lower dividend yield. A shock on demand, e.g. from x^u to x^d, makes the system evolve according to map $\tilde{f}_{ud}$ and moves the dividend yield from e^u to e^d. Since, by assumption, no constant trader is present, the excess supply of the risky asset when the shock hits is neither absorbed nor mitigated, and the pass-through of demand shocks onto market price is maximised. The price of the risky asset has to adjust for the continually injected new wealth by means of dividend payments (in analogous fashion as for fixed point $\tilde{\mathfrak{C}}$), and for the jumps in the market portfolio induced by demand shocks. The opposite occurs when demand jumps from x^d to x^u.

Given the dividend yield and investment fractions, the first equation in (2.3.5) provides the corresponding value of the return. A closer inspection of eqs. (2.3.8) and (2.3.9) reveals that the following relations hold:

$$r^{ud} < g < r^{du}. \tag{2.3.11}$$

When $x_{-1} = x = x^i$, $i \in \{u,d\}$, the dynamics temporarily resembles that of fixed point $\tilde{\mathfrak{C}}$. Irrespectively of the prevailing portfolio fraction, the return matches the dividend growth rate g for the aforementioned reason. Conversely, the return is $r^{ud} < g$ whenever the demand shock is negative and stochastic traders pass from a high to a low investment in the risky asset. The opposite occurs when the demand shock is positive, and $r^{du} > g$.

When a negative demand shock hits, the capital gain r^{ud} may even be negat-

[5] Fig. 2.2(b) is intended as a qualitative picture; it is not generally true that $\lim_{e_{-1}\to 0^+} \tilde{f}'_{du} > 1$ since the slope of the map in fact depends on the underlying value of g.

ive. In particular it holds

$$r^{ud} < 0 \iff g < \bar{g} = \frac{x^u - x^d}{x^d(1 - x^u)}. \tag{2.3.12}$$

The threshold $\bar{g}$ below which r^{ud} is negative inversely depends on the magnitude of the portfolio shift in terms of wealth, i.e. on the difference $x^u - x^d$, and is influenced by the overall position of x^u and x^d within the unit simplex. For a mild drop of around 1% of wealth invested in the risky security, $\min(\bar{g}) \approx 0.04$ (and note that the function $\bar{g}(x^u, x^d)$ is sharply convex). When a more economically meaningful deviation occurs, say $x^u - x^d = 0.1$, then $\min(\bar{g}) \approx 0.49$, meaning that a substantial dividend growth rate is required for r^{ud} to be positive; such a high rate is clearly unsustainable in the long run and at sharp odd with actual markets evidence. Therefore, in most conceivable cases, in response to a downward portfolio shock, stochastic traders have to 'fictionally' pay a significant premium in order to immediately sell part of their holdings. Conversely, during an upward portfolio shift, stochastic traders exert an *ab*-normal pressure on the demand of the risky asset, captured in turns by a higher realised return $r^{du} > g$, eventually driving the market-clearing price upwards. As we shall see, both movements are mitigated by the presence of constant traders, as long as the latter survive in the long-run.

Proposition 2.2 not only characterises the random fixed point but also shows that it is globally stable, implying that the dynamics of dividend yield and return converges to the values in (2.3.8) and (2.3.9) for all possible initial values. The following Corollary draws the implications of convergence to random fixed point $\tilde{\mathfrak{S}}$ and of the dynamics of stochastic traders demand $\{x_t\}$ on the dynamics of returns and dividend yields in the long-run.

Corollary 2.1. *Given the Markov process portfolio rule $\{x_t,\ t \in \mathbb{N}\}$ as per Assumption 2.2, at random fixed point $\tilde{\mathfrak{S}}$ returns and dividend yields evolve according to an irreducible, time-homogeneous, Markov chain characterised by the following transition probability matrix:*

$$\tilde{\mathcal{P}} = \begin{array}{c} \\ (g,e^u) \\ (r^{ud},e^d) \\ (r^{du},e^u) \\ (g,e^d) \end{array} \begin{array}{c} \begin{array}{cccc} (g,e^u) & (r^{ud},e^d) & (r^{du},e^u) & (g,e^d) \end{array} \\ \begin{pmatrix} 1-\pi^d & \pi^d & 0 & 0 \\ 0 & 0 & \pi^u & 1-\pi^u \\ 1-\pi^d & \pi^d & 0 & 0 \\ 0 & 0 & \pi^u & 1-\pi^u \end{pmatrix} \end{array}. \tag{2.3.13}$$

All states are positive-recurrent and there exists a unique invariant distribution $\tilde{\boldsymbol{\pi}}$ *satisfying the condition* $\tilde{\boldsymbol{\pi}} = \tilde{\boldsymbol{\pi}}\tilde{\mathcal{P}}$. *In particular it holds:*

$$\tilde{\boldsymbol{\pi}} = \left[\frac{\pi^u(1-\pi^d)}{\pi^u+\pi^d}, \frac{\pi^u\pi^d}{\pi^u+\pi^d}, \frac{\pi^u\pi^d}{\pi^u+\pi^d}, \frac{\pi^d(1-\pi^u)}{\pi^u+\pi^d} \right]. \tag{2.3.14}$$

2.4. Heterogeneous traders economy

We can now proceed with studying the effect of demand shocks in an economy with heterogeneous traders, the latter partitioned into constant group $\mathfrak{C}$ and stochastic group $\mathfrak{S}$. To this aim, we need to investigate the asymptotic survival of both groups. To this purpose, we shall define the following Terminology.

Terminology. *Group* $i \in \{\mathfrak{C}, \mathfrak{S}\}$ *is said to*

- ◇ survive *if its asymptotic wealth-share is strictly positive* $\mathbb{P}$*-almost surely, i.e. if* $\mathbb{P}\left\{ \limsup_{t\to\infty} \varphi_{n,t} > 0 \right\} = 1$;
- ◇ vanish *if its asymptotic wealth-share is zero* $\mathbb{P}$*-almost surely, i.e. if* $\mathbb{P}\left\{ \limsup_{t\to\infty} \varphi_{n,t} = 0 \right\} = 1$;
- ◇ dominate *if its asymptotic wealth-share is one* $\mathbb{P}$*-almost surely, i.e. if* $\mathbb{P}\left\{ \liminf_{t\to\infty} \varphi_{n,t} = 1 \right\} = 1$.

We are interested in characterising all the possible long-run outcomes. A first possibility is that constant traders hoard all aggregate wealth in the long-run, so that demand shocks become negligible. In other words, should we find that the stochastic group vanishes asymptotically, then the impact of demand shocks would only be transient. In order to account for this outcome, we introduce the concept of deterministic fixed point of a random dynamical system.

Definition 2.2. *A deterministic fixed point of system* (2.2.15) *is a vector* $(\varphi^*, r^*, e^*) \in \mathcal{D}$ *such that* $(\varphi^*, r^*, e^*) = \mathcal{F}_{x_{-1},x}(\varphi^*, r^*, e^*)$ *for all* (x_{-1}, x) *in* $\bigtimes_2\{x^u, x^d\}$.

Alternatively, the stochastic trader could dominate. To study this outcome, we adapt the concept of random fixed point given for homogeneous traders economies in Definition 2.1 to the current heterogeneous traders set-up, leading to the following definition.

Definition 2.3. *Call* $\mathcal{D}^*$ *the space of all random vectors* $(\Phi^*, R^*, E^*) : \bigtimes_2\{x^u, x^d\} \to \mathcal{D}$. *A random fixed point of system* (2.2.15) *is a random vector* $(\Phi^*, R^*, E^*) \in \mathcal{D}^*$ *such that* $(\Phi^*, R^*, E^*)_{x_{-1},x} = \mathcal{F}_{x_{-1},x}(\Phi^*, R^*, E^*)_{x_{-2},x_{-1}}$, *for all* (x_{-2}, x_{-1}) *and* (x_{-1}, x) *in* $\bigtimes_2\{x^u, x^d\}$.

It is possible to show that, under the stated assumptions, system (2.2.15) admits exactly two fixed points, one for each of the the aforementioned types.[6]

Proposition 2.3. *System* (2.2.15) *admits:*

- ◇ *a deterministic fixed point* $\mathfrak{C}$ *in which the constant group dominates and the stochastic group vanishes*

$$\mathfrak{C} = \left(0,\, g\,,\, g\,\frac{1-\overline{x}}{\overline{x}}\right), \tag{2.4.1}$$

- ◇ *a random fixed point* $\mathfrak{S}$ *in which the stochastic group dominates and the constant group vanishes*

$$\mathfrak{S} = \left(1,\, g\,\frac{x(1-x_{-1})}{x_{-1}(1-x)} + \frac{x-x_{-1}}{x_{-1}(1-x)},\, g\,\frac{1-x_{-1}}{x_{-1}}\right). \tag{2.4.2}$$

Proof. See Appendix A.3. □

Importantly for our purposes, at fixed point $\mathfrak{C}$ demand shocks do not pass through to market price, for the fraction of wealth (and therein of risky asset) detained by traders who are subject to these shocks is negligible. On the contrary, at fixed point $\mathfrak{S}$ demand shocks do lead to maximal pass-through, since now it is the fraction of traders who are not subject to demand shocks that is negligible. In what follows we shall investigate the convergence to this two scenarios as well as the case in which both groups of agents survive. In the latter case, the pass-through is endogenous in that its extent is directly linked to the traders' relative wealth dynamics.

In order to characterise the stability of both fixed points,[7] we restrict to their local asymptotic analysis for tractability reasons. Let us define $\rho^i\big|_j$ as the expected geometric (gross) growth rate of the wealth of group i when the system is at fixed point j, with $i, j \in \{\mathfrak{C}, \mathfrak{S}\}$. The following Lemma is useful.

Lemma 2.1. *The local asymptotic stability of fixed point* $j \in \{\mathfrak{C}, \mathfrak{S}\}$ *is entirely determined by the relative values of* $\rho^{\mathfrak{S}}\big|_j$ *and* $\rho^{\mathfrak{C}}\big|_j$.

Proof. See Appendix A.4. □

Since the relative wealth process is multiplicative, the geometric expected value of growth rates ultimately determines the relevant long-run outcome.

[6] With an abuse of notation, we denote by $i \in \{\mathfrak{C}, \mathfrak{S}\}$ the fixed point in which group i itself dominates.

[7] For random dynamical systems, a fixed point (either deterministic or random) is asymptotically stable when, for $\mathbb{P}$-almost all sequences of random shocks $\{x_t,\ t \in \mathbb{N}\}$, the path of states generated by the composition of maps $\tilde{\mathcal{F}}^{\mathfrak{S}}_{x_{t-1},x_t} \circ \cdots \circ \tilde{\mathcal{F}}^{\mathfrak{S}}_{x_0,x_1}$ converges to it, provided the initial state is chosen close enough.

2.4.1. Sufficient conditions for no endogenous pass-through

Here we investigate the conditions that prompt the constant group to dominate in the long-run, and consequently the stochastic group, together with their random component of aggregate demand, to vanish. By applying Lemma 2.1, the following Proposition provides sufficient conditions for the local asymptotic stability of fixed point $\mathfrak{C}$.

Proposition 2.4. *If parameters $\overline{x}$, x^u, x^d, g, π^u, and π^d are such that the following condition holds*

$$\left(1+g\frac{x^u}{\overline{x}}\right)^{\frac{\pi^u}{\pi^u+\pi^d}}\left(1+g\frac{x^d}{\overline{x}}\right)^{\frac{\pi^d}{\pi^u+\pi^d}} = \rho^{\mathfrak{S}}\big|_{\mathfrak{C}} < \rho^{\mathfrak{C}}\big|_{\mathfrak{C}} = 1+g, \tag{2.4.3}$$

then fixed point $\mathfrak{C}$ is locally asymptotically stable, whereas if $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}} > \rho^{\mathfrak{C}}\big|_{\mathfrak{C}}$ then $\mathfrak{C}$ is unstable.

Proof. See Appendix A.5. □

By restricting to a local analysis, we are evaluating the conditions for the stochastic group to survive when returns are determined by constant traders. In this case the stochastic traders hold an infinitesimal fraction of wealth and demand shocks are perfectly absorbed by the market. In other words, the pass-through is absent. The same happens for price increases that follow the recovery of the stochastic traders' positions. No such price upsurge occurs because the fraction of stochastic traders triggering them is negligible. Under these conditions, if $\overline{x} \geq x^u$ then the constant traders enjoy a higher growth for all realisations of the shock. Inequality (2.4.3) is always satisfied and the fixed point is locally asymptotically stable. Conversely, if $\overline{x} \leq x^d$ then the constant traders experience a lower growth with probability one, and the fixed point is unstable. In a locally asymptotically stable deterministic fixed point the survivor group must invest a higher share of wealth in the risky asset with respect to the vanisher. Intuitively, were it not the case, for the risky security yields a higher return with respect to the bond, an arbitrarily small redistribution in favour of the vanisher would bring the system further and further away from the initial fixed point, implying the latter is unstable. The results presented here are the stochastic generalisation of those already present in the literature (cf. Anufriev and Dindo, 2010, Proposition 5.2) regarding so-called 'fundamentalist' rules.

The interesting case is when $x^d < \overline{x} < x^u$ holds. Whenever $x_{-1} = x^u$ the wealth share φ of the stochastic traders increases; conversely, whenever $x_{-1} =$

x^d, φ shrinks. *Cœteris paribus*, the larger π^d, or the smaller π^u, the higher the likeliness that condition (2.4.3) is satisfied. In general, it is not possible to obtain a closed form solution of the inequality in (2.4.3) in terms of $\overline{x}$ since $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}$ is not algebraic. We are only able to analytically work out the following Special case.

Special case. $\pi^u = \pi^d$
When the transition probabilities coincide, implying that the average duration of high investment equals that of low investment, inequality (2.4.3) *translates into the following algebraic condition in terms of the constant trader's portfolio rule:*

$$\begin{aligned}\overline{x} &> \frac{x^u + x^d}{2} - \frac{(1+g)(x^u + x^d) - \sqrt{4gx^u x^d(2+g) + \left(x^u + x^d\right)^2}}{2(2+g)} \\ &= \frac{x^u + x^d}{2} - h(g),\end{aligned} \tag{2.4.4}$$

$$\text{with} \qquad h(g) \geq 0, \qquad h(0) = 0, \qquad h'(g) > 0.$$

The last equality shows that the RHS of inequality (2.4.4) *converges from below to the simple arithmetic average between x^u and x^d, as g approaches zero.*

This result implies that there would indeed be loss of generality should one adopt the so-called deterministic skeleton approach and substitute the $\{x_t\}$ process with its expected value.[8] The noise component incorporated into the stochastic traders' portfolio brings a detrimental effect to their own survival, at least in the $\pi^u = \pi^d$ case, since the constant traders can dominate even investing a portfolio fraction $\overline{x}$ that is strictly lower than the average of the stochastic ones'. The following Proposition provides a sufficient condition for the dominance of the constant trader even when $\pi^u \neq \pi^d$.

Proposition 2.5. *If $\overline{x} \geq \mathbb{E}[x_t]$ then fixed point $\mathfrak{C}$ is locally asymptotically stable.*

Proof. See Appendix A.6. □

Importantly, for demand shocks to pass through onto market prices, the stochastic traders must be, on average, more aggressive than the rest of the traders.

Finally, it is possible to show that there always exists exactly one value x' of $\overline{x}$ where a bifurcation occurs. Moreover, such value is bounded from below by x^d and from above by x^u.

[8] The expected value is computed according to the invariant distribution of the transition matrix in eq. (2.2.13) that in this special case reads $\left[\frac{1}{2}, \frac{1}{2}\right]$.

Proposition 2.6. $\exists!\ x' \in (0,1)$ *such that fixed point* $\mathfrak{C}$ *is locally asymptotically stable* $\forall\ \overline{x} > x'$ *and is unstable* $\forall\ \overline{x} < x'$*. Moreover, the following relation holds:*

$$x^d \ < \ x' \ < \ \mathbb{E}[x_t] < \ x^u. \tag{2.4.5}$$

Proof. See Appendix A.7. □

2.4.2. Sufficient conditions for maximal pass-through

Here we investigate the conditions that prompt the stochastic group to dominate, so that demand shocks, in the long run, entirely pass through to market prices. Following Corollary 2.1, we derive the condition for local asymptotic stability of fixed point $\mathfrak{S}$ in the following Proposition.

Proposition 2.7. *If parameters* $\overline{x}$, x^u, x^d, g, π^u, *and* π^d *are such that the following condition holds*

$$\left[1+\frac{g\overline{x}}{x^u}\right]^{\frac{\pi^u(1-\pi^d)}{\pi^u+\pi^d}} \cdot \left[1+g\frac{\overline{x}}{x^u}+\frac{(x^u-\overline{x})(x^u-x^d)}{x^u(1-x^u)}\right]^{\frac{\pi^u\pi^d}{\pi^u+\pi^d}}$$

$$\cdot \left[1+g\frac{\overline{x}}{x^d}+\frac{(\overline{x}-x^d)(x^u-x^d)}{x^d(1-x^d)}\right]^{\frac{\pi^u\pi^d}{\pi^u+\pi^d}} \cdot \left[1+\frac{g\overline{x}}{x^d}\right]^{\frac{\pi^d(1-\pi^u)}{\pi^u+\pi^d}} \ = \rho^{\mathfrak{C}}\big|_{\mathfrak{S}} < \rho^{\mathfrak{S}}\big|_{\mathfrak{S}} \ = 1+g, \tag{2.4.6}$$

then fixed point $\mathfrak{S}$ *is locally asymptotically stable, whereas if* $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}} \ > \ \rho^{\mathfrak{S}}\big|_{\mathfrak{S}}$ *then* $\mathfrak{S}$ *is unstable.*

Proof. See Appendix A.8. □

Similar to the previous case regarding fixed point $\mathfrak{C}$, the eigenvalue $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}$ is non-algebraic and therefore a closed form solution in terms of of $\overline{x}$ is in general unfeasible. The following Proposition derives sufficient conditions regarding stability and instability of fixed point $\mathfrak{S}$ in terms of the constant group portfolio rule.

Proposition 2.8. *Given* π^u, π^d, x^u *and* x^d*, a sufficient condition for local asymptotic stability of fixed point* $\mathfrak{S}$ *is*

$$\overline{x} \ \leq \ x^d \qquad \wedge \qquad g \ > \ \tilde{g} \ := \ \frac{x^u-x^d}{1-x^u}. \tag{2.4.7}$$

Conversely, a sufficient condition for instability of fixed point $\mathfrak{S}$ *is*

$$\overline{x} \ \geq \ x^u \tag{2.4.8}$$

Proof. See Appendix A.9. □

The first condition of Proposition 2.8 states that the stochastic group (locally) dominates, provided the constant group is always less exposed to the risky asset and the dividend growth rate is large enough. The value of g plays a role because when it is low the stochastic group doesn't earn enough during high investment to compensate for the losses due to negative demand shocks. We shall come back to this point in the next section.

The following Conjecture, based on a numerical investigation of the parameters space, is the analogous of Proposition 2.6 regarding fixed point $\mathfrak{S}$.

Conjecture 2.1. $\exists!\ x'' \in (0,1)$ *such that fixed point* $\mathfrak{S}$ *is locally asymptotically stable* $\forall\ \overline{x} < x''$ *and is unstable* $\forall\ \overline{x} > x''$.

Hint. See Appendix A.10.

2.4.3. Sufficient conditions for endogenous pass-through

So far we have investigated the conditions under which either no pass-through (Section 2.4.1) or maximal pass-through (Section 2.4.2) occurs asymptotically. In the present section we prove the existence of a non-degenerate region of the parameters space within which the pass-through is endogenous. Given that both groups of traders survive, their relative wealth dynamics is a mean reverting process, and the exact extent of the pass-through depends on the share of aggregate wealth detained by traders which are vulnerable to demand shocks. The following Proposition provides a sufficient condition for the emergence of endogenous pass-through.

Proposition 2.9. $\exists\ \hat{g} \in (0, \tilde{g})$ *such that* $\forall g < \hat{g}$ *the following relation holds:*

$$x'' < x^d < x'. \tag{2.4.9}$$

Proof. See Appendix A.11. □

While Proposition 2.9 only provides a sufficient condition on the dividend growth rate g that allows for the emergence of long-run heterogeneity, a numerical inspection of eqs. (2.4.3) and (2.4.6) validates the following Conjecture.

Conjecture 2.2. *The following relation holds* $\forall\ g > 0$*:*

$$x'' < x'. \tag{2.4.10}$$

Hint. See Appendix A.12.

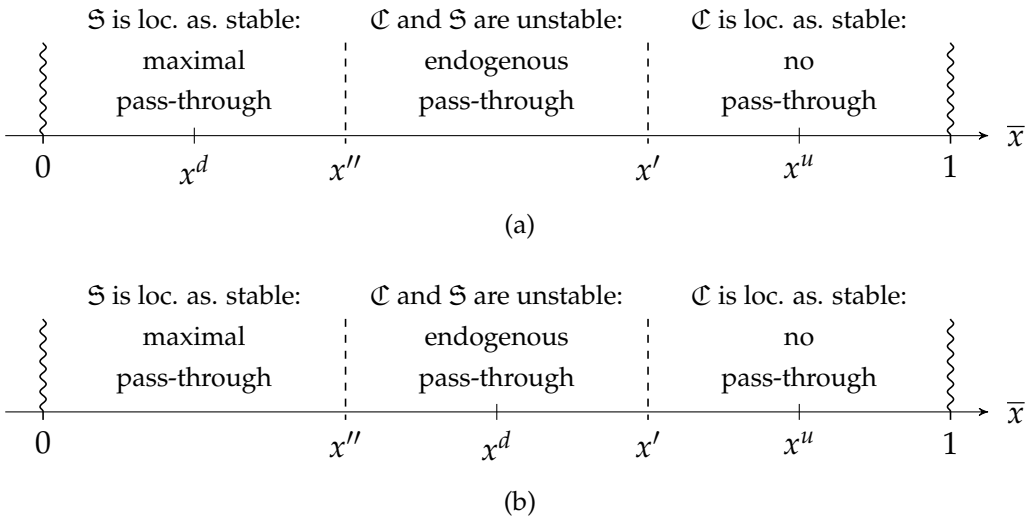

Figure 2.3.: Stability regions of fixed points $\mathfrak{C}$ and $\mathfrak{S}$ and extent of pass-through as a function of $\overline{x}$. (a): Regime with $x^d < x''$. (b): Regime with $x^d > x''$.

Fig. 2.3 portrays the relevant findings about the emergence of pass-through for all possible long-run outcomes. For $\overline{x} < x''$ there is maximal pass-through; for $\overline{x} \in (x'', x')$ the system exhibits endogenous pass-through; finally, for $\overline{x} > x'$ there is no pass-through. The position of the threshold value x'' may be either to the left, panel (a), or to the right of x^d, panel (b). A crucial parameter for the occurrence of either regime is the dividend growth rate. Following Proposition 2.8, if $g > \tilde{g}$ then for all $\overline{x} \leq x^d$ the stochastic trader dominates (locally) and the system is in the regime of panel (a). Long-run heterogeneity, and thus endogenous pass-through, can only occur when $\overline{x} > x'' > x^d$. As we shall see in the simulations in the next section, this regime is characterised by high growth (respectively, decline) of the stochastic group when it acts relatively more (respectively, less) aggressively with respect of the constant group. Losses due to buying-high and selling-low are relatively less severe. Following Proposition 2.9, if $g < \hat{g}$ then $x'' < x^d$ so that the regime of panel (b) occurs. Remarkably, when this is the case, the pass-through can be endogenous both when $\overline{x} > x^d$ but also when $\overline{x} \leq x^d$. As we shall show in the next section, a low dividend growth rate implies that, along the equilibrium paths of the regime of panel (b), losses during negative demand shocks are relatively more important than gains during periods of (locally) steady demand. For this reason, even with a strictly lower demand $\overline{x} < x^d$ the constant group is able to invade, albeit not to dominate, the stochastic group. Although we do not provide an explicit statement for

Description	Variable	Value
dividend rate of growth	g	0.05
stochastic investment up	x^u	0.7
stochastic investment down	x^d	0.3
probability down when up	π^d	0.01
probability up when down	π^u	0.01
initial wealth share	φ_0	0.5
initial return	r_0	null
initial yield	e_0	0.01

Table 2.1.: Parameters and initial conditions.

the cases $\hat{g} < g < \tilde{g}$, Conjecture 2.2 reveals that for any possible growth rate g the system is either in the regime of panel (a) or in the regime of panel (b).

2.5. Simulations and sensitivity analysis

In the previous Section we have outlined a pass-through taxonomy of all the long-run outcomes that system (2.2.15) admits. In this section, we explore the various possibilities by means of simulations. We also investigate the main determinants of the endogenous pass-through scenario in terms of the underlying parameters of the economy. First, it is useful to take a glimpse at how the dynamics in the latter case compares with the dynamics of the system featuring an asymptotically stable fixed point, either deterministic or random. Fig. 2.4 shows the portfolio (upper panel) and dividend yield dynamics (lower panel) in three typical simulation runs, one for each of the long-run scenarios depicted in Fig. 2.3. Once the relevant parameters of the model are set, the two threshold values x' and x'' can be computed numerically. In particular, for the parametrisation we use, reported in Table 2.1, they read $x' \approx 0.498$ and $x'' \approx 0.395$. Following Fig. 2.3, it is then sufficient to vary $\overline{x}$ to obtain a scenario of maximal pass-through (for $\overline{x} < x''$), one with endogenous pass-through (for $x'' < \overline{x} < x'$), or one in which (asymptotically) there is no pass-through whatsoever (for $\overline{x} > x'$). The upper panel of Fig. 2.4 shows a typical realisation of the Markov process $\{x_t\}$ (see Assumption 2.2) driving the stochastic group's investment, together with three horizontal lines representing the portfolio fraction of the constant group, one for each (in)stability scenario. For $\overline{x} = 0.55 > x'$, fixed point $\mathfrak{C}$ is locally asymptotically stable and the corresponding dividend yield dynamics (red)

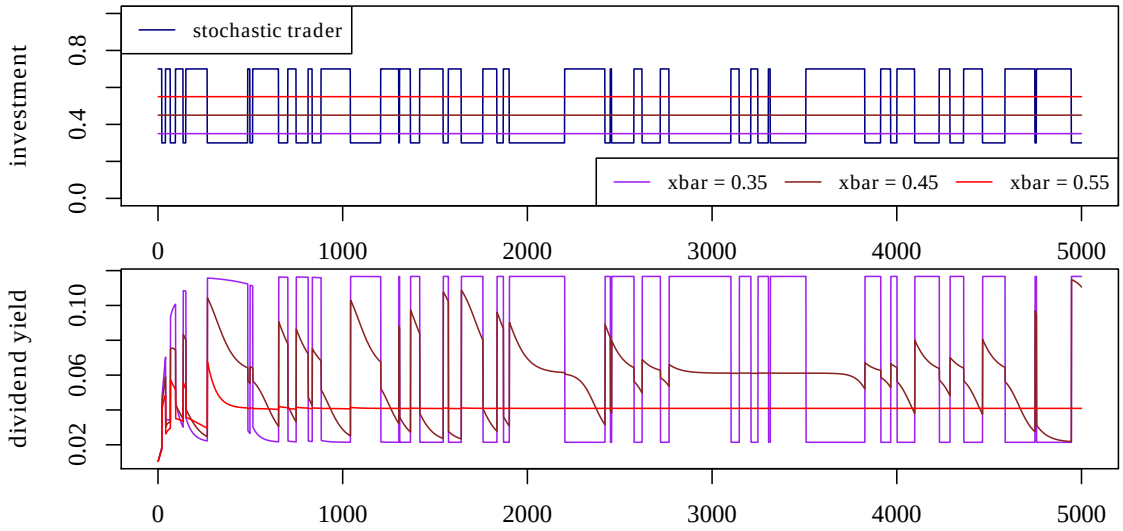

Figure 2.4.: Typical simulation runs within the first regime (cf. Fig. 2.3(a)) with either $\mathfrak{C}$ locally asymptotically stable (red), $\mathfrak{S}$ locally asymptotically stable (purple), or both $\mathfrak{C}$ and $\mathfrak{S}$ unstable (brown). Initialisation as of Table 2.1.

converges with oscillations to the constant value $e^* = g\frac{1-\overline{x}}{\overline{x}} \approx 0.041$, found in eq. (2.4.1). For $\overline{x} = 0.35 < x''$, fixed point $\mathfrak{S}$ is locally asymptotically stable and the corresponding dividend yield dynamics (purple) converges to the 4-state Markov process outlined in Corollary 2.1 and pictured in Fig. 2.2(a). Note that the values of the dividend yield are pairwise identical within the four states as found in eq. (2.3.9), and read $e^u = g\frac{1-x^u}{x^u} \approx 0.021$ and $e^d = g\frac{1-x^d}{x^d} \approx 0.117$, respectively. This mechanism can be more easily grasped in Fig. 2.2(b): regardless of the initial value e_0 of the yield, after a transient phase the dynamics converges to the boundary of the square with vertices (e^u, e^d), (e^u, e^u), (e^d, e^u) and (e^d, e^d), and eventually settles upon these vertices themselves, the jumps therein being triggered by the stochastic investment process $\{x_t\}$ with a one period lag. Finally, for $\overline{x} = 0.45 \in (x'', x')$, both admissible fixed points $\mathfrak{C}$ and $\mathfrak{S}$ are unstable and, since the constant group doesn't vanish, after an initial transient phase the dividend yield dynamics (brown) persistently fluctuates over a support strictly bounded by e^u from below and by $\max(e^d, e^*)$ from above (note that in this scenario $e^* \approx 0.061$). As will become clear shortly (this very scenario, also pictured in Fig. 2.5, is further analysed in the following Section), the reason why in some periods (e.g. between 1000 and 2000) the dynamics is richer than in others (e.g. around 3000) can be traced to the relative size of the traders in terms of wealth, and consequently in terms of market impact.

2.5.1. Endogenous pass-through

In this section we concentrate on those parametrisations that lead to endogenous pass-through. Following Fig. 2.3, we distinguish two main regimes. In the first, that typically occurs for high enough dividend growth rates, $x'' > x^d$, meaning that the constant traders need to invest $\overline{x} > x^d$ in order to invade, let alone dominate, the stochastic group and thus mitigate the extent of the pass-through. Conversely, in the second regime, that occurs only for low enough dividend growth rates, it holds $x'' < x^d$ and the constant traders are able to invade, although not to dominate, the stochastic group even when they invest a strictly less fraction of their wealth in the risky asset, $x'' < \overline{x} < x^d$. Albeit the long-run outcome of the two regimes is analytically analogous, the economic interpretation of the inherent trade-off is quite different.

Fig. 2.5 portrays a typical simulation run within the first regime (cf. Fig. 2.3(a)). The model is parametrised according to Table 2.1 and g is high enough to ensure that $x'' > x^d$ (the relevant threshold under this parametrisation is $\hat{g} \approx 0.010$). As found in the previous Section, the local stability thresholds are $x' \approx 0.498$ and $x'' \approx 0.395$. To ensure a situation of long-run heterogeneity, the value of $\overline{x}$ is selected inside the interval (x'', x').

The first panel depicts the portfolio fraction $\overline{x}$ of the constant group (green) and a realisation of the stochastic group portfolio process $\{x_t\}$ (blue). The second panel reports the dynamics of the wealth share φ_t of the stochastic group. At the beginning of the simulation run, the overall endowment is evenly split between the two groups. As trading unfolds, wide and sharp fluctuations in the distribution of aggregate wealth appear and persist indefinitely, as the system never converges towards a fixed point. We find that the value of the growth rate of the dividend is largely responsible for this wild dynamics. In general, the higher g, the wider the fluctuations of φ_t. The pattern of fluctuations of both dividend yields (third panel) and returns (fourth panel) is tightly coupled with the wealth share dynamics.[9] When φ_t is large, both dividend yields e_t and returns r_t are highly volatile; conversely, when φ_t is small, both e_t and r_t are relatively stable. Intuitively, when the stochastic group controls most of the wealth in the economy, it has a great impact on the market portfolio, and eventually on the clearing price. When this is the case, the portfolio shifts implied by the Markov

[9] The same is true also for the stationary component of the risky asset prices. Risky asset prices are not reported here but their dynamics can be inferred (and look indeed similar, up to the sign of their fluctuations) by the dividend yield dynamics, since the latter is growing at a constant rate of $1 + g$.

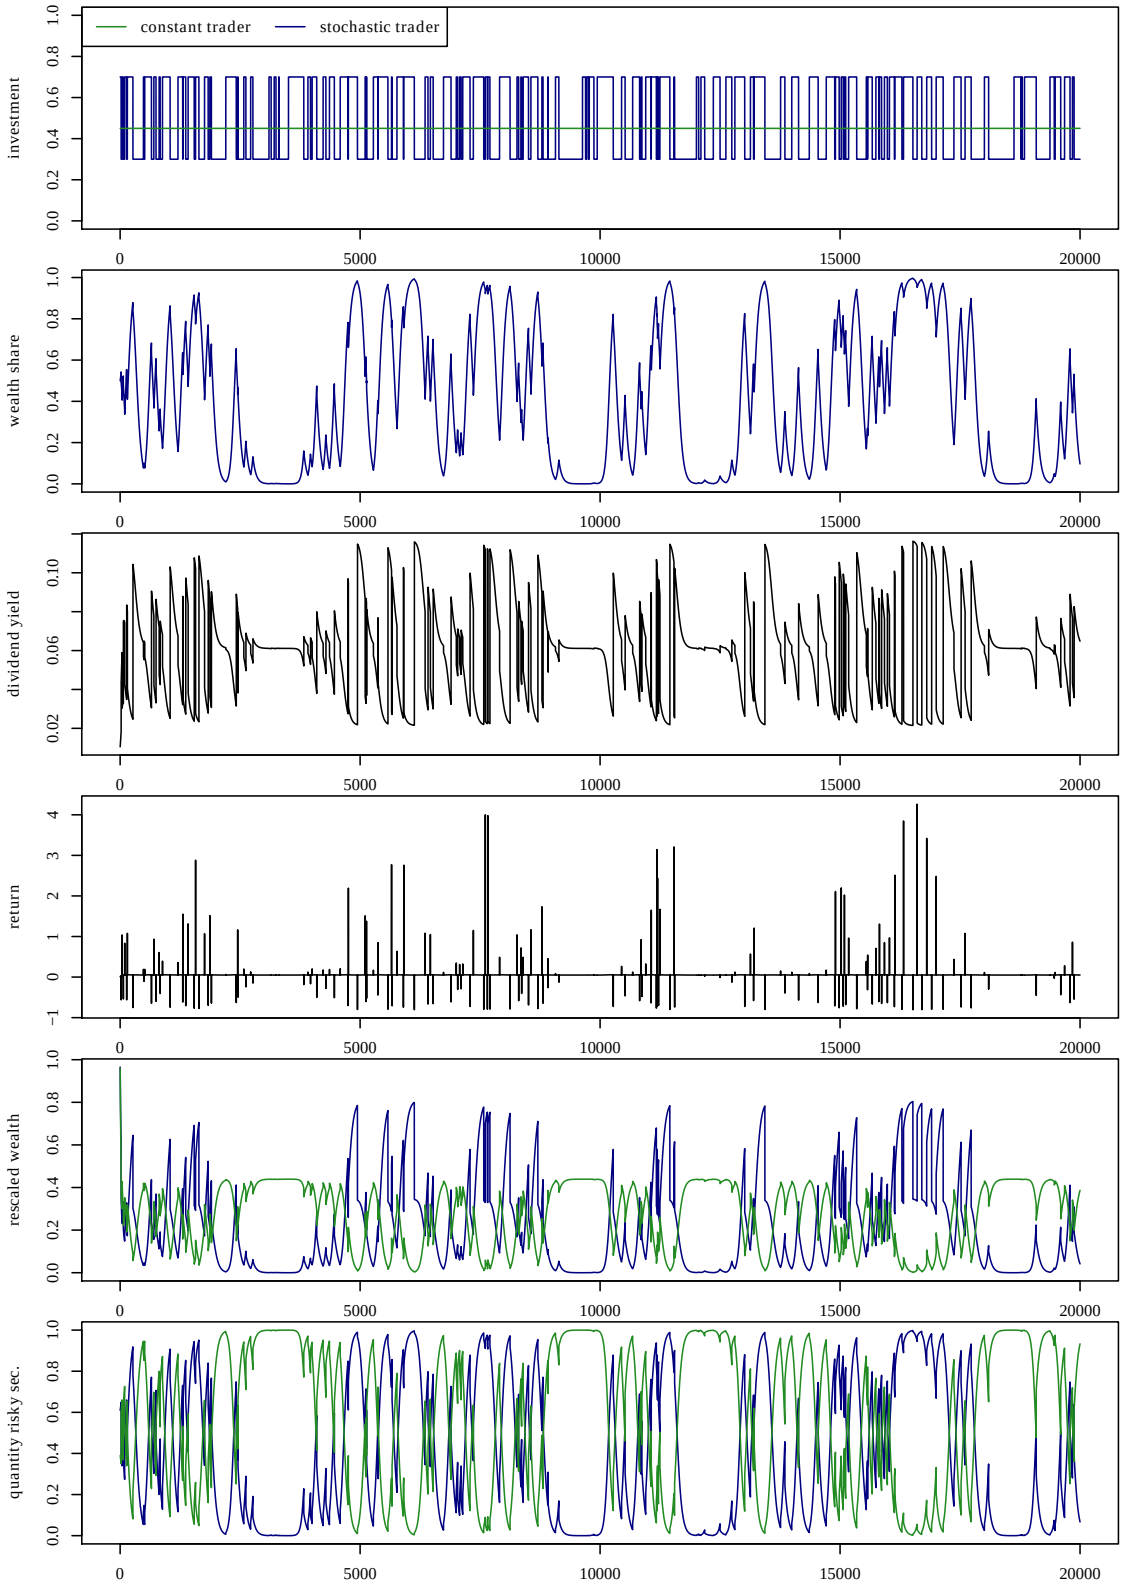

Figure 2.5.: Typical simulation run within the first regime (cf. Fig. 2.3(a)). Initialisation as of Table 2.1 and $\overline{x} = 0.45$. Green and blue series refer to the constant and stochastic group variables, respectively.

process $\{x_t\}$ are to a large extent incorporated into the prevailing price and the pass-through is substantial. This pattern closely resembles the famous volatility clustering stylised fact, i.e. the observation that "large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes" (Mandelbrot, 1963). A similar argument goes for the dynamics of traders' level of wealth (fifth panel)[10] and the actual exchanged quantity of the risky asset (sixth panel). When the stochastic group is predominantly influential, the system locally resembles fixed point $\mathfrak{S}$ where prices, and hence the market value of the portfolio, fluctuate according to the aforementioned random fixed point (cf. Fig. 2.2(a) and Corollary 2.1) while a null quantity of the asset is actually exchanged.

As representatives of the second regime, that is when $x'' < x^d$, we show the market dynamics in two distinct cases.[11] Fig. 2.6 pictures the same simulation of Fig. 2.5 except for a lower dividend growth rate, $g = 0.005 < \hat{g}$. Fig. 2.7 features both this new value of g and a lower investment fraction of the constant traders, now set to $\overline{x} = 0.25 < x^d$.

In both cases, fluctuations in the wealth share (second panel) in the absence of shock, and thus only due to the stream of dividend payments, appear dampened compared to Fig. 2.5. Intuitively, g has an influence on the speed of wealth adjustment for the group holding the greatest fraction of wealth in the risky security. In both these scenarios, the wealth share of each group remains, on average, relatively stable. In Fig. 2.6 the constant traders have the highest wealth share and thus the level of endogenous pass-through is smaller than in the other cases, as can be seen by comparing the dividend yield dynamics (third panel). In Fig. 2.7 the stochastic traders have the highest wealth share, leading to a larger pass-through. Leaving aside the size of fluctuations, the striking feature of this second regime with a lower dividend growth rate is that market volatility is more persistent.

A deeper analysis of the wealth dynamics in response to a local realisation of the Markov process $\{x_t\}$ is facilitated in Fig. 2.8, where we provide a close-up of all the three simulations presented earlier covering a shorter time span for visibility purposes. In all three scenarios, as long as $x_{t-1} = x_t = x^u$, the wealth of the stochastic group grows faster than that of the constant group, since $x^u > \overline{x}$; as a result, their wealth share increases. The speed of increment depends on the

[10] To get a more appreciable picture, the wealth series have been discounted by their non-stationary component $1 + g$.

[11] For the sake of comparability, all the simulations that we show throughout the chapter are driven by the same random seed.

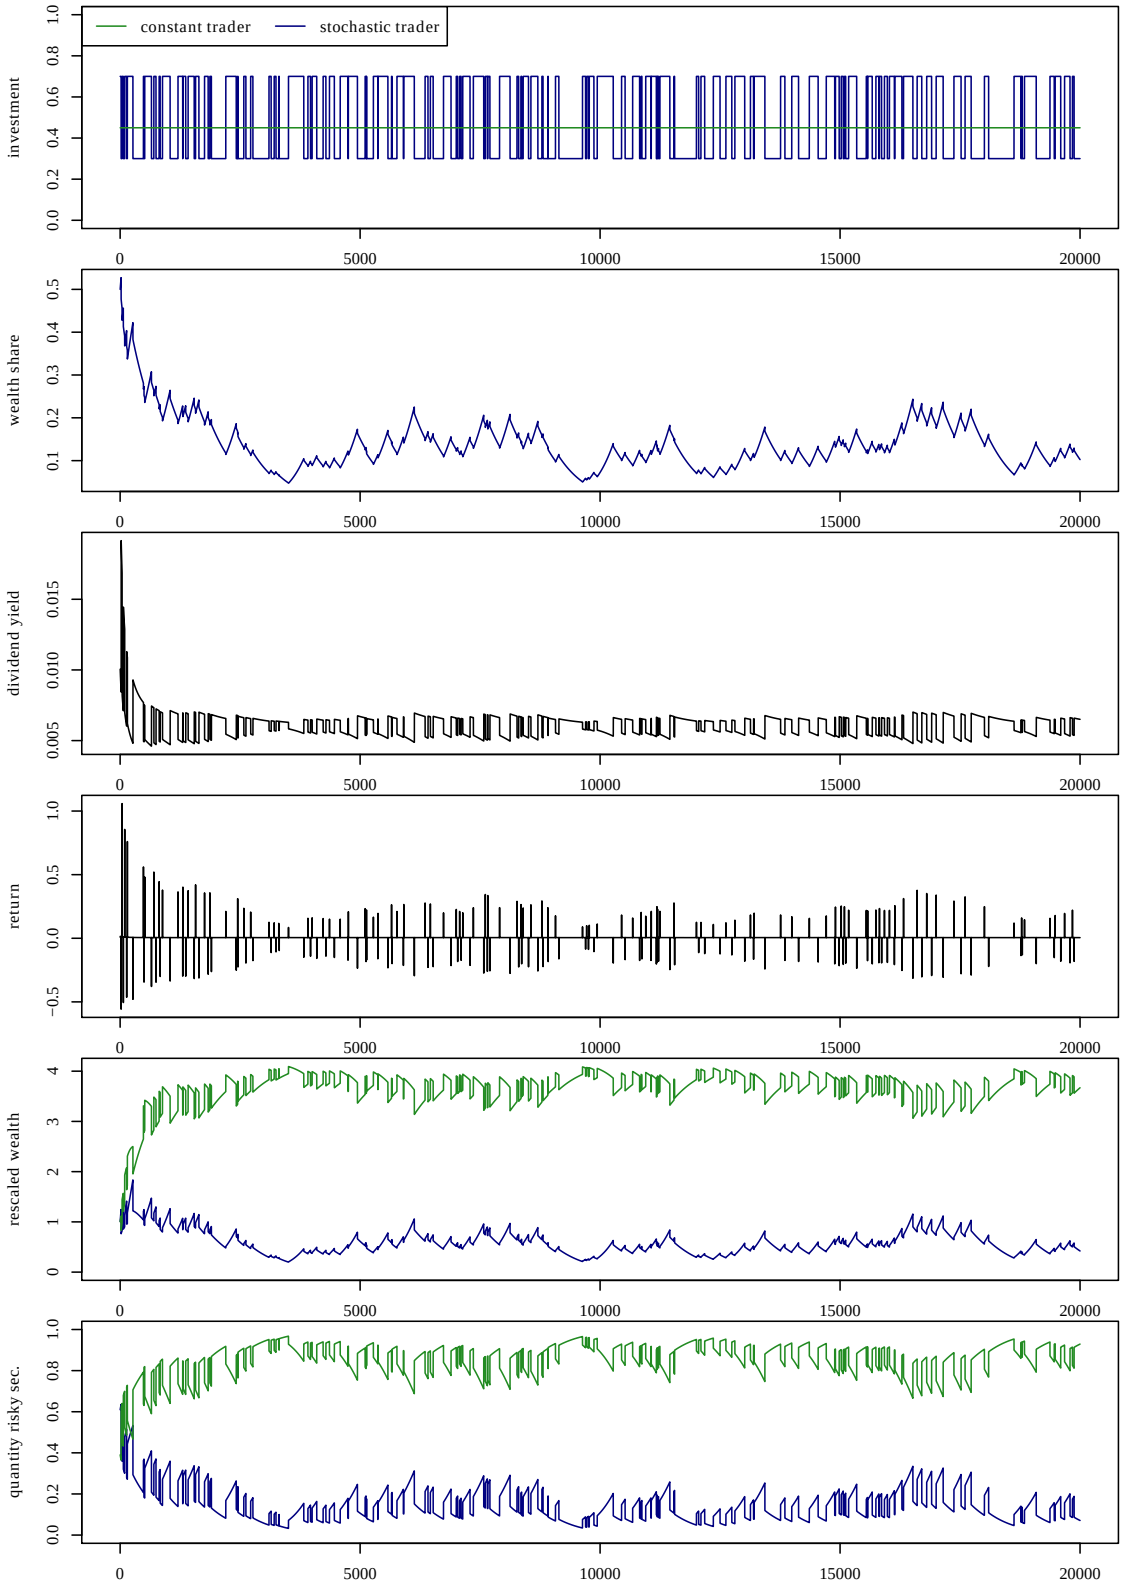

Figure 2.6.: Typical simulation run within the second regime (cf. Fig. 2.3(b)). Initialisation as of Table 2.1 except $g = 0.005$ and $\overline{x} = 0.45$. Green and blue series refer to the constant and stochastic group variables, respectively.

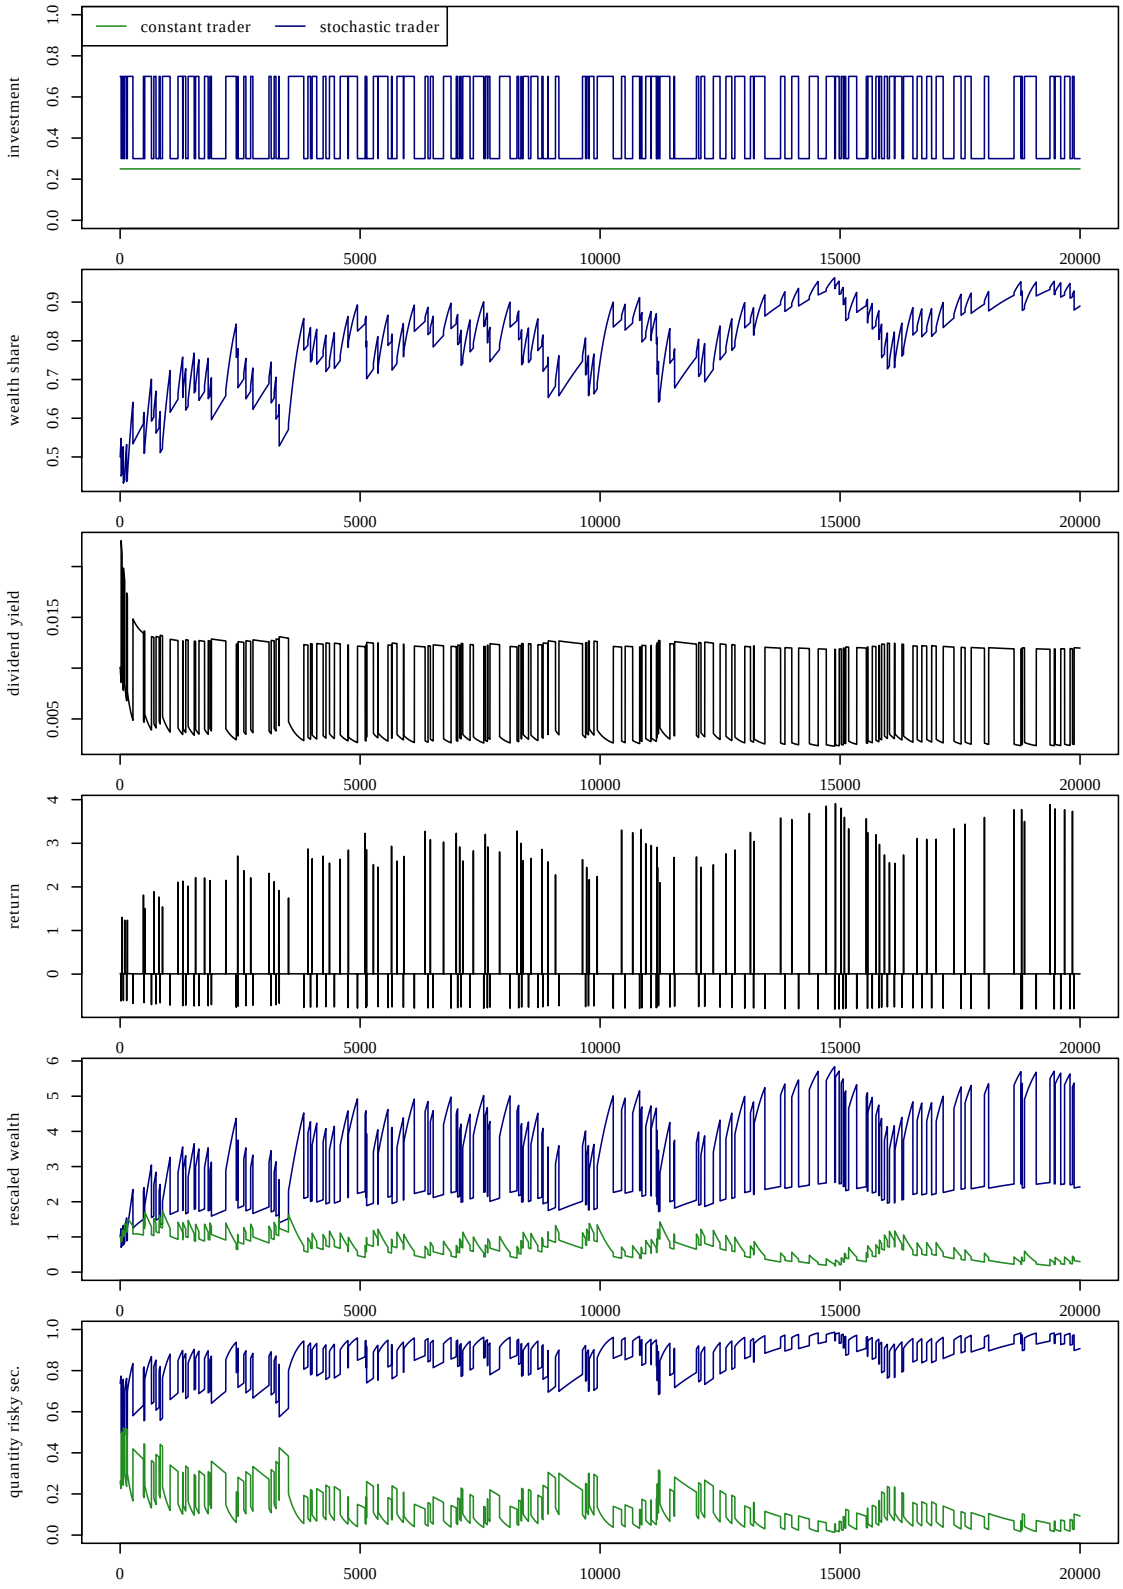

Figure 2.7.: Typical simulation run within the second regime (cf. Fig. 2.3(b)). Initialisation as of Table 2.1 except $g = 0.005$ and $\overline{x} = 0.25$. Green and blue series refer to the constant and stochastic group variables, respectively.

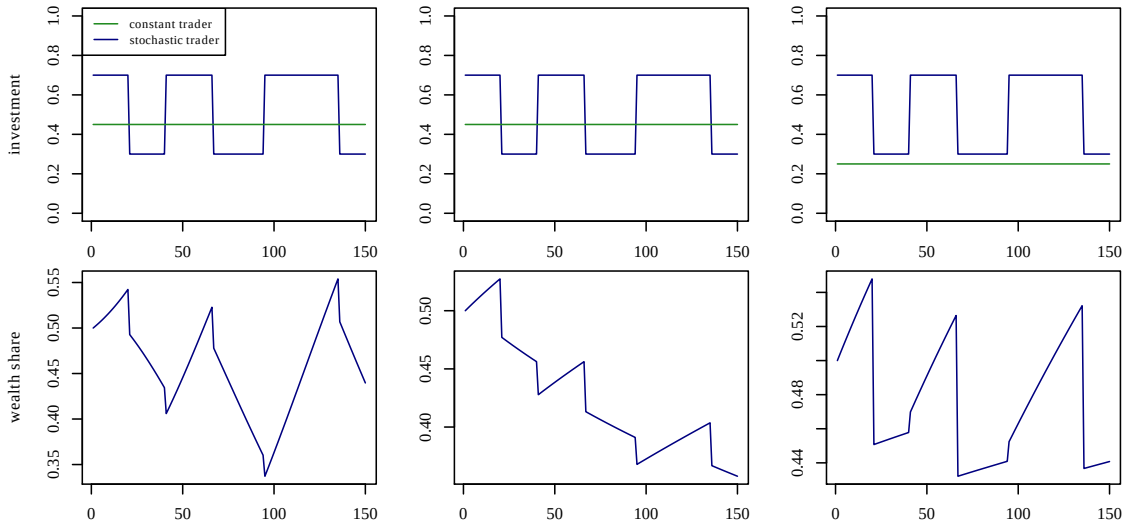

Figure 2.8.: Close-up of simulations in Fig. 2.5 (left), Fig. 2.6 (centre), Fig. 2.7 (right).

dividend growth rate g and it is higher in the first regime (left panel) than in the second regime (centre and right panels). When the stochastic portfolio suddenly shifts from high to low investment in the risky security, i.e. when $x_{t-1} = x^u$ and $x_t = x^d$, a downward pressure on the price is exerted and this consequently reduces the level of wealth for both groups, proportionally on their relative positions at time $t-1$. Again, since x^u is always greater than $\overline{x}$, the decrease hits the stochastic group to a larger extent, and in both cases the wealth share plummets accordingly. During the phases in which $x_{t-1} = x_t = x^d$, instead, the stochastic group is worse off in the first regime and in the first case of the second regime, since $x^d < \overline{x}$, and better off in the second case of the second regime, where the opposite relation holds true. Hence, their wealth share shrinks in the left and centre panels and expands in the right one. The differences between the first and the second panel are not due to the relative position of $\overline{x}$ with respect to x^d but rather to the size of g. For high g (left panel), the wealth of stochastic traders decreases quite sharply compared to that during negative demand shocks. For low g the decrease is of comparable size. The relative position of $\overline{x}$ ad x^d is instead crucial when the stochastic portfolio suddenly shifts from low to high investment in the risky security, i.e. when $x_{t-1} = x^d$ and $x_t = x^u$. The ensuing increase favours the stochastic group in the second scenario of the second regime (i.e. when $x^d > \overline{x}$, right panel) but penalises it in both other cases (left and centre panel).

2.5.2. Sensitivity analysis

We are left with studying to what extent changes in the values of the underlying parameters influence the asymptotic dynamics of the economy. Starting from the usual parametrisation of Table 2.1 except $g = 0.005$, Fig. 2.9 shows the evolution of the threshold values x' and x'' in terms of the constant trader portfolio $\overline{x}$ for changes in g (top-left panel), in the position of x_t within the unit simplex with constant jumps $x^u - x^d$ (top-right), in the dispersion of x_t around a constant mean (bottom-left), and in the transition probability π, under the Special case $\pi = \pi^u = \pi^d$ (bottom-right). In each plot, the relative position of the two thresholds splits the Cartesian plane in three distinct areas, corresponding to the stable regions of both fixed points $\mathfrak{C}$ and $\mathfrak{S}$, and their jointly unstable region where long-run heterogeneity and endogenous pass-through occur. The coloured points correspond to the parametrisation chosen for the simulations pictured in Fig. 2.5 (red), Fig. 2.6 (purple), and Fig. 2.7 (brown). In the top-left panel, we show the joint effect of $\overline{x}$ and g. For low levels of g (i.e. to the left of the vertical dotted line $g = \hat{g}$) the system is in the second regime of Fig. 2.3 and endogenous pass-through can occur both when $\overline{x} > x^d$ and when $\overline{x} \leq x^d$ (the horizontal dotted line in the picture represents the locus $\overline{x} = x^d$). For higher levels of g (to the right of the vertical dotted line) endogenous pass-through can occur only for $\overline{x} > x^d$. In fact, g has a (slightly) negative effect on the threshold x' and a positive effect on x''; the first result has been already shown in the aforementioned Special case and generalised in Proposition 2.5. Therefore, the higher g, the narrower the interval of $\overline{x}$ for which there is long-run heterogeneity. For even higher g, this interval shrinks even further, and at the limit it collapses into one point:

$$\lim_{g \to +\infty} x'(\cdot) = \lim_{g \to +\infty} x''(\cdot) = \mathbb{E}^{\boldsymbol{\pi}}_{\mathcal{G}}[x_t], \tag{2.5.1}$$

where $\mathbb{E}^{\boldsymbol{\pi}}_{\mathcal{G}}[\cdot]$ denotes the geometric expected value with respect to the invariant distribution $\boldsymbol{\pi}$ of the Markov process $\{x_t\}$. The top right panel shows what happens when the support of the stochastic portfolio x_t shifts within the unit simplex, keeping the range, i.e. the extent of the demand shock $x^u - x^d$, constant. Clearly, the width of the long-run heterogeneity corridor has a non-linear relation with $\mathbb{E}[x_t]$. In particular, the difference $x' - x''$ is especially large at the extrema of the simplex. Note also that $x''(\cdot)$ is not monotone, and for sufficiently high values of $\mathbb{E}[x_t]$ it is decreasing. When x'' crosses the bisector (the dotted line, along which $\overline{x} = x^d$) from above, a non-degenerate region between the two curves appear. For any point in this region, the constant group is able

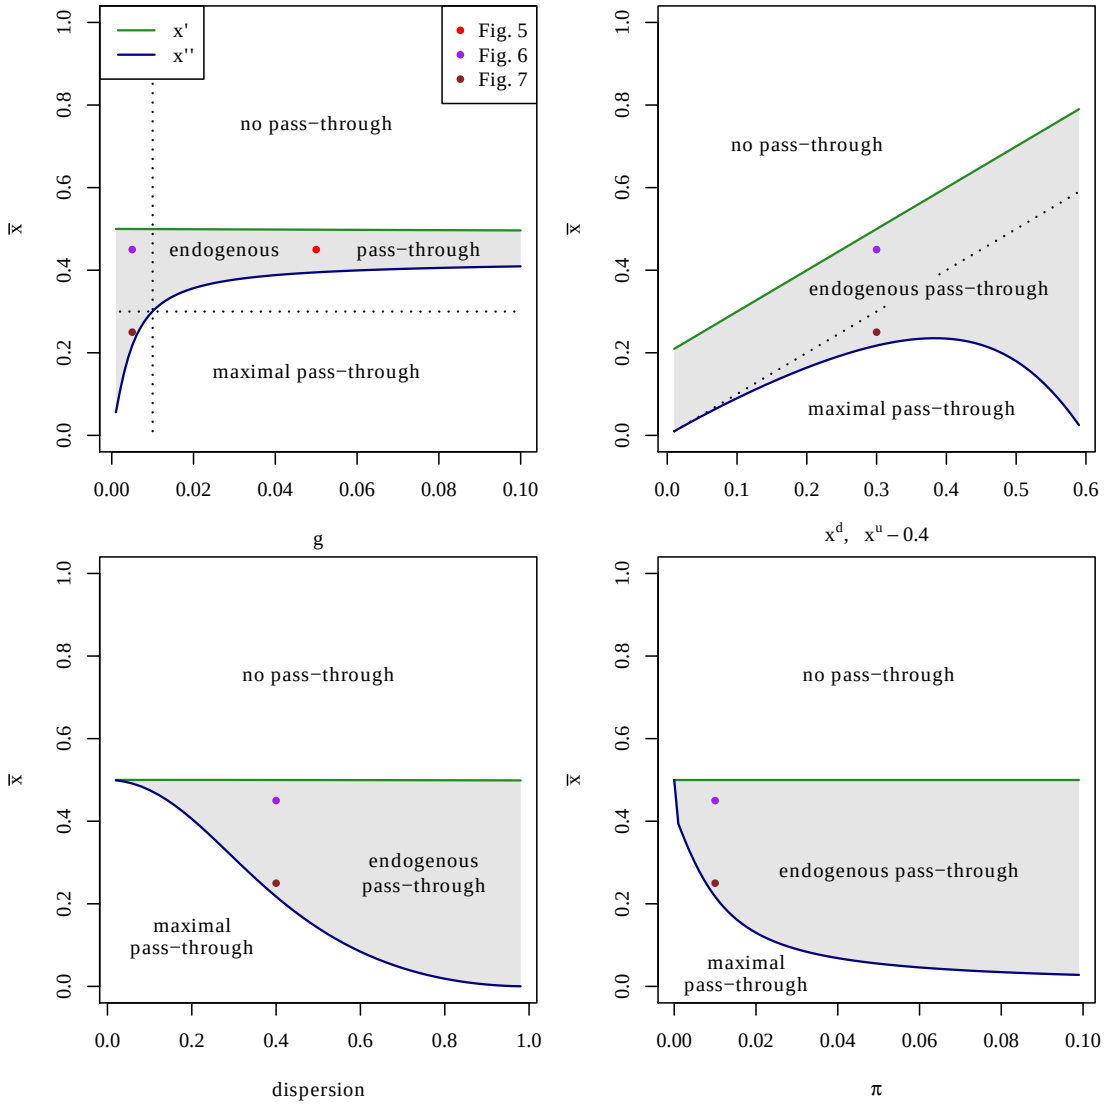

Figure 2.9.: Sensitivity analysis. Green and blue curves refer to the constant and stochastic group thresholds x' and x'', respectively. Coloured points correspond to simulations in Fig. 2.5 (red), Fig. 2.6 (purple), and Fig. 2.7 (brown). Initialisation as of Table 2.1 except for variables reported on the axes and $g = 0.005$.

to invade the stochastic group by investing a strictly less fraction of wealth in the risky asset; this is what we previously dubbed the 'second regime'. When we simulated the second scenario (Figs. 2.6 and 2.7) we opted for keeping x^u and x^d unchanged and reducing the dividend growth rate g that, following our discussion of the top-left panel, lowers the x'' threshold and enlarges the shaded area below the bisector line. As the top-right panel shows, however we could have kept both g and the jump $x^u - x^d$ unchanged, increased x^d close to 0.6, and eventually end up in a mathematically equivalent dynamics. In the bottom-left panel the thresholds are plotted against the symmetric range of $\{x_t\}$ around a constant $\mathbb{E}[x_t] = 0.5$. Both functions are decreasing but x'' is steeper, so that the interval of long-run heterogeneity widens as the dispersion increases. Intuitively, the more abrupt the change in the positions of the stochastic traders during switching phases, the more the effect of buying high and selling low penalises them and favours the constant group, which is able to dominate with a lower risky position $\overline{x}$. For similar reasons, the more frequent the switching phases of $\{x_t\}$, the more the stochastic traders are disadvantaged. This is highlighted in the bottom-right panel, where the range of $\{x_t\}$ is kept fixed, but the transition probability $\pi = \pi^u = \pi^d$ varies. Following Proposition 2.4, whenever $\pi^u = \pi^d$ the stability condition of fixed point $\mathfrak{C}$ does not depend on the value of π and therefore x' is a horizontal line. On the contrary, the stability of fixed point $\mathfrak{S}$ does depend on the specific value of π (see Proposition 2.7). In particular, the higher π, the more often demand shocks hit, so that losses induced by the buy-high and sell-low effect make the stochastic traders less likely to dominate.

2.6. Concluding remarks

We investigate asset prices and wealth-driven selection in a simple financial market where one of two groups of traders is subject to random demand shocks. When a positive (negative) shock hits, depending on the relative wealth of the two groups, the extra supply (demand) of risky asset is either fully absorbed by the market, or it is passed through onto market prices. We provide conditions on the economy parameters such that demand shocks are always fully absorbed in the long-run, when the stabilising constant agents dominate, as well as conditions under which demand shocks are entirely passed through, when the constant agents vanish. Moreover, there exist cases when both groups survive and demand shocks may or may not lead to sharp price movements. In this case we say that the pass-through is endogenous, its occurrence being coupled

with the traders' wealth dynamics. In order to derive these results we study the random dynamical system characterising the dynamics of relative wealth, risky asset returns, and dividend yield, as driven by an exogenous demand stochastic process. Our results are unobtainable from the study of the underlying deterministic skeleton alone. In particular, we show that simply taking into account the expected value of the stochastic portfolio is not sufficient to determine the long-run dynamics of the system.

Our model can be further extended in a number of directions. A first improvement would be to reconcile our framework with the typical investment strategies adopted within the HAMs literature, namely dependent on realised market outcomes, e.g. chartist and fundamentalist. Relatedly, it seems natural to extend our 2-state stochastic portfolio to a general n-state Markov process, with $n > 2$, in order to allow for more convincing (e.g. staircase) adjustments of holdings of the risky asset in response to a demand shock. This would indeed complicate the analysis since n possible levels of investment correspond to a 2^n-state random fixed point Markov process (see Corollary 2.1). In such a scenario, we believe that the role of trend-followers (and therein their size in terms of wealth and the magnitude of their portfolio response) might be crucial in spelling the patterns of shock propagation revealed by the empirical literature. In other words, trend-following behaviour could act as a microfoundation device for the spread of demand shocks that we assumed to be joint. On top of that, it would be appealing to further endogenise not just the response to an exogenous demand shock, but the emergence of the shocks themselves, for instance by relating them to the dynamics of the risky asset's fundamentals. This can in principle be done by explicitly taking into account the noise generated by the dividend's growth rate and connecting demand shocks to traders' expectations of it. Finally, an additional enhancement would be to allow for an arbitrary number of risky securities to be traded in the market, along the lines of the deterministic wealth selection model of Anufriev et al. (2012). This would enable the investigation of the spillover effect of shocks hitting the demand of one asset onto other assets.

This page intentionally left blank.

3

An agent-based model of intra-day financial markets dynamics[1]

3.1. Introduction

Huge improvements in information and communication technologies have substantially reduced the latency required to operate on financial markets in the last decades. This has fostered market activity at increasingly higher frequencies. Differently from traditional money managers, who generally hold their portfolio positions for a long period, ranging from a few days to even months, high-frequency traders aim at reaping profits from a large multitude of buy and sell operations that they execute *within* each trading day, rarely holding their positions overnight. These (very) short-term trading strategies have proved remarkably profitable even during periods of nearly unprecedented financial turmoil (see e.g. Aldridge, 2013). At the same time, the overall impact of these trading strategies on market dynamics is still unclear.[2] In addition, the increasing volumes of high-frequency traders in financial markets certainly impacts many of the stylised facts of intra-daily financial market dynamics. These statistical

[1] This chapter is a joint work with Mauro Napoletano (OFCE–Sciences Po, Sophia Antipolis, France). The manuscript is available as Staccioli, Jacopo and Mauro Napoletano (2018). 'An agent-based model of intraday financial markets dynamics'. LEM Working Papers series n. 12/2018. URL: http://www.lem.sssup.it/WPLem/2018-12.html. The authors acknowledge the financial support of European Union's Horizon 2020 grant No. 640772 – Project Dolfins.

[2] The trend of progressively shortening the time needed to collect real-time information and post a new order has been in place for many decades, starting with the introduction of high-speed telegraph service and later boosted by the availability of powerful computer systems. Nevertheless, a full and agreed understanding of the functioning, potential benefits, and disadvantages of high-frequency trading has yet to be reached (see also Aldridge, 2013; Jacob Leal et al., 2016; SEC, 2014).

properties are still begging for a sound theoretical framework (see Cont, 2011 for a discussion).

We propose a parsimonious agent-based model of a financial market that is able to *jointly* reproduce many of the empirically validated stylised facts. These include properties related to returns (leptokurtosis, absence of linear autocorrelation, volatility clustering), trading volumes (volume clustering, correlation between volume and volatility), and timing of trades (number of price changes, autocorrelation of durations between subsequent trades, heavy tail in the distribution of such durations, order-side clustering).

In the last few decades, the still flourishing literature on agent-based models (some of the milestones include Arthur et al., 1997; Levy et al., 1994; Lux, 1995, 1998; Lux and Marchesi, 2000) has proved invaluable for investigating and replicating the statistical properties of financial markets that are hardly reconcilable with the representative agent paradigm.[3] However, the vast majority of the proposed models typically focuses only on a subset of the whole ensemble of recognised stylised facts, and in particular on the facts that are time-scale invariant. These generally include properties related to rates of return, such as leptokurtosis, absence of linear autocorrelation, and volatility clustering. Other stylised facts concerning the timing of orders posting and trades execution are often neglected. This is partially due to the acknowledged difficulty of defining a reasonable mapping from the iterations of a computer simulation to proper calendar time (see e.g. Cioffi-Revilla, 2002). A notable exception is the work of Kluger and McBride (2011), who propose a model that replicates the intra-day U-shaped seasonality in market activity, i.e. the tendency of exchanged volumes to peak during the early morning just after market opening and late afternoon just short of market closing, leaving a trough around lunch-time.

To the best of our knowledge, no previous study has ever addressed the simultaneous emergence of all the stylised facts that shape financial dynamics, i.e. including those at the intra-day level. We therefore attempt at filling this gap, by proposing at the same time a methodological solution to the time mapping problem and by identifying the building blocks of the model which are responsible for the emergence of the solicited stylised facts.

Our model relies on three main ingredients. The first consists of a behavioural specification of traders, which is typical of many established models in the lit-

[3] Some of these models, often dubbed "heterogeneous agent models" are low-dimensional and mathematically tractable; others, are too complex to be investigated analytically and rely on extensive numerical simulations. See Hommes (2006) and LeBaron (2006) for a discussion. Our work aims at contributing to the latter strand of literature.

erature. Indeed, we assume that traders are of two types: fundamentalists and chartists. Fundamentalists only take into account the fundamental value of the security (which we shall assume constant across time and common knowledge), by buying the asset if it is undervalued and selling it if it is overvalued. Chartists instead rely on the recent history of price changes to set their orders, by extrapolating the underlying trend if they are followers or counteracting it if they are contrarians. This specification is justified by empirical surveys of financial practitioners' behaviour (see e.g. Frankel and Froot, 1990).

The second ingredient, which to our knowledge has never appeared in any previous contribution, amounts to a realistic scheduling of trading events. More precisely, we borrow the exact time structure of a trading day on a real financial market, namely the EURONEXT, and we design our simulations according to the sequence and durations of the different phases therein (see Euronext, 2017). The latter consists, in chronological order, of a morning order accumulation phase, an opening batch auction, a lengthy phase of real-time order matching according to a continuous double auction, a pre-closing order accumulation phase, and a closing batch auction. Imposing a strict and realistic schedule on the unfolding of events enables to devise a sound and plausible correspondence between simulation iterations (which we shall identify with seconds) and calendar time. Microstructure details about the central order book also comply with EURONEXT specifications.

The last ingredient of the model is an endogenous mechanism for traders participation. We assume that traders (of either type) are more willing to engage in trading whenever the price change (of either sign) realised in the immediate past is high enough. The intuition is that large realised (absolute) returns signal the possibility of reaping further profit in the future. Note that in the following we shall not impose any short-sale restriction. A similar scheme is devised in Ghoulmie et al. (2005), Aloud et al., 2013 and Jacob Leal et al. (2016). In spite of being extremely simplistic, we find that this activation mechanism proves crucial for matching our target stylised facts, specifically those related to the timing of trades execution.[4]

[4]Another conceivable ingredient, commonly adopted in financial models akin to ours, is a switching scheme between the fundamental and chartist strategies. In many contributions this is known to foster volatility clustering (see e.g. Kirman and Teyssière, 2002; Lux and Marchesi, 2000, and for a discussion Cont, 2007). Complementary simulation analyses that we carried out however indicate that this component is irrelevant in our setting, in which volatility clustering arises purely from the interaction of heterogeneous traders and is especially influenced by the trend-following momentum on behalf of chartists.

The next Section provides an overview of the various stylised facts that characterise high frequency financial dynamics. Section 3.3 describes in detail the various assumptions of our model. Section 3.4 reports the results of numerical simulations under different scenarios. Finally, Section 3.5 concludes.

3.2. Stylised facts

We begin by describing the various statistical properties that characterise the intra-day dynamics of many financial markets and that our model aims at reproducing. Some of these properties are recognised to apply across different time scales while others are intra-day specific and thus require a proper calendar setting to be analysed (see Cont, 2001, 2011, for a more detailed account). The former properties are mainly related to asset returns and have already been studied and replicated in a number of agent-based models lacking of a rigorous definition of calendar time. The latter instead require a more explicit architecture in terms of microstructure. In what follows, we present each stylised fact, distinguishing between time-invariant and intra-daily facts.

3.2.1. Time-invariant stylised facts of financial markets

SF1 – Leptokurtic returns The unconditional distribution of returns is characterised by a heavier tail with respect to the Gaussian distribution (Fama, 1965; Kon, 1984). The magnitude of excess kurtosis is typically inversely related to the time scale of analysis. This finding stands at sharp odds with the normality assumption adopted in a number of models, most notably the Black-Scholes formula for derivatives pricing (see e.g. Hull, 2017).

SF2 – Absence of autocorrelation of (raw) returns The time series of (raw) rates of return exhibits a statistically significant serial correlation for a very short amount of time, quickly decaying to zero afterwards. Intuitively, should there be more predictable autocorrelation structure, this information could be used to perform *'statistical arbitrage'* with positive profits (Mandelbrot, 1971).

SF3 – Volatility clustering While the linear autocorrelation of returns displays very little structure, the autocorrelation of non-linear functions such as the absolute value or the squared value of returns is usually positive and tends to decay at a much slower pace. Therefore, while the signs of future returns are not readily predictable, their magnitudes are, and tend to cluster in time, giving rise to

prolonged periods of low volatility followed by periods of high volatility (Andersen and Bollerslev, 1997; Mandelbrot, 1963). This clearly suggests that the series of returns is not independent.

SF4 – Leverage effect The leverage effect or asymmetric volatility (Black, 1976) captures the asymmetric tendency of volatility to be higher during price drops rather than during price surges. This translates into the negative correlation between price volatility – e.g. absolute returns – and the (raw) returns of the asset (Aït-Sahalia et al., 2013; Bollerslev et al., 2006; Bouchaud et al., 2001).

SF5 – Autocorrelation of volumes The quantities exchanged during successive trades display significant positive serial correlation (Campbell et al., 1993; Engle, 2000; Gallant et al., 1992). This is true across different time aggregation units and both for indices and individual stocks.

SF6 – Correlation between volumes and volatility Price variability and trading volumes display positive correlation (Foster, 1995; Tauchen and Pitts, 1983). The underlying idea is that the flow of information acts as a common determinant of both changes in prices and traded quantities.

3.2.2. Intra-daily stylised facts of financial markets

SF7 – Number of price changes per day In a cross-section perspective, the number of price-changing trades per day is clearly related to the degree of liquidity of the market and is typically linked to the capitalisation of the underlying security. Over time, moreover, there is a tendency of reduction in the time needed to execute a market order, fostering the submission of an increasingly larger number of orders, eventually leading to an increasing frequency of actual trades. Nowadays, for blue-chips in highly liquid markets and in the absence of 'disruptive' fundamental news, this number is often around 10,000, with a substantial degree of variance (Bonanno et al., 2000; McInish and Wood, 1991).

SF8 – Autocorrelation of durations between subsequent trades Within continuous double auctions, the actual timing of transactions is endogenous since a freshly submitted order might not find a compatible crossing order already stored in the book. Therefore, the time intervals between subsequent transactions is both random and tightly linked to the previous history of orders post-

ing. Empirically, these durations display positive autocorrelation – translating in clustered periods of frequent transactions followed by periods of sporadic transactions (Cont, 2011).

SF9 – Fat-tailed distribution of durations between subsequent trades The distribution of the durations defined in SF6 reveals a heavier tail with respect to an exponential distribution, that would be instead expected if traders submitted their orders in a non-correlated timely fashion (Raberto et al., 2002).

SF10 – Order-flow clustering The arrival of orders over time to the central order book is clustered with respect to the side of intended transaction: buy orders are more likely to follow previous buy orders, while sell orders are more likely to follow sell orders (Biais et al., 1995).

SF11 – U-shaped activity Market activity throughout the day displays a strong seasonality, with peaks of exchanged quantities in the early morning after market opening and in the late afternoon in the vicinity of market closing, and a relative more tranquil period in the hours around lunch-time (Jain and Joh, 1988; Lockwood and Linn, 1990).

In what follows, we aim at developing a simple and parsimonious model which is nonetheless capable of *jointly* reproducing all the aforementioned stylised facts, with the exception of the intra-day volume seasonality[5] (SF11), which is unobtainable by construction in our setting as will be clear later, and of the leverage effect (SF4), for which we believe a more complex behavioural specification is needed.

3.3. The model

Consider an order-driven financial market in which a single long-lived stock is traded by a population of heterogeneous agents. In line with the empirical literature on practitioners' behaviour in financial markets pioneered, among others, by Frankel and Froot (1990), Allen and Taylor (1990), Taylor and Allen (1992),

[5]Kluger and McBride (2011) provide an agent-based model that reproduces the U-shaped nature of intra-day volumes, although they don't discuss the whole ensemble of the stylised facts listed above.

and more recently by Menkhoff (2010), we consider two trading strategies: fundamentalist and chartist. A fundamentalist trader believes that the price of a security will quickly revert to its fundamental value; a chartist (or technical) trader, instead, believes that the future price of a security can be predicted using the trend of past realised market outcomes. Since we are interested in modelling short-term dynamics, we assume that the security pays no dividend and there is no "fundamental" news circulating during this time span. In this sense, besides an additive i.i.d. noise component incorporated in both strategies, the dynamics of prices and returns is endogenously determined by the interaction of the two strategies with the market microstructure, and observed volatility is actually excess volatility.

3.3.1. Timing and market setting

Since we are interested in describing the dynamics of a generic stock at a well-defined time scale – the intra-day level – we need to devise a mechanism that maps the iterations of our agent-based model to proper calendar time. This is a notoriously daunting and controversial task within the agent-based literature (see e.g. Cioffi-Revilla, 2002). To address this issue, we impose a strict global schedule to the sequence of events. In particular, we design our simulations to closely replicate the timing structure of an existing stock market, namely the EURONEXT. A typical trading day on the EURONEXT exchange unfolds as follows (Euronext, 2017):

at 7:15am the trading day starts with the pre-opening phase in which orders accumulate on the central order book without any transactions taking place;

at 9:00am a (batch) opening auction takes place, matching the orders submitted during the pre-opening phase and determining the opening price;

from 9:00am to 5:30pm the market operates according to a continuous double auction, and the introduction of a new order immediately generates one or more transactions if there are matching orders on the opposite side of the book. This phase is dubbed the 'main trading session';

at 5:30pm pre-closing phase starts, in which matching of orders is discontinued and, as in the pre-opening phase, orders accumulate with no transaction taking place;

at 5:35pm the closing auction takes place, matching the orders submitted during the pre-closing phase and determining the closing price of the day.

from 5:35pm to 5:40pm orders can be entered for execution at the closing price only. This phase is dubbed 'trading at last'.

With the exception of the trading in the last phase[6], we model our trading day according to the schedule above, and we identify a single iteration of the model with a calendar second. Hence, the pre-market phase corresponds to 6,300 time steps (1 hour and 45 minutes), the main trading session to 30,600 time steps (8½ hours), and the pre-closing phase to 300 time steps (5 minutes). A whole trading day consists of 37,200 simulation steps of our model.

At every time step some of the traders are activated (see the next section). They proceed in forming their expectations about the future performance of the security, and submit limit orders accordingly. When an order is submitted, it is either stored on the central order book or matched (if possible), depending on the current phase of the trading day. If matched, the order gives rise to one or more trades, the relevant quantities are exchanged, and a new price is disseminated. The central order book follows the usual price-time priority rule.

3.3.2. Traders' participation

We devise two alternative mechanisms for traders participation, one exogenous and one endogenous. In the first, a single randomly selected trader is activated at each time step. This activation scheme is similar to the one employed by Chiarella and Iori (2002). In the second, we follow Jacob Leal et al. (2016) and we assume that traders' activation is endogenous in the following sense: at every time step all traders decide whether they are willing to submit an order by comparing the last recorded price change (in absolute value) to a trader-specific and time-varying threshold, drawn from a common distribution with positive support. In particular, trader i is active at time t if $|r_\tau| > \delta_{i,t} \sim |\mathcal{N}(0, \sigma_\delta^2)|$, where $\tau < t$ denotes the last time in which a trade occurred. If multiple agents are active at time t, they engage in trading in randomised order. If no trader is endogenously activated at time t, then with a certain probability $\phi > 0$ the mechanism falls back to the baseline activation scheme, and a randomly selected trader is asked to submit an order. While the first exogenous mechanism is useful as a baseline scenario to describe and test the functioning of the model, we discover that the second endogenous mechanism is better suited for replicating our stylised facts. In particular, the endogenous activation allows for both

[6] We don't model this phase since by construction has no influence on the price of the security, and is therefore deemed irrelevant with respect to our objective stylised facts.

crowded and uneventful periods in which either many or no orders are submitted, and contribute to clustering of volumes, of trade durations, and of the order-flow.[7]

3.3.3. Traders' behaviour

Traders form expectations about the future (log) return over a certain time horizon h as follows:

$$\hat{r}^F_{i,t+h} = w^F_i \cdot \log\left(\frac{p^F}{p_t}\right) + \varepsilon_t \tag{3.3.1}$$

$$\hat{r}^C_{i,t+h} = w^C_i \cdot \log\left(\frac{p_t}{p_{t-h}}\right) + \varepsilon_t \tag{3.3.2}$$

The superscript F (respectively, C) identifies the fundamentalist (respectively, chartist) strategy. The variable $p^F > 0$ denotes the fundamental price of the security, which is common knowledge among all traders. The term $h \in \mathbb{N}_+$ measures the horizon the trader operates within, and $\varepsilon_t \sim \mathcal{N}(0, \sigma^2_\varepsilon)$ is a common i.i.d. noise component.

The coefficients $w^F_i \sim |\mathcal{N}(0, \sigma^2_F)|$ and $w^C_i \sim \mathcal{N}(\mu_C, \sigma^2_C)$ are trader-specific and capture the "aggressiveness" of the underlying strategy. More specifically, the weight w^F_i quantifies how quickly the price of the stock is expected to revert to its fundamental value. In contrast, the weight w^C_i measures the extent to which traders believe the future return over period h will match its past figure. From eq. (3.3.2), it is also clear that all chartists use only the last realised return over the time-span h to form their expectation.

The above assumptions about chartists' expectations help in containing the dimensionality of the model[8] and stand at variance with previous works (e.g. Pellizzari and Westerhoff, 2009), which instead assume a weighted moving average (typically exponentially or linearly) over *multiple* past returns. However, given our intra-daily setting, we believe that the short memory of chartists mimics more closely the fast response of high-frequency traders to suddenly realised signals. Notice also that we admit an imbalance between trend followers and contrarians, depending on the value of the mean μ_C of the distribution of chartists' weights w^C_i.

[7] Pellizzari and Westerhoff (2009) introduce a similar rule, based on past profits.

[8] The gains in terms of parsimony are due to the fact that we don't need to quantify the memory of the traders (or even worse, a distribution thereof), and a rate of decay of the importance of remote past history.

Once a trader has formed her expectation about the future return, she submits a limit order to the central order book. A limit order, $\ell_{i,t}$, is a tuple {price, quantity, validity} such that: price equals the expected prevailing price at the end of period $t+h$, rounded to the nearest tick; quantity is always fixed to one unit, carrying a positive (respectively, negative) sign if the order is to be stored on the buy (respectively, sell) side of the central order book, depending on whether the trader expects the price to increase or decrease; validity, namely the time after which the order expires and is automatically deleted from the central order book, is set to equal the horizon of the expectation. We assume that all traders have unlimited access to external credit at a zero interest rate, so that they can either short-sell or leverage-buy the stock without bound. In other terms, traders don't face a budget constraint; nevertheless, they are prevented from borrowing an infinite amount of money by the unitary quantity rule. To sum up, a limit order $\ell_{i,t}$ submitted by trader i (either fundamentalist or chartist) at time t takes the form:

$$\ell_{i,t} = \left\{ \texttt{round}\left(p_t \cdot \exp(\hat{r}_{i,t+h}),\, tick\right), \quad \texttt{sign}(\hat{r}_{i,t+h}), \quad t+h \right\} \tag{3.3.3}$$

where $\texttt{round}(\cdot)$ denotes the rounding function, *tick* is the minimum price increment/decrement (a parameter of the market), and $\texttt{sign}(\cdot)$ is the sign function, which takes value 1 if the expected return is positive, -1 if it is negative, and zero otherwise.

We do not model order cancellation as an element of a trader's strategy. However, we introduce the following automatic cancellation rule: when a trader submits a new order, all other orders already submitted by the same trader and stored on the book that are inconsistent with the new expectation are automatically cancelled. These include all orders stored on the opposite side of the book and those orders whose underlying price is deemed unfavourable give current expectations. For example, when a buy (respectively, sell) order is issued at price $\tilde{p}$, all sell (respectively, buy) orders, and all buy (respectively, sell) orders whose price is greater (respectively, less) than $\tilde{p}$, are automatically cancelled. The first condition ensures that a trader never trades with herself, i.e. it rules out the possibility that two orders submitted by the same trader are matched together. The second condition ensures that in case a trader is currently willing to buy (sell) the security at a certain price, she is no longer willing to buy (sell) at a higher (lower) price, as per orders submitted under possibly different beliefs.

It is important to note that no reference whatsoever to any specific time of the day appears in either eq. (3.3.1) or eq. (3.3.2). In other words, none of the traders knows "what the time is" when asked to submit an order, and behaves

identically throughout every phase and instant of the trading day. This implies that, by construction, our model is unable to reproduce SF11, and that any spike in market activity observed in our series has the same probability of occurring during morning, lunch, or afternoon time.

3.4. Numerical simulations

In spite of the very simple behavioural rules that we assume, the complexity associated with the endogenous nature of a limit order book dynamics prevents us from studying the system analytically and to come up with a closed form solution. We thus follow the standard practice in agent-based models of numerically simulating the system and then performing the relevant statistical analysis on the generated time-series.

We start by fixing a few parameters and design principles that are kept stable across our simulations. The market is populated by $N = 1{,}000$ traders; the fundamental price of the stock is constant and equals $p^F = 100$, while the *tick* value, i.e. the smallest possible increment or decrement of the price, equals 0.001. At the beginning of the simulation the price is set to equal its fundamental value, $p_0 = p^F$, and all chartists are provided with a history of past prices between $t = -h$ and $t = 0$ that evolves (backwards) as a pure random walk whose increments are given by the same noise component ε_t present in eqs. (3.3.1) and (3.3.2). Finally, we fix the horizon of traders' expectations h to 1,000 seconds (simulation time steps); incidentally, this value equals the expected duration between two consecutive activations of a same trader within the exogenous activation scheme, given the number of traders N.

In order to perform the statistical analysis needed to validate our model against the stylised facts listed in Section 3.2, we rescale the relevant time series by pooling the stream of trade messages into homogeneous time windows of one calendar minute each.[9] The minute-by-minute price (respectively, volume) series corresponds to the average (respectively, sum) of the underlying trading prices (volume) during that minute. Following Section 3.3.1, the main trading session consists of 510 minutes.

We simulate the model[10] under three different scenarios. In Section 3.4.1 we

[9] This is necessary because, at the finest level of granularity, our simulations yield time series of the relevant quantities that are *irregular* since, by construction, trade emerges endogenously when at least two crossing orders are stored on the central order book.

[10] The simulation is coded in `C++11` and largely exploits the object-oriented programming paradigm, defining *classes* for `traders`, for the central order `book`, and for the `order` data

only include purely noise traders. This allows us to evaluate the impact of market microstructure on the generation of stylised facts. In Section 3.4.2 we investigate the effects of the interplay between fundamentalists and chartists on the one hand, and market microstructure on the other hand, under the baseline exogenous activation. Finally, in Section 3.4.3 we add a further element of complexity by assuming that traders follow the endogenous participation scheme described in Section 3.3.2. Finally, in Section 3.4.4 we perform some complementary sensitivity analyses. The results that we show correspond to averages across 100 Monte Carlo simulations of a fully fledged trading day (see Section 3.3.1). All confidence intervals are set at the 95% level.

3.4.1. Noise traders only

The first simulation scenario, which we dub NT, is useful to properly disentangle the effects implied by the market microstructure details on the generation of market statistical properties from those implied by our assumptions about traders' behaviour. Noise traders do not condition their investment on any market-related variable; rather, they *"trade on noise as if it were information"* (Black, 1986). Given our formulation, we set all the w_i's in eqs. (3.3.1) and (3.3.2) to zero, such that the expected return for each trader will only depend on the i.i.d. noise component ε_t.

Table 3.1 summarises the specific parametrisation. By setting δ_t to infinity we rule out endogenous activation, and by setting $\phi = 1$ we ensure that exactly one trader is activated at every time step t.

Fig. 3.1 pictures the relevant plots under this scenario. Panel (a) shows the evolution of the minute-by-minute market price for a typical trading day while panel (b) reports its log-differences. The average number of price changes per day under this scenario is 13578. If microstructure effects were completely irrelevant, given our limit price function (3.3.3), then the time series of realised

structure. The code supports the execution of fully parallel Monte Carlo simulations, using the `OpenMP` framework. Random number generation relies on the 32-bit Mersenne Twister, as implemented in the C++ Standard Library (`std::mt19937`). Parameters and initialisation for all the Monte Carlo simulations are passed through a single `json` file during run-time, so that the code needs not be (re)compiled every time a new scenario is simulated. The file is parsed using the `jsoncpp` library. Each Monte Carlo simulation returns a `SQLite` database file containing the associated initialisation and a stream of messages from the central order book, each corresponding to a successful transaction (each message reports the current POSIX timestamp, bid, ask, transaction price, quantity, and depth of the book for both sides). The output databases are then imported and analysed using `R`.

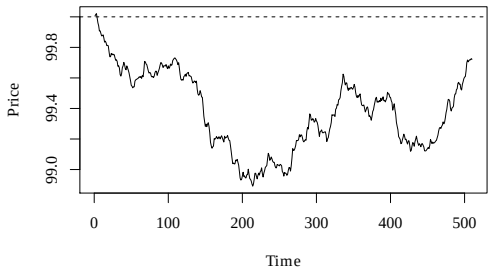

(a) : price of a typical trading day

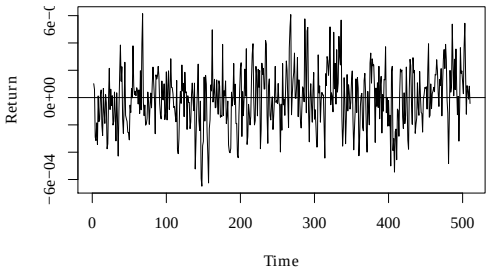

(b) : returns of a typical trading day

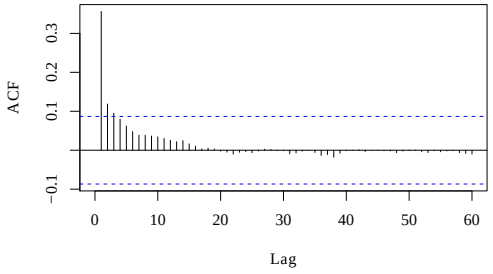

(c) : ACF of (raw) returns

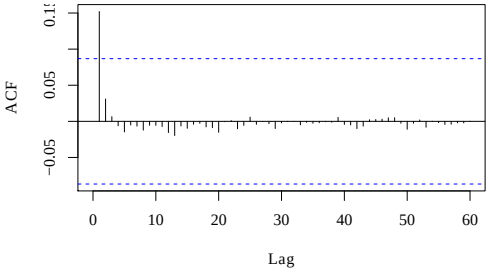

(d) : ACF of absolute returns

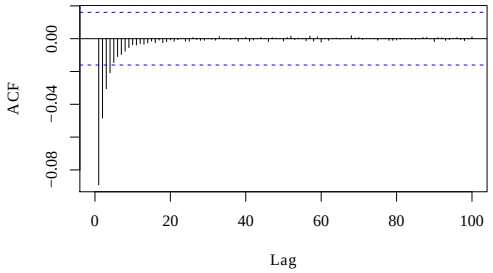

(e) : ACF of inter-trade times

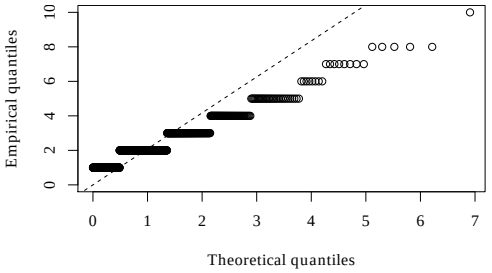

(f) : Q-Q plot of inter-trade times

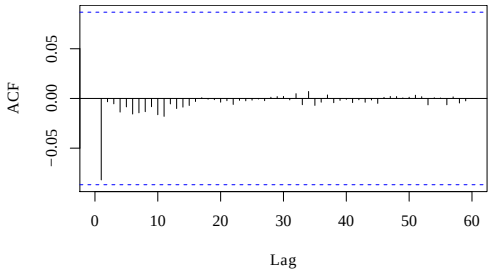

(g) : ACF of exchanged volumes

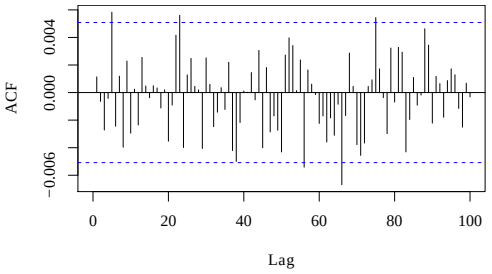

(h) : ACF of order book imbalance

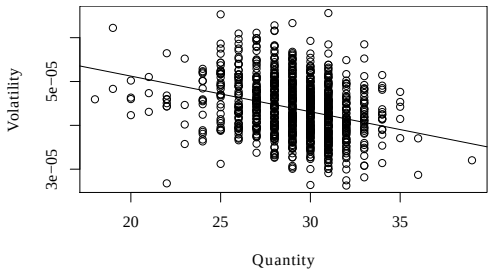

(i) : volume/volatility correlation

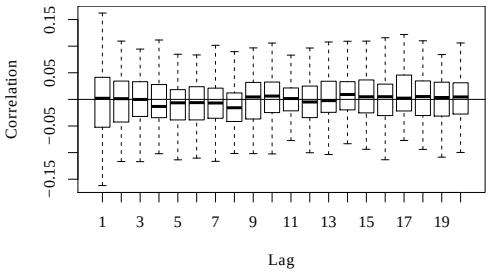

(j) : leverage effect correlation coefficients

Figure 3.1.: Main stylised facts under the noise traders scenario.

Parameter	Value
number of traders	$N = 1{,}000$
fundamental price	$p_F = 100$
initial price	$p_0 = p_F$
smallest price change	$tick = 0.001$
horizon/order validity	$h = 1{,}000$
noise process	$\varepsilon_t \sim \mathcal{N}\left(\mu_\varepsilon = 0,\ \sigma_\varepsilon^2 = \texttt{5e-5}\right)$
fundamentalist weight	$w_i^F = 0$
chartist weight	$w_i^C = 0$
activation threshold	$\delta_t \to +\infty$
activation fallback probability	$\phi = 1$

Table 3.1.: Parameters and initial conditions for the NT scenario.

returns should share the same statistical properties of the i.i.d. series of expected returns. In contrast, we find that the Ljung-Box statistic strongly rejects (p-value < 0.001) the null hypothesis of independence. This is also visible in panel (c), which shows the autocorrelation function of price returns. Positive autocorrelation for the first lag is substantial, and for the second lag is very close to the confidence threshold. The Augmented Dickey-Fuller (ADF) test doesn't reject (p-value < 0.001) the presence of a unit root within the price series. Prices are therefore well approximated by a random walk, although its increments are not independent. Moreover, the (absolute) kurtosis of the sample distribution of returns, $\kappa \approx 3.43$, is only negligibly higher than that of expected returns, that by construction equals 3. We conclude that the EURONEXT microstructure setup does force a time dependence character into the resulting series, although this lasts for just under a couple of minutes.

Furthermore, panel (d) pictures the autocorrelation function of the absolute value of returns. Its rate of decay is very high and only the first lag is significant; we conclude that volatility clustering is not present in this scenario. Panels (e) and (f) relate to the properties of time durations between subsequent trades. Panel (e) shows the autocorrelation function of such durations, whereas panel (f) pictures a quantile-quantile plot of their distribution, compared to a fitted exponential distribution. The autocorrelation function is negative for the first few lags, and the distribution has a tail that is thinner than that of an exponential distribution. This suggests that there is no correlation structure in either the

exchanged volumes of the asset (panel (g)), nor in the clustering of buy and sell orders stored in the book (panel (h)). Finally, panel (i) shows the presence of a negative relationship between exchanged volumes and volatility, instead of the predicted positive correlation. Likewise, panel (j) suggests that any leverage effect is absent in our series.

The first column on the right of Table 3.4 compares these results with our objective stylised facts. Only two facts are matched as a consequence of the interactions implied by the market microstructure. It is evident that more structure on the behaviour of the traders is needed in order to obtain a more realistic dynamics.

3.4.2. Fundamentalists and chartists

In this scenario (FC), we move a step forward by switching on our fundamentalist and chartist specifications, according to eqs. (3.3.1) and (3.3.2). On the one hand, fundamentalist traders anchor the price dynamics to a neighbourhood of the fundamental price p^F. On the other hand, chartists tend to exacerbate or to counteract the prevailing trend, depending on their being trend followers or contrarians. Accordingly, the stronger the magnitude of trend following behaviour, the wider the divergence of price from p^F (either upwards or downwards) should be. The parametrisation we propose, reported in Table 3.2, yields a price dynamics characterised by a unit root, as in the previous section (the ADF test doesn't reject the null with p-value < 0.001). Note that we set the value of $\mu_C > 0$. Reasonably, the overall sentiment among the crowd of chartists generates a self-reinforcing dynamics, rather than a self-opposing one.[11]

Fig. 3.2 shows the relevant plots under this scenario. As expected, the evolution of the price series is more "centred" around the fundamental value p^F with respect to the NT scenario thanks to the fundamentalists' anchoring behaviour (panel (a)). However, the presence of chartists introduces a persistence character in the dynamics of returns: the presence of both trend followers and contrarians is crucial because their effect on the autocorrelation function of returns cumulates in absolute value, but cancels out when the sign is taken into account. This is clearly visible in panels (c) and (d).[12] Intuitively, while we allow a slight im-

[11] Moreover, as will become clear later in Section 3.4.4, we find that this assumption fosters the fat-tailedness character of inter-trades durations, and thus helps in replicating stylised fact SF9.

[12] In a separate experiment (not shown) we set $w_i^C \sim \pm\,|\mathcal{N}(\mu_C, \sigma_C^2)|$, i.e. we include either trend-followers or contrarians but not both. In this case we find that the autocorrelation function

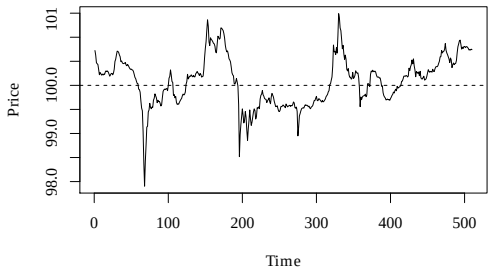

(a) : price of a typical trading day

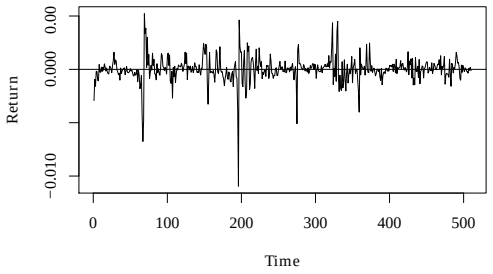

(b) : returns of a typical trading day

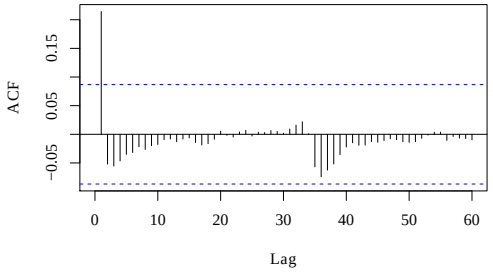

(c) : ACF of (raw) returns

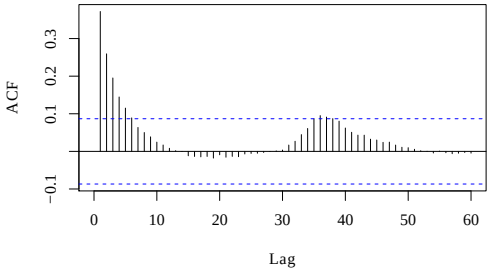

(d) : ACF of absolute returns

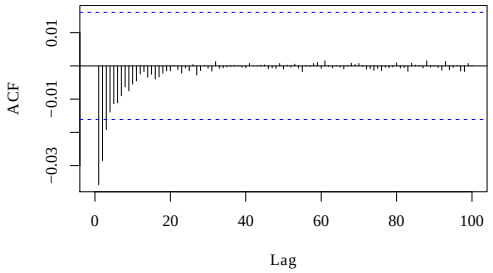

(e) : ACF of inter-trade times

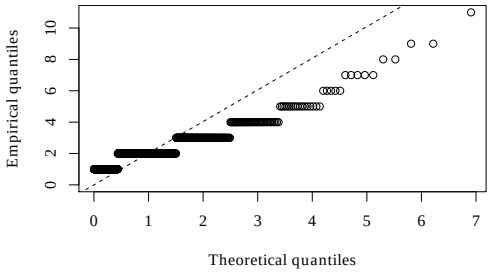

(f) : Q-Q plot of inter-trade times

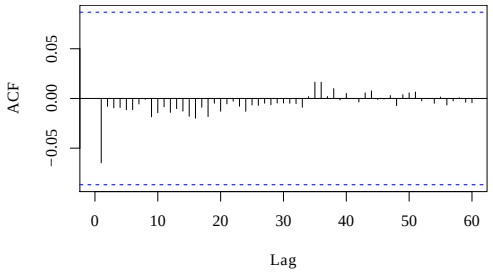

(g) : ACF of exchanged volumes

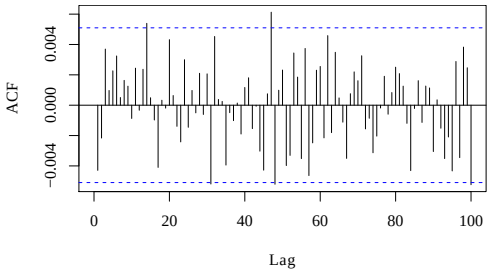

(h) : ACF of order book imbalance

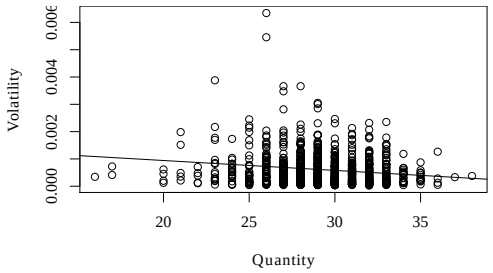

(i) : volume/volatility correlation

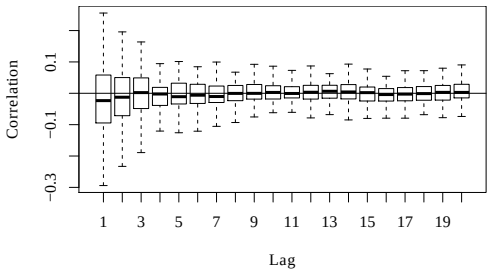

(j) : leverage effect correlation coefficients

Figure 3.2.: Main stylised facts under the fundamentalists vs. chartists scenario.

Parameter	Value
number of traders	$N = 1{,}000$
fundamental price	$p_F = 100$
initial price	$p_0 = p_F$
smallest price change	$tick = 0.001$
horizon/order validity	$h = 1{,}000$
noise process	$\varepsilon_t \sim \mathcal{N}\left(\mu_\varepsilon = 0,\ \sigma_\varepsilon^2 = \texttt{5e-5}\right)$
fundamentalist weight	$w_i^F \sim \left\lvert\mathcal{N}\left(\mu_F = 0,\ \sigma_F^2 = 0.001\right)\right\rvert$
chartist weight	$w_i^C \sim \mathcal{N}\left(\mu_C = 0.01,\ \sigma_C^2 = 0.1\right)$
activation threshold	$\delta_t \to +\infty$
activation fallback probability	$\phi = 1$

Table 3.2.: Parameters and initial conditions for the **FC** scenario.

balance between followers and contrarians, a larger imbalance would have the effect of adding memory to the autocorrelation function of (raw) returns, which is contradicted by empirical evidence. The average number of price changes, 14187, is in line with the previous scenario. In contrast, the kurtosis of minute returns increases to 17.45, thus replicating SF1.[13]

Finally, the statical properties reported in panels (e), (f), (g), (h), (i), and (j) are qualitatively similar to the **NT** case. This indicates that the timing structure of orders submission and matching is not substantially influenced by the presence of the new behavioural specification; exchanged volumes display no persistence character either.

The second column on the right of Table 3.4 summarises the list of stylised facts reproduced with the introduction of fundamentalist and chartist strategies. The improvement with respect to the noise traders scenario is clear: volatility clustering and leptokurtosis of price returns are now correctly matched. However, more structure is needed if one wants to reproduce also orders' timing and clustering properties.

of returns (panel (c)) and of absolute returns (panel (d)) look very similar and thus fail to validate our target stylised facts SF2 and SF3.

[13]The kurtosis decreases with the time window and reverts back to 3, i.e. to statistical *normality*, for 15-minute returns.

Parameter	Value
number of traders	$N = 1{,}000$
fundamental price	$p_F = 100$
initial price	$p_0 = p_F$
smallest price change	$tick = 0.001$
horizon/order validity	$h = 1{,}000$
noise process	$\varepsilon_t \sim \mathcal{N}\left(\mu_\varepsilon = 0,\ \sigma_\varepsilon^2 = \texttt{5e-5}\right)$
fundamentalist weight	$w_i^F \sim \left\lvert\mathcal{N}\left(\mu_F = 0,\ \sigma_F^2 = 0.001\right)\right\rvert$
chartist weight	$w_i^C \sim \mathcal{N}\left(\mu_C = 0.01,\ \sigma_C^2 = 0.1\right)$
activation threshold	$\delta_t \sim \left\lvert\mathcal{N}\left(\mu_\delta = 0,\ \sigma_\delta^2 = 0.3\right)\right\rvert$
activation fallback probability	$\phi = 1/3$

Table 3.3.: Parameters and initial conditions for the **EA** scenario.

3.4.3. Endogenous activation

In this final scenario, which we label **EA**, we assume that fundamentalists and chartists endogenously activate according to the scheme outlined in Section 3.3.2. Table 3.3 summarises the specific parametrisation that we employ in this scenario. The ultimate goal is to retain the properties encountered in the previous scenarios and, in addition, to replicate those properties related to the duration and clustering of orders and those about the volumes of trade.

The endogenous activation scheme captures one ever more common high-frequency nature of financial markets (see Easley et al., 2012). A crowd of traders, many of which are algorithmic machines, typically responds very quickly to a newly posted signal and engages in trading for a while until coordination on a new price has emerged.[14]

Fig. 3.3 pictures the relevant plots under this scenario. Panels (a) to (d) are qualitatively similar to those of scenario **FC**, suggesting that the good proper-

[14] In principle, such a signal can arise either from *within* the order book, e.g. as a disruptive newly submitted order, or from *outside*, in which case it is related to fundamental news about the asset. Empirically, it has been shown that only a fraction of realised volatility is attributable to freshly available news about dividends, prospective earnings, or other crucial balance sheet and macroeconomic variables (see e.g. Cutler et al., 1989; Shiller, 1981). In our model no news is ever released and all traders agree on a constant fundamental value; thus, all the signals come from within the order book, and are the result of sheer trading activity by the traders. The totality of the generated volatility is *excess volatility*.

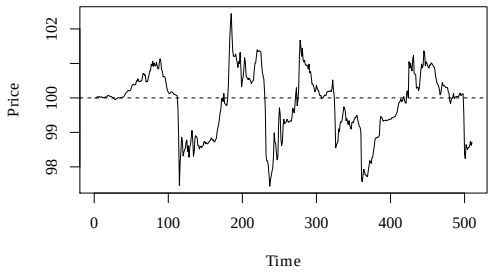

(a) : price of a typical trading day

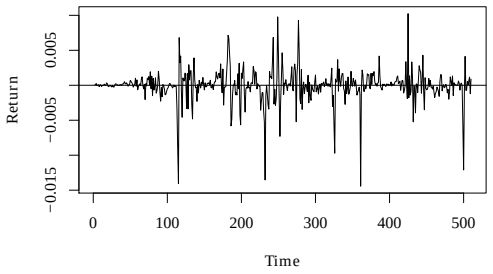

(b) : returns of a typical trading day

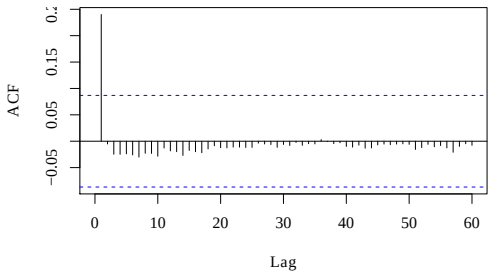

(c) : ACF of (raw) returns

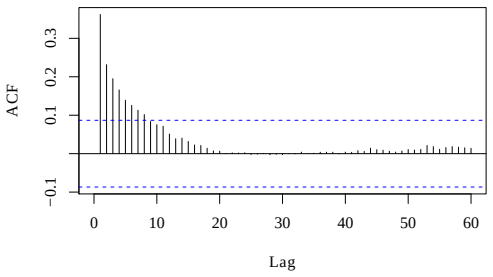

(d) : ACF of absolute returns

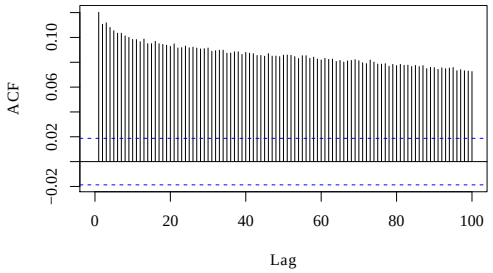

(e) : ACF of inter-trade times

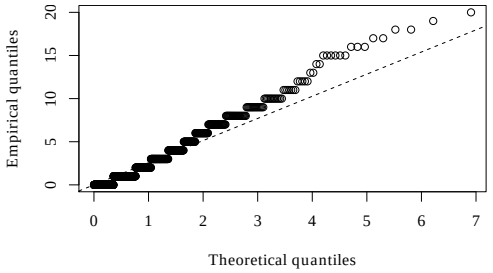

(f) : Q-Q plot of inter-trade times

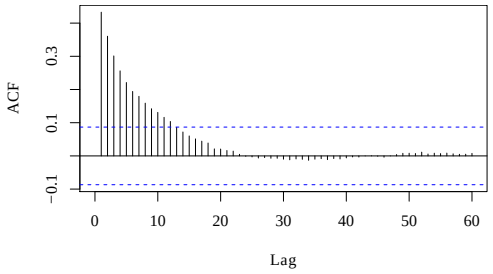

(g) : ACF of exchanged volumes

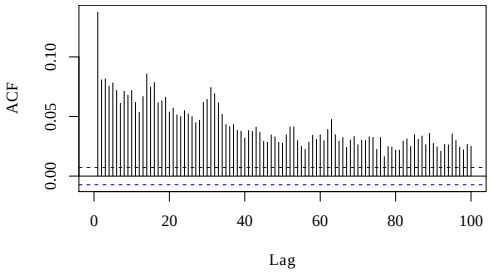

(h) : ACF of order book imbalance

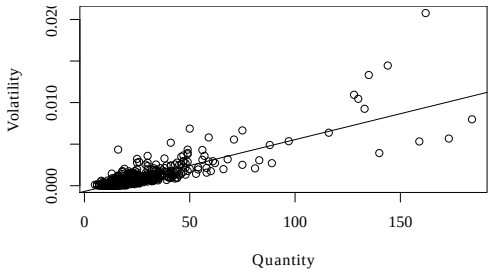

(i) : volume/volatility correlation

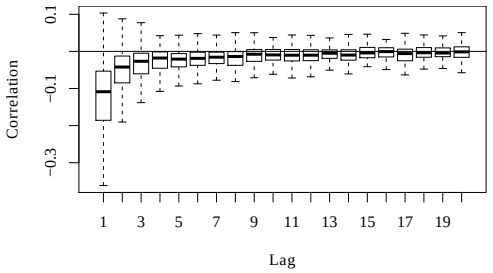

(j) : leverage effect correlation coefficients

Figure 3.3.: Main stylised facts under the fundamentalists vs. chartists scenario with endogenous activation.

ties about price and returns generated in the latter setting have not been compromised by the new activation assumption. Leptokurtosis has increased to a minute-by-minute figure of $\kappa \approx 72.33$, decreasing to around 6.5 for 15-minute returns, and 3.7 for 30-minute returns. The average number of price changes, 11272, has decreased as a result of the new participation scheme, but is still a perfectly acceptable level for liquid traded securities (Cont, 2011).

The main benefits of endogenous activation are noticeable in the subsequent panels of Fig. 3.3. For the first time, panel (e) shows a strong and very slowly decaying autocorrelation in inter-trade durations (SF8), and the quantile-quantile plot in panel (f) suggests that the tail of their distribution is fatter than exponential (SF9). Moreover, both volumes (panel (g)) and order-flow (panel (h)) are clustered (matching respectively SF5 and SF10). Panel (i) shows a positive and significant relationship between volumes and volatility (p-value < 0.001) (as per SF6). An analogously significant relationship holds also for pooled series at 15-minute and 30-minute level. Finally, the boxplot in panel (j) suggests a slight improvement with respect to the previous scenarios: the correlation coefficient for the first 10 lags is negative and increasing for the majority of our Monte Carlo simulations. Nonetheless, since the 'whiskers' of the plot (denoting the $\pm\ 1.5 \cdot IQR$ markers of the underlying distribution) are very spread apart, we conservatively consider SF4 as not matched.

The rightmost column of Table 3.4 suggests that most of the stylised facts described in Section 3.2 are successfully reproduced by this version of the model. In particular, the emergence of the properties about orders' duration and clustering is very much linked to the dynamics induced by endogenous activation in this scenario. Indeed, the level of the variance of the distribution of agents' activations thresholds, $\sigma_\delta^2 = 0.3$, is such that, on average, exactly one trader is endogenously activated at time t in response to a realised absolute return $|r_{t-1}| \approx 0.000375$, whereas the average absolute return in the **FC** scenario is approximately 0.0003. This means that most of the time traders are not endogenously activated, and the fallback exogenous activation scheme takes over, with probability $\phi = 1/3$. However, due to the leptokurtic nature of returns (SF1), there exist periods in which a much larger-than-average price change takes place, and a multitude of traders are willing to submit new orders at the same time. Moreover, the price change generated by such turbulent event is likely to be itself larger than the δ_t threshold for a number of traders, possibly triggering a new wave of crowded endogenous activation in the next period, ultimately lengthening the duration of the price adjustment process.

		scenario		
stylised fact		**NT**	**FC**	**EA**
SF1	leptokurtic returns	✗	✓	✓
SF2	no linear autocorr.	✓	✓	✓
SF3	volatility clustering	✗	✓	✓
SF4	leverage effect	✗	✗	✗
SF5	autocorrelation of volumes	✗	✗	✓
SF6	volume/volatility correlation	✗	✗	✓
SF7	number of price changes per day	✓	✓	✓
SF8	autocorrelation of durations	✗	✗	✓
SF9	fat-tailed durations	✗	✗	✓
SF10	order-flow clustering	✗	✗	✓
SF11	U-shaped activity	✗	✗	✗

Table 3.4.: Replication of the target stylised facts within all the simulated scenarios.

3.4.4. Sensitivity analysis

In this section we briefly discuss the effect of varying, one at a time, the main parameters of the model[15] in a neighbourhood of the parametrisation used in the most complete scenario, i.e. **EA** scenario, with endogenous traders' activation (cf. Table 3.3). For each of the discussed parameters, we present plots showing the change in the relevant statistics of the simulation and we relate it to the stylised facts of financial markets.

Changes in the variance of the fundamentalists' weights, σ_F^2

In the limit of $\sigma_F^2 \to 0$, the price becomes less "anchored" to the fundamental value. In this extreme case the persistence of volatility (absolute returns), of volumes, of trade durations, and of the order flow, the number of price changes, and the leptokurtic signature of returns are maximised. The plot in panel (a) of Fig. 3.4 shows the number of statistically significant lags (at the 95% confidence level) as a function of of σ_F^2 for each of the autocorrelation functions calculated in the **EA** simulation scenario (cf. Fig. 3.3). Autocorrelations computed on minute-

[15] We also experimented with changes in other parameters of the model besides the ones discussed in this section. We do not report the results of these additional sensitivity analyses here. However, they are available from the authors upon request.

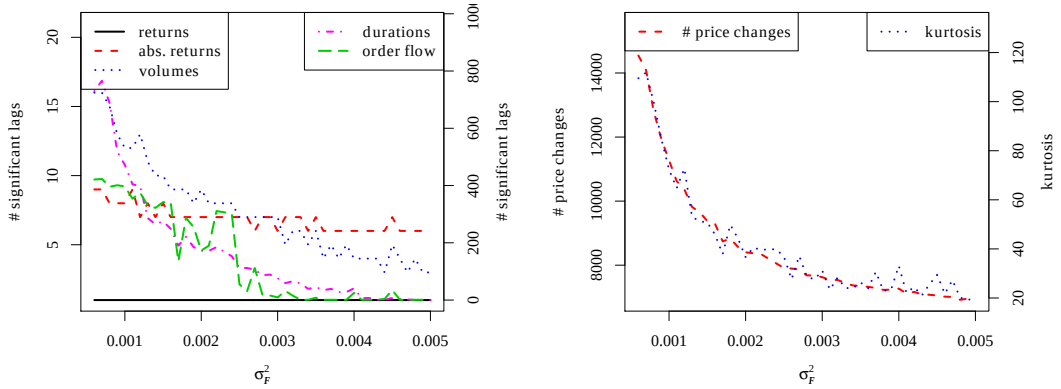

(a) : autocorrelation functions

(b) : # of price changes and kurtosis

Figure 3.4.: Sensitivity analysis for σ_F^2. Other parameters as per the **EA** scenario.

by-minute data (namely returns, absolute returns, and volumes) refer to the left scale. Those computed on tick data (namely durations and order flow) are on the right scale. In addition, panel (b) depicts the effect of σ_F^2 upon the number of price changes per day (left scale) and the minute returns' kurtosis (right scale). All measures are averages across 100 Monte Carlo simulations for each value of the underlying parameter.

From the analysis of the above mentioned figures it is clear that as σ_F^2 grows, fundamentalists counteract the effect of technical traders and all the considered statistics either decrease or remain unchanged. It is also clear from the picture that σ_F^2 has not effect whatsoever on the persistence of raw returns.

Changes in the variance of the chartists' weights, σ_C^2

At very low values of σ_C^2 and with $\mu_C > 0$ the distribution of chartists' weights is significantly skewed towards trend-following strategies.This causes raw returns to be autocorrelated for several lags (see Fig. 3.5(a), left scale). As σ_C^2 increases the distribution of the chartists' strategies tends to be more balanced between trend-followers and contrarians. This decreases the autocorrelation of raw returns without significantly influencing the persistence of either absolute returns and of traded volumes.

At the same time, the extent of chartists' disagreement about future returns increases with σ_C^2, because the weights of followers and contrarians are located farther and farther away from one another. *Cœteris paribus*, this increases the occurrence of abnormal returns (cf. Fig. 3.5(b)). In addition, abnormal returns and the mechanism of endogenous activation (see Sections 3.3.2 3.4.3) encourage

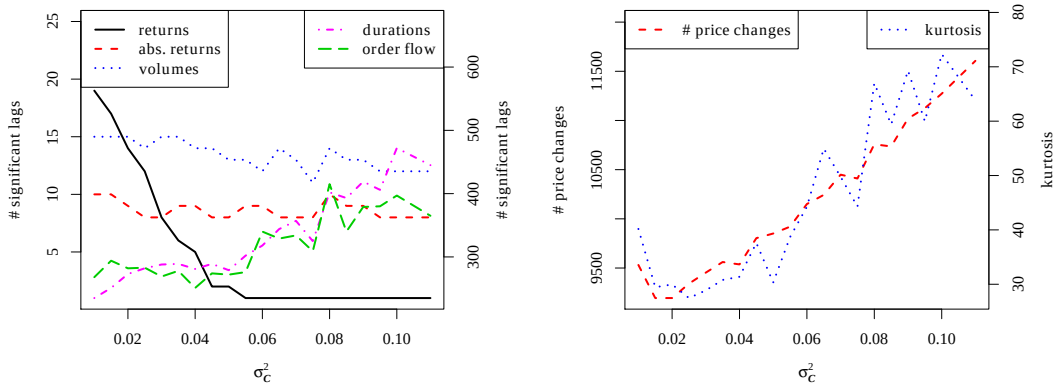

(a) : autocorrelation functions

(b) : # of price changes and kurtosis

Figure 3.5.: Sensitivity analysis for σ_C^2. Other parameters as per the **EA** scenario.

several traders to post their orders simultaneously. This increases the number of price changes per day and generates persistence in inter-trade durations and in order flows (see Fig. 3.5(a), right scale).

Changes in the variance of traders' activation thresholds, σ_δ^2, and of the fallback probability ϕ

These two parameters jointly regulate both the total amount of trade that takes place in the market and the timing structure thereof. First, note that, by definition $\delta_{i,t} \sim |\mathcal{N}(0,\sigma_\delta^2)|$, and therefore $\mathbb{E}[\delta_{i,t}]$ is an increasing function of σ_δ^2. In the limit of $\sigma_\delta^2 \to \infty$), $\delta_t \to \infty$ and activation is never endogenous and the average amount of trade and of price changes in the market is *cœteris paribus*, a monotonically increasing function of ϕ (cf. the dashed red line in Fig. 3.6(b)). If $\phi = 0$ then, trivially, no trader is ever activated and no trade takes place; if instead $\phi = 1$ the amount of trade is maximised, under the assumption of uniform activation: exactly one trader is activated in every period (cf. scenarios **NT** and **FC**) and both volumes and the number of transactions (a superset of the number of price changes) are bounded from above by t.

Furthermore, when σ_δ^2 is large $\delta_{i,t}$ is also large on average. In this case inter-trade times exhibit little serial correlation (cf. Fig. 3.7(a)). In the opposite limit, $\delta_t = 0$, all traders are instead active at every time step (regardless of ϕ), and the number of transactions is maximised and bounded from above by $N \cdot t$. Finally, lower average values of δ_t cause larger crowds of traders to participate in response to a given signal, boosting the number of price changes (cf. Fig. 3.7(b)).

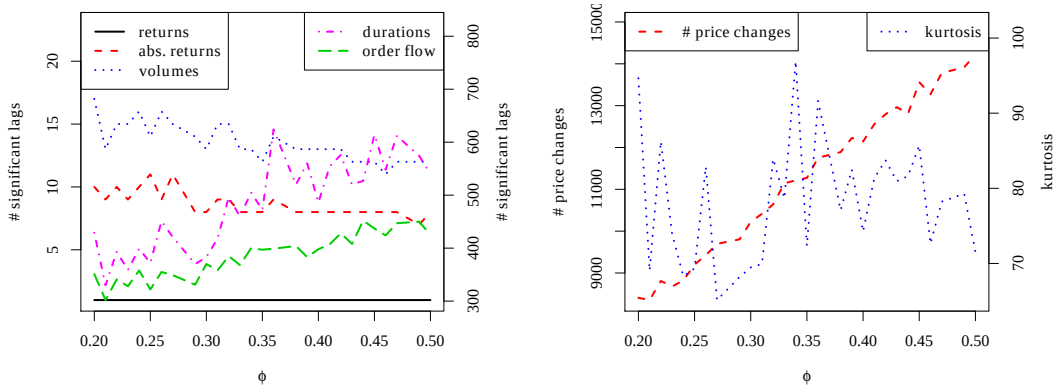

(a) : autocorrelation functions (b) : # of price changes and kurtosis

Figure 3.6.: Sensitivity analysis for ϕ. Other parameters as per the **EA** scenario.

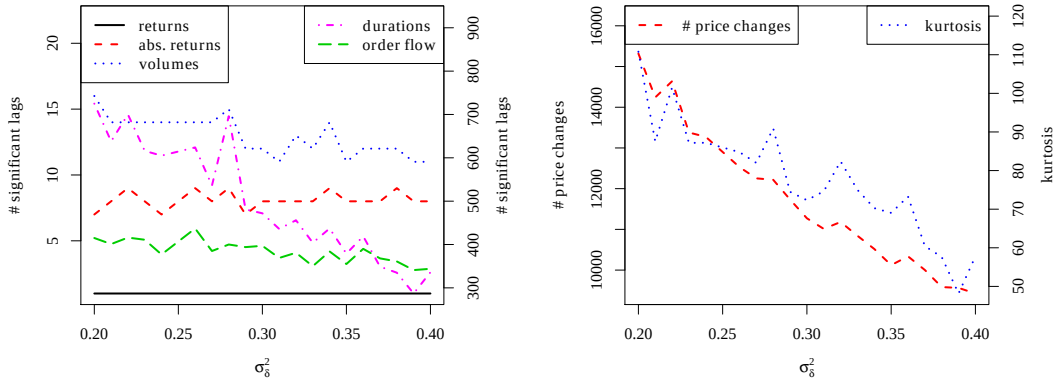

(a) : autocorrelation functions (b) : # of price changes and kurtosis

Figure 3.7.: Sensitivity analysis for δ. Other parameters as per the **EA** scenario.

As a result, both the order-flow and trading times tend to cluster (see again Fig. 3.7(a)).

Changes the investment horizon and orders' validity: h

The shorter the order's validity h, the fewer the orders stored on the book at all times. Accordingly, the number of price changes per day decreases as a function of h (cf. Fig. 3.8(b)).

Furthermore, the parameter h also sets the memory span of chartists (cf. eq. (3.3.2)). An increase in this parameter thus leads the latter to exacerbate small trends in prices. This is because cumulative returns over longer periods tend to be larger than returns over short periods in presence of some positive

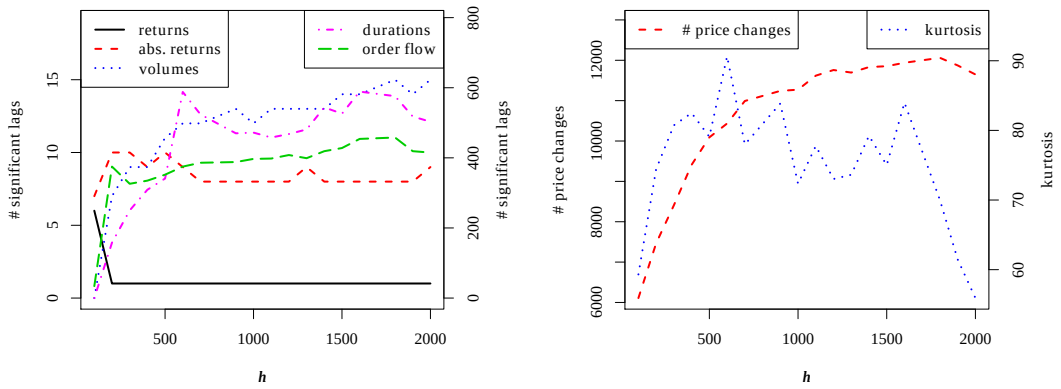

(a) : autocorrelation functions

(b) : # of price changes and kurtosis

Figure 3.8.: Sensitivity analysis for h. Other parameters as per the **EA** scenario.

autocorrelation. Accordingly, the higher h, the more likely expected returns between time t and $t + h$ are large in absolute terms, and the farther is the limit price of newly submitted orders with respect to the current price. This effect, combined with the endogenous activation mechanism, contributes to the persistence of absolute returns, volumes, orders' duration and sides, which are all observed for small increases in h when starting from a low base value of $h = 100$ (see Fig. 3.8(a)). It also explains the increase in the returns kurtosis displayed in (cf. Fig. 3.8(b)).

Nevertheless, the above effects quickly vanish with further increases in h. Indeed, the number of significant lags in (raw) returns autocorrelation quickly decays, eventually stabilising at a low value of 1 already for h around 100. In addition, the degrees of persistence in absolute returns and in order flow and durations also converge to a stationary value. Finally, the kurtosis of price returns evolves non-monotonically with h, eventually returning to the same low value observed at $h = 100$. Notice that in the **EA** scenario we set $h = 1000$. It follows that increases in the value of h in that scenario have no significant effects on the autocorrelation functions of the main market variables we consider.

3.5. Concluding remarks

The distinctive statistical properties that shape financial market dynamics at daily and intra-daily frequencies have been typically attributed to the specific patterns of information release and its diffusion among the population of traders. We show that many such properties can be simultaneously reproduced

in a framework wherein fundamental news are absent and information, originating from within the financial market (as the by-product of trading activity) is common knowledge. We build a parsimonious agent-based model in which trading and its statistical properties emerge endogenously out of the interaction between fundamentalist and chartist strategies on the one hand, and a realistic market microstructure specification on the other hand.

A novel element that we introduce is the definition of simulation time in terms of a strict schedule that we borrow from the microstructural specification of a real stock market, namely the EURONEXT. We believe this plausibly relates each iteration of our numerical simulations to proper calendar time, and enables us to investigate which properties apply within a specific time-window and how they evolve at different time-scales. We also devise a simple endogenous activation scheme that encourages traders participation in an increasing fashion with realised profit opportunities.

We find that our assumptions regarding the underlying microstructure introduce a slight dependence in the series of returns, which quickly fades away within a couple of minutes. We also find that the fundamentalist vs chartist framework is suitable for replicating the empirically validated dependence properties of returns (leptokurtosis, absence of linear autocorrelation, and volatility clustering). Nonetheless, the introduction of our endogenous participation scheme proves crucial for the emergence of the persistence character in the timing structure of market activity. Under this scenario we are able to simultaneously reproduce, along with the stylised facts just mentioned, the fat-tailed and serially correlated nature of durations between trades, and the clustering of both volumes and order-flow.

We believe that our framework can be fruitfully extended in several directions. First, our model cannot reproduce, by construction, the U-shaped pattern of intra-day market activity. More stringent assumptions regarding the traders' budget constraint or the introduction of a time feedback that puts pressure on traders close to the end of the trading day (e.g. due to margin requirements) could be useful in this respect. Similarly, a more structured specification of chartists' behaviour might unveil a more asymmetric response of volatility with respect to price drops and surges (leverage effect). Finally, in this chapter we only considered the ability of the model to qualitatively replicate the main statistical properties of financial markets. However, one could further fine-tune the calibration of the parameters of the model by exploiting actual financial data.

This might allow one to perform quantitative experiments on regulatory policies affecting market microstructure or trading behaviour.

This page intentionally left blank.

4

A 2-step functional principal component analysis of intra-day volatility trajectories[1]

4.1. Introduction

The extent to which a certain asset or investment is expected to yield dispersed uncertain returns constitutes a top concern in the decision making of virtually any financial actor, ranging from practitioners to policymakers. To mention a few, returns volatility enters the decisions of investors and risk managers when considering the purchase of a financial instrument, of banks and financial institutions when issuing and pricing a new derivative, and of central banks and financial regulators who are committed to maintaining orderly operations within the financial system. Although the traditional use of volatility as a proxy for risk has inevitably attracted a negative connotation, volatility itself can be regarded as a commodity and consequently can be priced and exchanged as such. Tradable instruments replicating the implied volatility of underlying assets and the so-called volatility arbitrage have increasingly gained popularity in recent times, especially among high-frequency traders. The value attached to understanding the emergence and evolution of volatility over time is therefore tautologically justified by the enormous amount of wealth at stake.

In this contribution we study intra-day volatility trajectories from two major financial markets along the lines of recent developments in functional data analysis (FDA). Volatility is not an observable phenomenon and as such it can not be directly inferred from discretely observed tick data. Strictly speaking, the theoretically observable quadratic variation would require continuous monitoring of the underlying log-price process. Within traditional volatility models

[1] This chapter is a joint work with Matteo Barigozzi (Department of Statistics, London School of Economics and Political Science, London, UK).

(see e.g. Bandi and Phillips, 2003; Florens-Zmirou, 1993; Renò, 2008), it is generally assumed that the volatility process is functionally linked to some observed state variable, such as prices or returns. Most of these models, however, target a single realisation of the volatility process whereas we aim at modelling repeated intra-day realisations thereof. FDA comes to rescue as it provides rigorous statistical tools that are naturally adapted to the analysis of collections of smooth continuous curves, all defined over the same support. In our framework these curves consist of individual intra-day volatility trajectories observed over multiple trading days, each sharing the same duration.

To the best of our knowledge, Müller et al. (2011) is the first and only piece of research that applies this methodology to the investigation of intra-day volatility curves. They first devise a general diffusion model with drift for log-returns of a single financial security that targets repeated realisations of the underlying process. This model is then discretised and applied to real data on the S&P500 index, yielding a sample of intra-day time series which are then transformed into continuous functional objects by means of a smoothing algorithm. Finally, the resulting collection of curves is studied in a principal component fashion.

Our work proceeds along similar steps as Müller et al. (2011) although we extend their contribution in a number of directions. First, we generalise their framework, focussed on a single market index, in order to allow for the joint investigation of an arbitrarily large collection of assets, e.g. all the constituents of the underlying index. This upgrade, which serves as our baseline model, dramatically complicates the analysis since it brings an additional dimension to the collection of curves, beyond the daily repeated realisations. We tackle the issue by devising a 2-step procedure based on functional principal component analysis: the first step reduces the dimensionality across days while the second, applied to the output of the first step, reduces the dimensionality across assets, yielding a single set of curves that contain most of the information embedded in the original trajectories. In particular, we focus on the relative contribution of the mean and of the first functional principal component, in predicting the original volatility trajectories of the individual assets. Our second improvement consists of fitting a CAPM-inspired factor model to the data in order to separate the market component of the constituents' volatility, namely the part that correlates with the index as a whole, from the residual idiosyncratic component, and study these distinct parts both independently and jointly. The analysis of the 1-dimensional market component is similar to the exercise carried out in Müller et al. (2011), while the 2-dimensional idiosyncratic component requires

the aforementioned 2-step dimension reduction procedure. The output of the latter allows the analysis of the so-called common idiosyncratic volatility (CIV), i.e. the part of idiosyncratic returns volatilities that is common across assets, and to quantify its contribution, if any, in shaping the original volatility trajectories.[2] Finally, we apply our extended theoretical framework to the empirical investigation of two major financial market indices for which we have high-frequency tick data available for all the underlying constituents, namely the S&P500 and the EURONEXT 100. While we proceed with a parallel and independent analysis of the two markets, a preliminary international comparison of the underlying volatility patterns is drawn.

Our results are as follows. First, we show that our baseline model outperforms the one proposed by Müller et al. (2011) in predicting the original, asset-specific volatility trajectories. This result is not surprising since the latter model, which corresponds to the market volatility branch of our extended model, exploits far less information with respect to the former. Second, more importantly, we show that our extended model, based on the CAPM-inspired distinction of market and idiosyncratic volatilities greatly improves the prediction performance with respect to the baseline model, even though it exploits the very same original information. Third, we show that within the second step of our 2-step procedure the loss of information due to dimension reduction is virtually nil. This implies that the resulting single curve, representing the CIV, incorporates all the information relevant for the prediction of the original volatility trajectories that is embedded in all the asset-specific functional principal components. Moreover, we show that in all the various models presented, the contribution given by the first functional principal component, even discarding higher orders' components, substantially improves the prediction of the original trajectories with respect to using the mean in isolation. Finally, a visual analysis of the various functional principal components involved in our procedure and their corresponding loading coefficients gives some insight about some of the 'stylised facts' that financial markets' data typically exhibit, such as the U-shaped intra-day activity motive and volatility clustering, and the patterns of interna-

[2] Herskovic et al. (2016) consider a large panel of stocks over the period 1926–2010. After estimating a factor model on daily returns (using either the value-weighted market portfolio, the 3 Fama and French factors, or the first five principal components of the cross section of returns) they find a substantial degree of common variation in the residuals, both across firm-size quantiles and industry group. Barigozzi and Hallin (2016) obtain similar results for the S&P100 index and find that the magnitude of the CIV is more pronounced during periods of financial turmoil, such as the great financial crisis of 2008–2009.

tional substitution and complementarity in place between the American and European stock markets under study.

A number of contribution are close to ours in spirit. Kokoszka et al. (2014) propose a functional dynamic factor model of intra-day cumulative returns, a stationary proxy for intra-day prices, for a selection of US stocks. Hays et al. (2012) propose a similar model for zero-coupon bond yield curves. Aït-Sahalia and Xiu (2018) study the high-frequency covariance structure of the constituents of the S&P100 index by conducting principal component analysis on a weekly basis.

The remainder of the paper is organised as follows. Section 4.2 outlines the whole theoretical architecture of our framework, which is then exploited for the empirical applications on market data in Section 4.3. In particular, Section 4.3.1 investigates what we call the baseline model, while sections 4.3.2 and 4.3.3 are devoted, respectively, to the market and idiosyncratic volatility branches of our extended model. The prediction performance of the various models is quantitatively assessed in Section 4.4. Finally, Section 4.5 concludes and lays down a few conceivable extensions.

4.2. Model

We first devise in the next subsection a diffusion model with drift for log-returns from which we obtain the distinct measures of market volatility and individual assets' idiosyncratic volatility. In Section 4.2.2 we briefly give some background on functional data analysis and we construct the functional curves describing the aforementioned volatility measures. Finally, we outline in Section 4.2.3 the dimension reduction techniques, based on functional principal component analysis, which define our models for the original volatility trajectories.

4.2.1. A diffusion model with drift for repeated volatility trajectories

Müller et al. (2011) propose the following stochastic diffusion model with drift for log-returns:

$$\mathrm{d}\log\overline{X}_j(t,\omega) = \overline{\mu}_j(t,\omega)\,\mathrm{d}\,t + \overline{\sigma}_j(t,\omega)\,\mathrm{d}\,W_j(t,\omega) \tag{4.2.1}$$

where the arguments $\omega \in \Omega$ highlight the stochastic nature of the underlying processes, $\overline{X}_j(t,\omega)$ denotes the intra-day price of the S&P500 index at time $t \in$

$[0,T]$ on day $j = 1,\dots,J$, $\overline{\mu}_j(t,\omega)$ and $\overline{\sigma}_j(t,\omega)$ are i.i.d. copies of the usual drift and variance processes, and $W_j(t,\omega)$ are independent standard Wiener processes.

In our contribution we extend eq. (4.2.1) in two ways. First, instead of considering uniquely the reference index, we study, together with the index itself, all the distinct securities $i = 1,\dots,I$ the index consists of. Second, we distinguish the *market* component, i.e. the part of the return that correlates with the index, from the *idiosyncratic* component of the return, i.e. the residual part that doesn't correlate with the index.

For every asset i in the index, we estimate the following factor model of individual log-returns (we will drop the explicit dependence on ω hereafter for the only purpose of notational convenience):

$$\mathrm{d}\log X_{ij}(t) = \beta_i \cdot M_j(t) + e_{ij}(t) \tag{4.2.2}$$

from which we filter out the common market component $M_j(t) := \mathrm{d}\log \overline{X}_j(t)$ and obtain the idiosyncratic returns component $e_{ij}(t)$. As already pointed out, the β_i parameters' estimates can be interpreted as the usual CAPM betas. We then apply the same stochastic diffusion model with drift of eq. (4.2.1) to both the market index returns $M_j(t)$ and the residual idiosyncratic returns $e_{ij}(t)$, thereby obtaining $I+1$ distinct measures of volatility. Accordingly, we assume the following augmented stochastic diffusion model with drift:

$$\mathrm{d}\log X_{ij}(t) = \beta_i \cdot \underbrace{\left[\overline{\mu}_j(t)\,\mathrm{d}t + \overline{\sigma}_j(t)\,\mathrm{d}W_j(t)\right]}_{\text{market volatility}} + \underbrace{\mu_{ij}(t)\,\mathrm{d}t + \sigma_{ij}(t)\,\mathrm{d}W_{ij}(t)}_{\text{idiosyncratic volatility}} \tag{4.2.3}$$

As customary in the literature, we further assume that market data is observable at short regular time intervals Δ, justifying the asymptotic assumption $\Delta \to 0$, and we define the generic discretised version of log-returns and diffusion terms as follows:

$$X^{\Delta}_{ijt} = \frac{1}{\sqrt{\Delta}} \log\left(\frac{X_{ij}(t+\Delta)}{X_{ij}(t)}\right) \tag{4.2.4}$$

$$W^{\Delta}_{ijt} = \frac{W_{ij}(t+\Delta) - W_{ij}(t)}{\sqrt{\Delta}} \tag{4.2.5}$$

where the subscript $t = 1,\dots,T/\Delta$ denotes regularly Δ-spaced intra-day times and the market diffusion term W^{Δ}_{jt} is defined as the straightforward single-asset restriction of eq. (4.2.5). The discretised version of eq. (4.2.3) therefore reads:

$$\begin{aligned} X^{\Delta}_{ijt} = {} & \frac{\beta_i}{\sqrt{\Delta}} \cdot \left[\int_t^{t+\Delta} \overline{\mu}_j(v)\,\mathrm{d}v + \int_t^{t+\Delta} \overline{\sigma}_j(v)\,\mathrm{d}W_j(v)\right] \\ & + \frac{1}{\sqrt{\Delta}} \cdot \left[\int_t^{t+\Delta} \mu_{ij}(v)\,\mathrm{d}v + \int_t^{t+\Delta} \sigma_{ij}(v)\,\mathrm{d}W_{ij}(v)\right] \end{aligned} \tag{4.2.6}$$

Following Müller et al. (2011, Lemma 1), under suitable regularity assumptions[3] eq. (4.2.6) can be approximated by

$$X_{ijt}^{\Delta} \approx \beta_i \cdot \overline{\sigma}_{jt} W_{jt}^{\Delta} + \sigma_{ijt} W_{ijt}^{\Delta} \tag{4.2.7}$$

which we redefine for notational convenience as

$$X_{ijt}^{\Delta} = \beta_i \cdot M_{jt}^{\Delta} + e_{ijt}^{\Delta} \tag{4.2.8}$$

Note that eq. (4.2.8) essentially constitutes a discretised form of eq. (4.2.2). The usual smooth generic log-volatility process $V(t) = \log\left[\sigma(t)^2\right]$ can be therefore approximated by the discrete series

$$\overline{V}_{jt}^{\Delta} = \log\left[\left(M_{jt}^{\Delta}\right)^2\right] \tag{4.2.9}$$

for market volatility on trading day j and

$$V_{ijt}^{\Delta *} = \log\left[\left(e_{ijt}^{\Delta}\right)^2\right] \tag{4.2.10}$$

for the $I \times J$ idiosyncratic volatilities.

4.2.2. Functional data analysis

The field of functional data analysis (FDA) provides a collection of statistical techniques that are well suited for analysing processes that are continuous in their very nature and can be sampled at high frequencies (for an introduction we refer the reader to Ramsay and Silverman, 2005). Volatilities of financial securities arguably satisfy these requirements. While volatility itself is not a directly observable phenomenon, it can be approximated, starting from high-frequency (observable) prices, using the aforementioned procedure (see eqs. (4.2.3) to (4.2.10)).

A continuous measure of the underlying volatility processes can then be approximated by a Λth-order Fourier series over the interval $[0,1]$ as follows:

$$\overline{V}_j(t) \approx \overline{\alpha}_j + \sum_{\lambda=1}^{\Lambda} \overline{\gamma}_{j\lambda} \cdot \cos\left(\frac{2\pi\lambda t}{T}\right) + \sum_{\lambda=1}^{\Lambda} \overline{\delta}_{j\lambda} \cdot \sin\left(\frac{2\pi\lambda t}{T}\right) \tag{4.2.11}$$

$$V_{ij}^{*}(t) \approx \alpha_{ij}^{*} + \sum_{\lambda=1}^{\Lambda} \gamma_{ij\lambda}^{*} \cdot \cos\left(\frac{2\pi\lambda t}{T}\right) + \sum_{\lambda=1}^{\Lambda} \delta_{ij\lambda}^{*} \cdot \sin\left(\frac{2\pi\lambda t}{T}\right) \tag{4.2.12}$$

[3] The assumptions on the underlying generic processes $\mu(t)$ and $\sigma(t)$ include uniform Lipschitz continuity of order 1, boundedness, and smoothness and boundedness of derivatives of $\sigma(t)$.

Coefficients α's, γ's and δ's are chosen to minimise the distance with the original discrete series $\overline{V}_{jt}^{\Delta}$ and $V_{ijt}^{\Delta *}$, e.g. by penalised least squares.

The $[(I+1)\times J]$-dimensional collection of $\overline{V}_j(t)$'s and $V_{ij}^*(t)$'s begs for some dimensionality reduction to be analysed properly. Within the FDA toolbox, the functional counterpart of the usual principal component analysis (PCA) comes at hand. While PCA constructs a K-dimensional summary of a p-dimensional random vector, with $p > K$, functional principal component analysis (FPCA) approximates a generic continuous, and thus intrinsically infinite-dimensional real valued curve $Y(t)$ with the finite-dimensional object

$$Y(t) \approx \mu(t) + \sum_{k=1}^{K} \xi_k \cdot \phi_k(t) \tag{4.2.13}$$

where $\mu(t)$ is the mean function of $Y(t)$, $\phi_k(t)$ are the eigenfunctions of $Y(t)$, defined over the same domain, and ξ_k are the associated loading coefficients. This result is based on the Karhunen-Loève theorem which states that a centred continuous stochastic process $Y(t) : [a,b] \to \mathbb{R}$ admits the infinite representation

$$Y(t) = \sum_{k=1}^{\infty} \xi_k \cdot \phi_k(t)$$

where ξ_k are pairwise uncorrelated random variables and $\phi_k(t)$ are continuous real-valued functions on $[a,b]$ and pairwise orthogonal in $L^2([a,b])$.

4.2.3. Functional principal component analysis of volatility curves

When the underlying functional objects consist of Fourier series as in our case (cf. eqs. (4.2.11) and (4.2.12)), FPCA amounts to performing the usual PCA eigen-decomposition on the collection of vectors of Fourier coefficients. For market volatilities $\overline{V}_j(t)$ this correspond to decomposing the 2-dimensional array $\{\overline{\alpha}_j, \overline{\gamma}_{j1}, \ldots, \overline{\gamma}_{j\Lambda}, \overline{\delta}_{j1}, \ldots, \overline{\delta}_{j\Lambda}\}$ and for idiosyncratic volatilities $V_{ij}^*(t)$ the 3-dimensional array $\{\alpha_{ij}^*, \gamma_{ij1}^*, \ldots, \gamma_{ij\Lambda}^*, \delta_{ij1}^*, \ldots, \delta_{ij\Lambda}^*\}$. We set $K = 1$, i.e. we focus uniquely on the *first* functional principal component of each of the volatility measures. While this assumption seems restrictive, the first component is by construction the most informative about the original process and, as we shall show in the following, it is convenient for the subsequent analysis of the common idiosyncratic volatility. The estimate of the market volatility reads:

$$\widehat{\overline{V}}_j(t) = \overline{\mu}(t) + a_j \cdot \overline{\phi}(t) \tag{4.2.14}$$

The eigenfunction $\overline{\phi}(t)$ captures the principal mode of variation of the index volatility and coefficients a_j represent their daily loadings. Except for the $K = 1$ restriction, this estimate is analogous to the one performed in Müller et al. (2011). In a similar fashion, the estimates for the I idiosyncratic volatilities read:

$$\widehat{V}^*_{ij}(t) = \mu^*_i(t) + b_{ij} \cdot \phi^*_i(t) \tag{4.2.15}$$

Here we obtain I principal eigenfunctions, one for each of the index constituents. In order to investigate the common idiosyncratic volatility, we perform a second FPCA on the $\phi^*_i(t)$ principal eigenfunctions, obtaining

$$\widehat{\phi}^*_i(t) = c_i \cdot \psi^*(t) \tag{4.2.16}$$

The new eigenfunction $\psi^*(t)$ should capture, if present, any residual comovement in idiosyncratic volatilities, i.e. in the volatility of single securities after the market component has been filtered out. From eqs. (4.2.15) and (4.2.16) it is possible to estimate the idiosyncratic volatilities $V^*_{ij}(t)$ by using only the (partial) information contained in the common idiosyncratic volatility eigenfunction $\psi^*(t)$:

$$\widehat{\widehat{V}}^*_{ij}(t) = \mu^*_i(t) + (b_{ij} \cdot c_i) \cdot \psi^*(t) \tag{4.2.17}$$

For the sake of completeness, we also estimate the following benchmark model in which we include the information about each single security but without distinguishing the market component from the idiosyncratic components of volatility:

$$\widehat{V}_{ij}(t) = \mu_i(t) + d_{ij} \cdot \phi_i(t) \tag{4.2.18}$$

$$\widehat{\phi}_i(t) = h_i \cdot \psi(t) \tag{4.2.19}$$

$$\widehat{\widehat{V}}_{ij}(t) = \mu_i(t) + (d_{ij} \cdot h_i) \cdot \psi(t) \tag{4.2.20}$$

where the target measure of volatility is the functional counterpart of

$$V^{\Delta}_{ijt} = \log\left[\left(X^{\Delta}_{ijt}\right)^2\right] \tag{4.2.21}$$

In the following section we fit all the aforementioned functional model to real financial data in order to compare, in Section 4.4, their relative performance in capturing the actual trajectories of intra-day volatility.

4.3. Data and empirical application

In this application we estimate the models outlined in the previous section on real data from two separate market indices, the S&P500 and the EURONEXT 100. Both indices list liquid blue chips traded on American and European stock exchanges, respectively. We have available tick data about trades that took place between 24th May and 8th December 2017 for all the constituents of the indices, amounting to 1,559,379,085 data points across 139 trading days for the S&P500 and 109,946,543 data points across 143 days for the EURONEXT 100. A trading day consists of different phases with distinct microstructural pricing mechanisms, such as order accumulation periods with no price announcements at the beginning and end of the trading day, each followed by a batch auction, and an open market session in which continuous double auction pricing is in place, and a new price emerges as soon as a trade is successfully executed. Since our interest is on intra-day volatility, we focus uniquely on the latter phase, lasting $T = 6½$ hours for the S&P500 and $T = 8½$ hours for the EURONEXT 100. We therefore perform some data cleaning on the original dataset along the lines of Gallo and Brownlees (2006) and we single out regular trades from the open market session. We drop two trading days from the S&P500 due to early closing[4] and one from the EURONEXT 100 due to a brief market-wide suspension, along with a few securities that experienced trading suspensions in our reference period, ending up with $I = 502$ stocks (out of 505) and $J = 137$ days for the S&P500 index and $I = 93$ stocks (out of 100) and $J = 142$ days for the EURONEXT 100. Starting from irregularly spaced tick data, we construct regular time series of log-returns at $\Delta = 5$-minute intervals as per eq. (4.2.4), totalling $T/\Delta = 78$ figures per day for the S&P500 and $T/\Delta = 102$ for the EURONEXT 100.[5] These regular time series are shown for later reference in Fig. 4.1. Our chosen measures of volatility are then computed according to eq. (4.2.21) for the baseline model and to eqs. (4.2.9) and (4.2.10) for our extension, after fitting the factor model in eq. (4.2.8) and singling out the idiosyncratic component $e_{ij}^{\Delta}(t)$ from the market component $M_j^{\Delta}(t)$. Since with have available actual observations of the reference market index, i.e. of the value-weighted market portfolio, we fit the linear

[4] Early closing occurred on 3rd July (the day before Independence day) and on 24th November (the day after Thanksgiving).

[5] Following Aït-Sahalia et al. (2005) a 5-minute interval seems a good choice and is the *de facto* standard in the literature. Note however that the time series will undergo substantial smoothing when constructing the relevant functional objects and any residual microstructure noise is likely to be lost in the process.

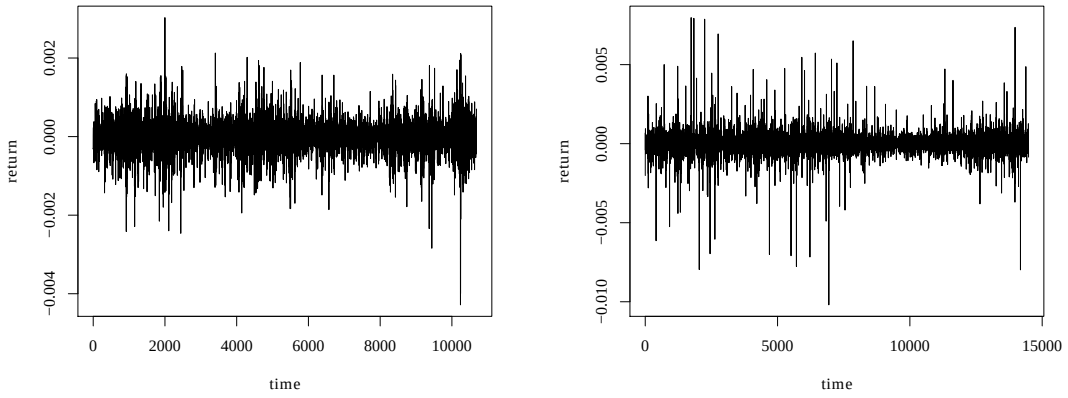

Figure 4.1.: 5-minute log-returns of the S&P500 (left) and EURONEXT 100 (right) indices.

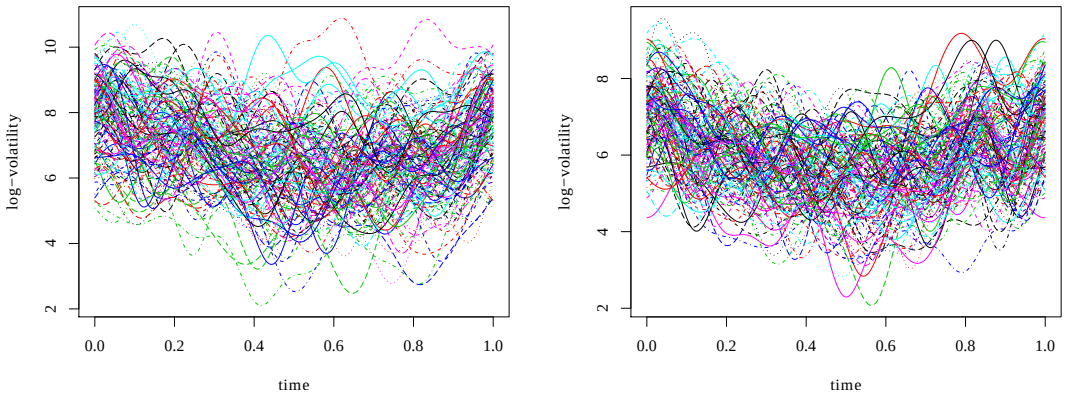

Figure 4.2.: Intra-day volatility curves $\overline{V}_j(t)$ of the index for all available trading days j; S&P500 (left) and EURONEXT 100 (right).

model in eq. (4.2.8) using these data as the independent variable. The obtained series serve as the input of our functional data analysis.

To construct the functional counterpart of the volatility series, we fit a $\Lambda = 7$-th order Fourier series (i.e consisting of 1 constant term and 14 sines and cosines Fourier bases) to the empirical data by means of penalised least squares, as implemented in the fda package[6] (see also Ramsay et al., 2009). An identical procedure is also followed by Hörmann et al. (2014). Fig. 4.2 shows the obtained functional volatility trajectories $\overline{V}_j(t)$ for all the available trading days $j = 1, \ldots, J$ of the S&P500 (left panel) and the EURONEXT 100 (right panel) indices. Although the picture looks confusing at a first glance since a large multitude of curves are plotted at once (137 and 142, respectively), a clear tendency is apparent. Most of the trajectories overlap in a U-shaped fashion, suggesting

[6] https://cran.r-project.org/package=fda

that volatility is higher during early and late hours, i.e. in the vicinity of market opening and closing, and lower around halfway in the trading day. Such a seasonal effect is well known in the financial literature and constitutes a fairly robust 'stylised fact' (see e.g. Jain and Joh, 1988; Lockwood and Linn, 1990). Another feature that becomes evident at a second glance is that the U-shaped property holds on average but does not apply to every single trading day. Many 'outliers' can be identified in both markets, displaying crests and troughs at virtually any time throughout the trading session. These are likely related to the timing of release of fundamental news about the constituents of the indices. It is worth noting however that our choice of Fourier bases for constructing the functional trajectories, as opposed to e.g. kernel-based methods, emphasises the oscillatory motion of the curves and their peaks therein.

The remainder of this section is organised as follows. In the next subsection the baseline model is fitted to data for the sake of comparison with our extended model, which is estimated separately in its market volatility component (Section 4.3.2) and idiosyncratic volatility component (Section 4.3.3).

4.3.1. Baseline model

The baseline model exploits all the available information, in the sense that all the data concerning single constituents of the index is included. However, we don't single out the underlying market and idiosyncratic volatilities, and therefore these measures remain essentially confounded. Following the principal component decomposition in eq. (4.2.18), we proceed by identifying the mean functions $\mu_i(t)$ and the eigenfunctions $\phi_i(t)$, one for each of the constituents i of the index, and the loading coefficients d_{ij} of these latter, one for every asset i and for every trading day j. The curves in Fig. 4.3 denote the individual constituents' mean functions. As expected, the U-shaped tendency mentioned in the previous section is evident and applies on average to all securities. With respect to the S&P500 (left panel) and with the exception of a single outlier which experienced a substantially lower volatility throughout our reference period (corresponding to News Corp, an American entertainment and mass media company), the overall bundle of trajectories looks quite compact, implying a roughly uniform distribution of average volatility among the various assets. Differently, the curves relative to the EURONEXT 100 appear more disperse, suggesting a more fat-tailed distribution thereof. The thin lines in Fig. 4.4 picture all the relevant principal eigenfunctions $\phi_i(t)$. By construction the eigenfunctions are defined up to their sign; this in general brings up an identification problem. However,

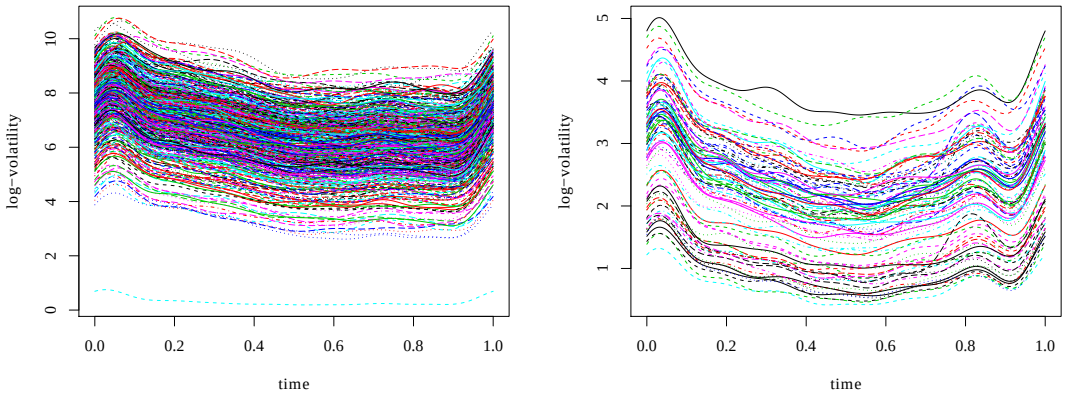

Figure 4.3.: Mean functions $\mu_i(t)$ obtained by estimating the baseline model; S&P500 (left) and EURONEXT 100 (right).

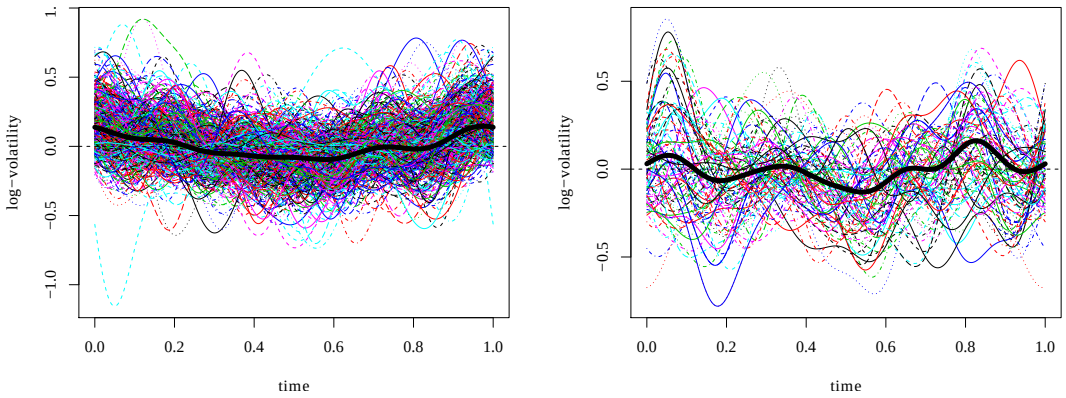

Figure 4.4.: Eigenfunctions $\phi_i(t)$ obtained by estimating the baseline model and their principal eigenfunction $\psi(t)$ (thick curve); S&P500 (left) and EURONEXT 100 (right).

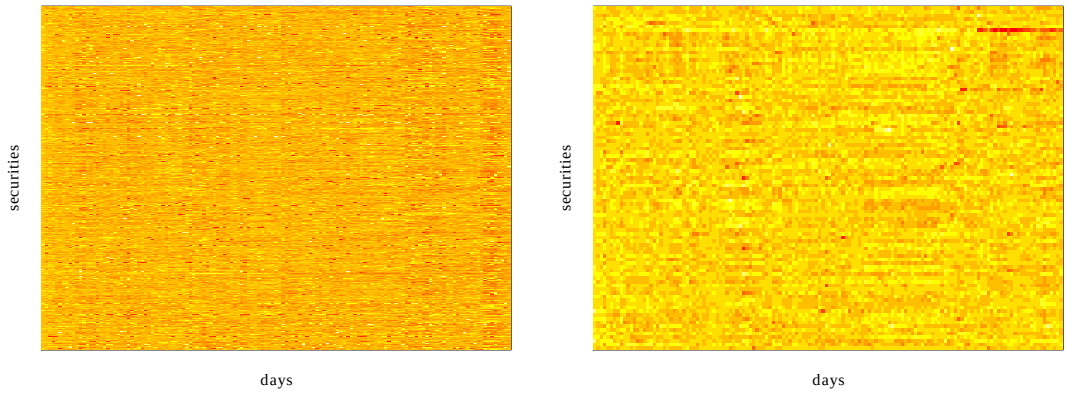

Figure 4.5.: Heatmap of the loading coefficients d_{ij} within the baseline model; S&P500 (left) and EURONEXT 100 (right). Darker shades denote larger magnitudes.

since we are only interested in the first principal eigenfunction of every security and given that the original volatility trajectories are defined in $\mathbb{R}_{++}$, we keep untouched those eigenfunctions whose constant term (of the underlying Fourier bases) is positive and swap the sign of those with a negative constant term. At a first glance, the U-shaped pattern recognised in the market volatility curves (Fig. 4.2) is still present, meaning that the mean curves alone are not sufficient to capture the whole underlying motion. This is especially true for the S&P500 index. The loading coefficients d_{ij} are represented by the heatmap in Fig. 4.5, where darker colours denote larger magnitudes. The picture for the S&P500 (left panel) looks quite uniform, although some vertical shades can be identified, corresponding to brief periods of market turbulence also visible in Fig. 4.1 (cf. e.g. the rightmost part of the left panel). Note however that Fig. 4.1 portrays high-frequency 5-minute returns, while each pixel of the heatmap in Fig. 4.5 corresponds to one trading day, implying that days exhibiting an outstanding peak in the returns series need not be necessarily more volatile all in all. The picture for the EURONEXT 100 (right panel) instead contains some horizontal structure, suggesting that, after the mean is taken into account, there are securities that are systematically more volatile (darker rows) than others (paler rows), at least for certain prolonged periods. The most visible example of this is given by the reddish segment in the north-east corner, denoting an asset (Altice, a Netherlands-based multinational telecoms company) that consistently experiences a substantially higher volatility for about a month in a row.

We proceed by further reducing the (still) high-dimensional collection of eigenfunction in Fig. 4.4 in order to grasp their main underlying dynamics. We

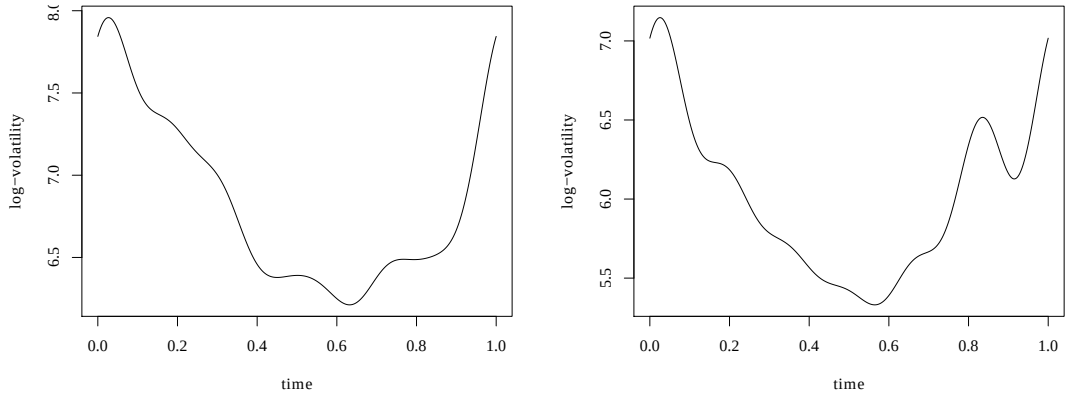

Figure 4.6.: Mean function $\overline{\mu}(t)$ obtained by estimating the market volatility model; S&P500 (left) and EURONEXT 100 (right).

thus proceed with the second step of our procedure, namely another functional principal component decomposition, according to eq. (4.2.19). In this way we are able to identify a single eigenfunction $\psi(t)$ summarising the principal mode of variation of the I previous eigenfunctions. The principal eigenfunction $\psi(t)$, represented as the thick line in Fig. 4.4, captures in a single curve what the mean functions $\mu_i(t)$ do not. Again, the usual U-shaped motif is present. Given the lower number of constituents (93 vs. 502) of the EURONEXT 100 index (right panel), their picture looks more noisy than for the S&P500 (left panel). It remains true, however, that in both markets the lowest trough falls around halfway in the trading day and the two highest crests happen in the vicinity of market opening and market closing.

4.3.2. Market volatility

The market volatility decomposition in eq. (4.2.14) pinpoints a single couple of curves, namely the mean function $\overline{\mu}(t)$ and the eigenfunction $\overline{\phi}(t)$, together with a vector of loading coefficients a_j of the latter, from the collection of intra-day volatility curves already shown in Fig. 4.2. As expected, the mean function, shown in Fig. 4.6, shares the U-shaped dynamics of the market volatility trajectories in Fig. 4.2. The noticeable spike on the right branch of the EURONEXT 100 curve (right panel) is most likely related to the American market's morning opening. The principal eigenfunction $\overline{\phi}(t)$, pictured in Fig. 4.7, shares with the mean function the fact of displaying its lowest point halfway throughout the trading day, but exhibits less of a 'U' shape. Given the lower dimensionality of the market volatility model, it is feasible to visualise the predicted volatil-

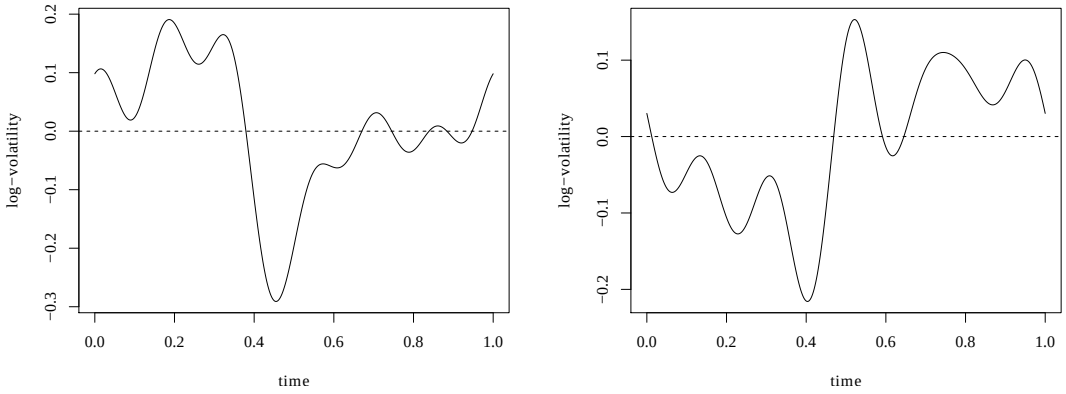

Figure 4.7.: Eigenfunction $\overline{\phi}(t)$ obtained by estimating the market volatility model; S&P500 (left) and EURONEXT 100 (right).

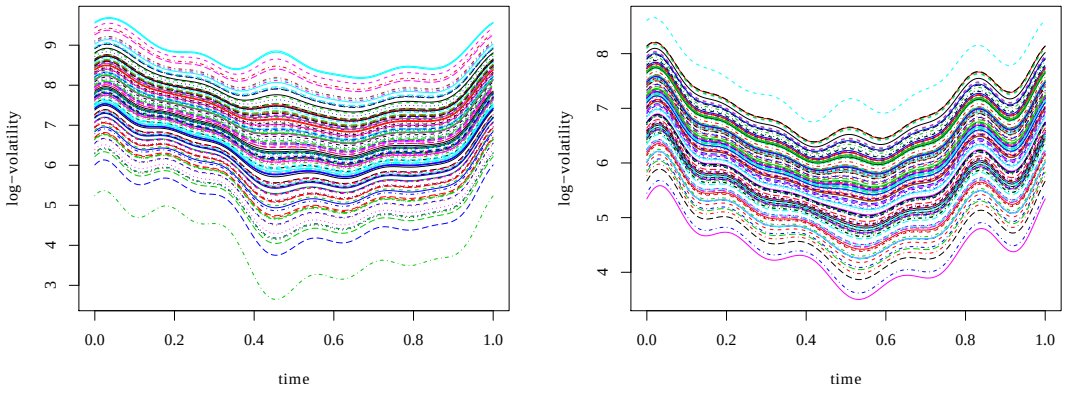

Figure 4.8.: Predicted market volatility trajectories $\widehat{\overline{V}}_j(t)$ by the market volatility model; S&P500 (left) and EURONEXT 100 (right).

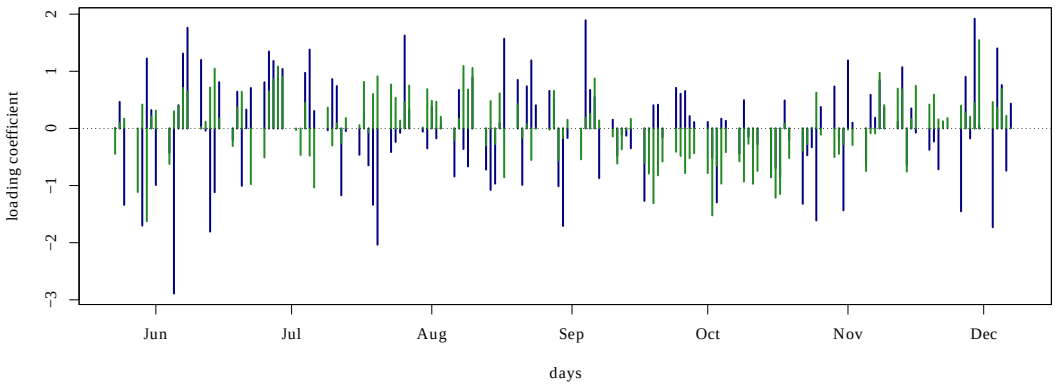

Figure 4.9.: Market volatility loading coefficients a_j; S&P500 (blue) and EURONEXT 100 (green).

ity trajectories in a single picture, reported in Fig. 4.8. Following eq. (4.2.14) the predicted $\widehat{\overline{V}}_j(t)$'s appear as parallel displacements of the eigenfunction $\overline{\phi}(t)$, whose magnitude for each day j is given by the corresponding loading coefficient a_j, summed to the mean function $\overline{\mu}(t)$ reported in Fig. 4.6. Higher (respectively, lower) curves represent trajectories with a larger (smaller) overall intra-day volatility. For the S&P500 (left panel) one curve visibly outlies the bundle from below, corresponding to a day (specifically, 6th June 2017) of exceptionally low volatility at the market index level. Conversely, the EURONEXT 100 picture (right panel) displays an outlier above the stack, coinciding with a day of exceptionally high volatility (specifically, 1st December 2017). Such extreme values are better appreciable in Fig. 4.9, reporting the series of (dated) loading coefficients a_j for the two indices. This picture also reveals some interesting insights about the evolution of market volatility over both time and space. For instance, it is clear that for the EURONEXT 100 (green bars) there exist prolonged periods in which the coefficients are either greater than zero (e.g. from mid-July to early August) and hence displaying a consistently higher than average volatility, or less than zero (e.g. from mid-September to mid-October), in which volatility is consistently below average. This pattern reflects the well known volatility clustering property of financial markets, in which the signs of future returns are not readily predictable although their magnitudes are, and tend to cluster in time, giving rise to prolonged periods of low volatility followed by periods of high volatility (see e.g. Andersen and Bollerslev, 1997; Mandelbrot, 1963). To a lesser extent, i.e. for less consecutive days, this is also true for the S&P500 (e.g. between late June and early July). This shouldn't come as a surprise since the American stock market is generally regarded as being relatively more liquid

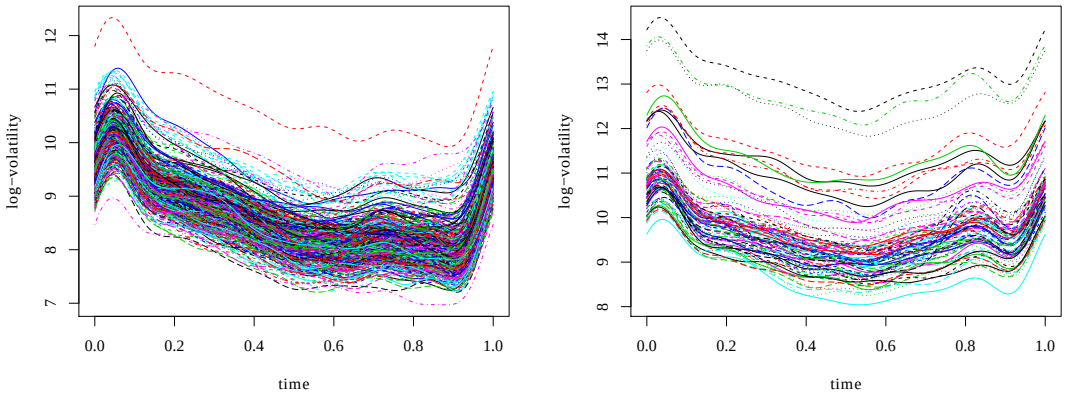

Figure 4.10.: Mean functions $\mu_i(t)$ obtained by estimating the idiosyncratic volatility model; S&P500 (left) and EURONEXT 100 (right).

than its European counterpart. Fig. 4.9 also allows a preliminary international comparison of the two financial markets. Specifically, there are periods in which the loading coefficients for the two indices share the same sign and periods in which they don't. In the first, the American and European stock markets act as complements, while in the second they act as substitutes for volatility.

4.3.3. Idiosyncratic volatility

The second part of our extended model only exploits information that concerns the idiosyncratic component of returns, i.e. after the market return is filtered out of each constituent of the index, following eq. (4.2.8). We estimate a functional principal component decomposition, analogous to the one carried out in the baseline model (see Section 4.3.1), on the idiosyncratic volatility measures $V^*_{ij}(t)$ defined in eq. (4.2.10). Following eq. (4.2.15) we obtain the mean functions $\mu^*_i(t)$ and eigenfunctions $\phi^*_i(t)$, one for each of the constituents i of the index, and their loading coefficients b_{ij}, one for every asset i and for every trading day j. The mean functions, pictured in Fig. 4.10 still retain the U-shaped signature for both indices. The S&P500 (left panel) displays an outlier standing above the remaining, relatively uniform, bundle of curves (corresponding to Newmont Mining Corporation, an American gold mining company). The EURONEXT 100 (right panel) presents a more dispersed structure, similar to Fig. 4.3. The range of the eigenfunctions, pictured as thin lines in Fig. 4.11, has expectedly shrunk with respect to the baseline counterpart (cf. Fig. 4.4) since a sizeable part of the overall volatility, namely the market volatility captured by the CAPM betas of eq. (4.2.2), has been excluded. Differently from the eigen-

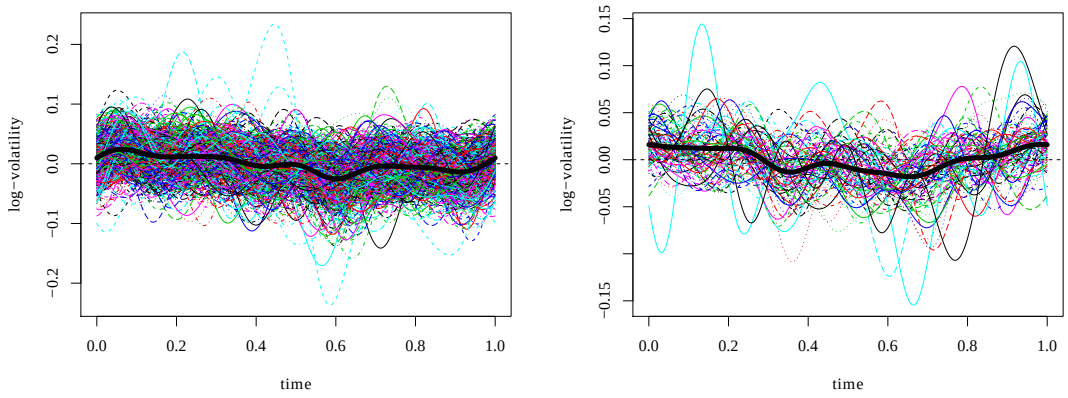

Figure 4.11.: Eigenfunctions $\phi_i^*(t)$ obtained by estimating the idiosyncratic volatility model and their principal eigenfunction $\psi^*(t)$ (thick curve); S&P500 (left) and EURONEXT 100 (right).

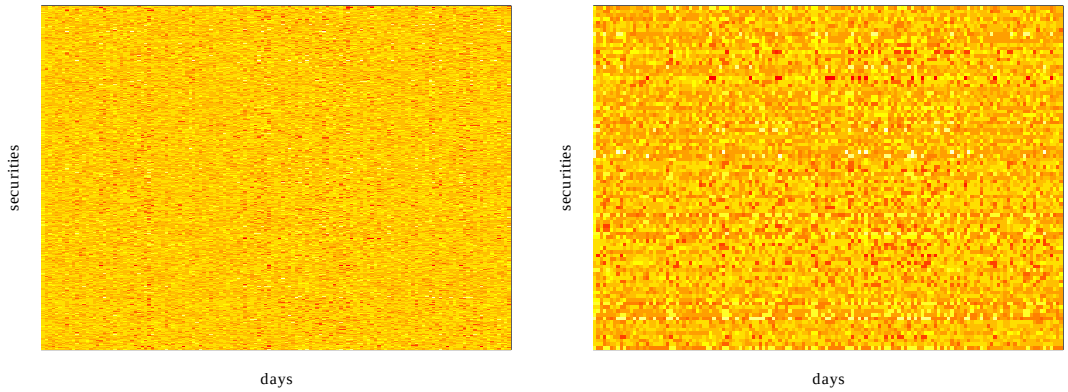

Figure 4.12.: Heatmap of the loading coefficients b_{ij} within the idiosyncratic volatility model; S&P500 (left) and EURONEXT 100 (right). Darker shades denote larger magnitudes.

functions of the baseline model in Fig. 4.4, the U-shaped pattern is much less visible to the naked eye and a few outliers stand out, suggesting that a handful of constituents display a markedly different dynamics with respect to majority. The loading coefficients b_{ij} are represented by the heatmap in Fig. 4.12. As expected, most of the vertical structure present in Fig. 4.5 and capturing variation across assets has vanished. Although the heatmap looks visually closer to random noise for both indices, not all of the common variation has disappeared. The second step of our procedure, involving the functional principal component decomposition of the idiosyncratic eigenfunctions $\phi_i^*(t)$, aims specifically at eliciting this residual common part of idiosyncratic volatility trajectories. Proceeding along the same lines of the baseline model, we obtain a single principal

eigenfunction $\psi^*(t)$ representing the common idiosyncratic volatility of the constituents of the indices. Represented as the thick line in Fig. 4.11, it still retains a good deal of the U-shaped motion that is present at the market level, in the sense that in both markets the lowest trough falls around halfway in the trading day and the two highest crests happen in the vicinity of market opening and closing. This also implies that the linear CAPM residuals of eq. (4.2.2) display some correlation with the regressor. In spite of being the result of a 2-step procedure that, by construction, throws a lot of information away at every step, we show in the next section that the CIV curve $\psi^*(t)$ retains a very valuable extent of information about the individual volatility trajectories $V_{ij}(t)$ which can be successfully exploited for the sake of their prediction.

4.4. Relative performance of the different models

Each of the models discussed in sections 4.3.1 to 4.3.3 aims at reducing the dimensionality of an originally large dataset comprising multiple daily volatility trajectories (as in the market volatility model), possibly for each of the constituents of a market index (as in the baseline and the idiosyncratic volatility models). In doing so, it singles out a much smaller set of curves, consisting of mean functions and eigenfunctions, that attempt to incorporate most of the information contained in the original trajectories. To test the relative performance of the various models, we set up a number of linear regressions whose dependent variable consist of the original volatility trajectories of single constituents of the market index, and the regressors are given by a combination of the underlying mean functions and eigenfunctions, added in an incremental fashion. We estimate each regression for every constituent asset i and for every available trading day j. From each experiment we obtain $i \times j$ adjusted R^2 coefficients, i.e. totalling 68,774 for the S&P500 and 13,206 for the EURONEXT 100, of which we compute averages across days and compare the resulting empirical distribution.

For the baseline model (see Section 4.3.1) we estimate the following linear regression experiments:

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \mu_i(t) + \varepsilon_{ij}(t) \tag{4.4.1}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \mu_i(t) + \nu_{2ij} \cdot \phi_i(t) + \varepsilon_{ij}(t) \tag{4.4.2}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \mu_i(t) + \nu_{2ij} \cdot \psi(t) + \varepsilon_{ij}(t) \tag{4.4.3}$$

The first experiment in eq. (4.4.1) only includes the assets' mean functions $\mu_i(t)$; in the second and third regressions we add either the principal eigenfunctions

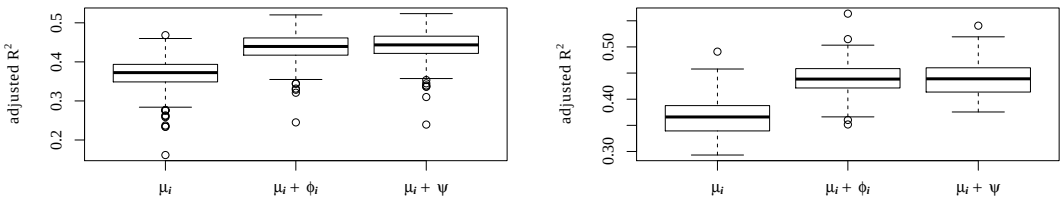

Figure 4.13.: Prediction performance of the baseline model; S&P500 (left) and EURONEXT 100 (right).

$\phi_i(t)$ (eq. (4.4.2)) or their common principal eigenfunction $\psi(t)$ (eq. (4.4.3)). Fig. 4.13 reports the boxplot of the adjusted R_i^2 coefficients in each of the regression experiments (underlying regressors are specified on the horizontal axis). It is clear that, while the mean functions alone constitute good predictors of the original volatility trajectories, explaining on average about 37.00% of the overall variance in the S&P500 and 36.59% in the EURONEXT 100 (these and subsequent numerical figures about adjusted R^2 coefficients are rounded up to their fourth decimal digit), adding the principal eigenfunctions $\phi_i(t)$ to the equation gives an edge in terms of performance, increasing the explained variance to, respectively, about 43.73% and 43.85%. This is not surprising since, by construction, the infinite series of eigenfunctions (of which by assumption we include only the first) capture the residual part of information that is not embedded into the mean function. The second fact discernible from Fig. 4.13 is that including the single principal eigenfunction $\psi(t)$ in place of the I eigenfunctions $\phi_i(t)$ does not make the prediction considerably better or worse (the average adjusted R^2 coefficient is 0.4423 for the S&P500 and 0.4392 for the EURONEXT 100). This means that a single curve, the second-order principal eigenfunction $\psi(t)$, captures virtually all the information carried by the multiple ϕ's. It also appears that the lower dimensional eq. (4.4.3) scores slightly better than its higher-dimensional counterpart (eq. (4.4.2)).

Regarding our extended model, we first evaluate the performance of each of its branches, namely the market volatility and idiosyncratic volatility models, and then we test their explanatory power in conjunction. For the lower-dimensional market volatility model (see Section 4.3.2) we fit the following two regressions

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \overline{\mu}(t) + \varepsilon_{ij}(t) \tag{4.4.4}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \overline{\mu}(t) + \nu_{2ij} \cdot \overline{\phi}(t) + \varepsilon_{ij}(t) \tag{4.4.5}$$

The first experiment in eq. (4.4.4) only includes the market's mean function $\overline{\mu}(t)$,

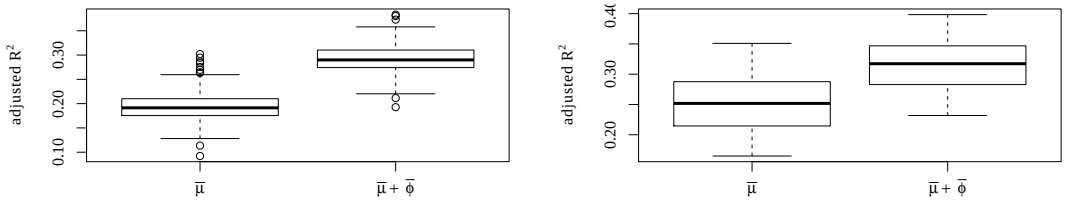

Figure 4.14.: Prediction performance of the market volatility model; S&P500 (left) and EURONEXT 100 (right).

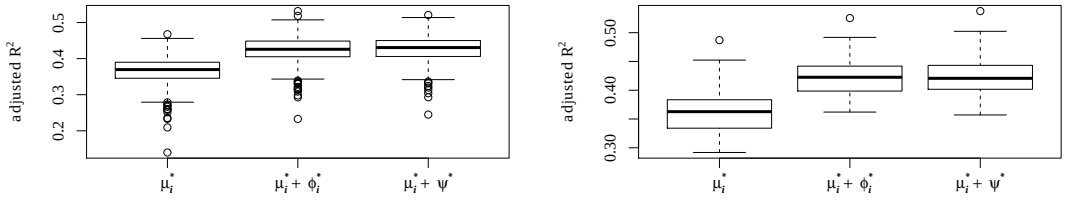

Figure 4.15.: Prediction performance of the idiosyncratic volatility model; S&P500 (left) and EURONEXT 100 (right).

while the second experiment in eq. (4.4.5) also includes the principal eigenfunction $\overline{\phi}(t)$. The boxplots in Fig. 4.13, as expected, suggest that including both curves into the equation gives a better prediction, explaining about 29.15% of the individual volatility trajectories for the S&P500 (left panel) and 31.74% for the EURONEXT 100 (right panel), as opposed to, respectively, 19.44% and 25.30% for the model that exploits the mean function alone. Note that in these experiments we are only using aggregate information about the index to predict the volatility curves of its constituents and that a lower R^2 coefficient is therefore to be expected with respect to the baseline model.

The linear regressions that we use to test the idiosyncratic volatility model (see Section 4.3.1) are the 'starred' analogous of those used for the baseline model (cf. eqs. (4.4.1) to (4.4.3)):

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \mu_i^*(t) + \varepsilon_{ij}(t) \tag{4.4.6}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \mu_i^*(t) + \nu_{2ij} \cdot \phi_i^*(t) + \varepsilon_{ij}(t) \tag{4.4.7}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \mu_i^*(t) + \nu_{2ij} \cdot \psi^*(t) + \varepsilon_{ij}(t) \tag{4.4.8}$$

The usual boxplots, reported in Fig. 4.15,suggest a very similar picture to the baseline ones in Fig. 4.13, both qualitatively and quantitatively. The interpretation is identical, so we just report the relevant numerical figures: from left to right, the average adjusted R^2 coefficients read 0.3667, 0.4250 and 0.4273 for the

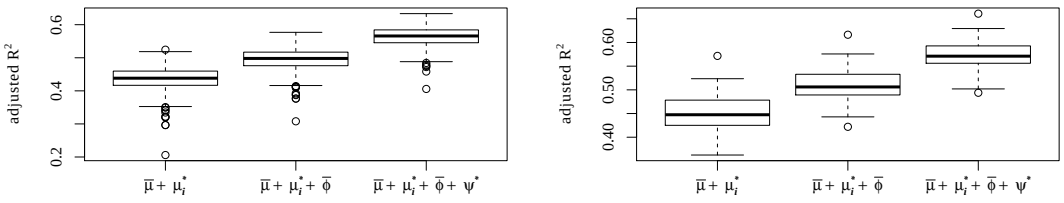

Figure 4.16.: Prediction performance of the whole extended model; S&P500 (left) and EURONEXT 100 (right).

S&P500 (left panel), and 0.3617, 0.4214 and 0.4225 for theEURONEXT 100 (right panel).

We receive a twofold message from Figs. 4.13 to 4.15. First, that in all three instances there is a good deal of information about the original volatility trajectories that is not captured by their mean functions and that including even only the first-order eigenfunctions sensibly increases the performance of the predictions. Second, with respect to Figs. 4.13 and 4.15, that the information contained in the I first-order eigenfunctions can be captured, essentially in its entirety, by a single curve, thereby enabling to further reduce the dimensionality of the problem at virtually no loss.

As a last step, we merge together the market and the idiosyncratic volatility models and estimate the following regression featuring the functional curves obtained therein in a bottom up fashion:

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \overline{\mu}(t) + \nu_{2ij} \cdot \mu_i^*(t) + \varepsilon_{ij}(t) \tag{4.4.9}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \overline{\mu}(t) + \nu_{2ij} \cdot \mu_i^*(t) + \nu_{3ij} \cdot \overline{\phi}(t) + \varepsilon_{ij}(t) \tag{4.4.10}$$

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \overline{\mu}(t) + \nu_{2ij} \cdot \mu_i^*(t) + \nu_{3ij} \cdot \overline{\phi}(t) + \nu_{4ij} \cdot \psi^*(t) + \varepsilon_{ij}(t) \tag{4.4.11}$$

Fig. 4.15 reports the corresponding boxplots. The linear model in eq. (4.4.9), including the market mean function $\overline{\mu}(t)$ and the I idiosyncratic mean functions $\mu_i^*(t)$, explains 43.58% of total variation in the S&P500 and 45.08% in the EURONEXT 100 (right panel). These figures are greater than their counterparts in the two separate market and idiosyncratic volatility models, meaning that both curves are useful in conjunction to predict the volatility trajectories for each asset i. On top of that, eq. (4.4.10) includes the market principal eigenfunction $\overline{\phi}(t)$ and scores a better performance of respectively 49.49% and 51.06% explained variance at the cost of adding a single curve. Finally, eq. (4.4.11) assesses the performance of the whole extended model by adding the second-order principal eigenfunction $\psi^*(t)$. The explained variance, 56.34% for the S&P500 and

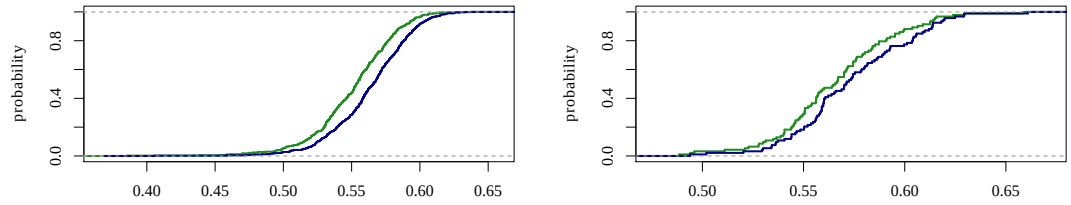

Figure 4.17.: Empirical cumulative density of the adjusted R^2 coefficients distribution of regressions in eq. (4.4.11) (navy) and of analogous regressions featuring $\phi_i^*(t)$ in place of $\psi^*(t)$ (green); S&P500 (left) and EURONEXT 100 (right).

57.38% for the EURONEXT 100 is the highest of all the tested models and displays a substantial advantage with respect to the baseline model in spite of exploiting the very same initial information.[7]

For the sake of completeness, we also estimate the following linear model featuring the I principal eigenfunctions $\phi_i^*(t)$ in place of the single second-order principal eigenfunction $\psi^*(t)$:

$$V_{ij}(t) = \nu_{0ij} + \nu_{1ij} \cdot \overline{\mu}(t) + \nu_{2ij} \cdot \mu_i^*(t) + \nu_{3ij} \cdot \overline{\phi}(t) + \nu_{4ij} \cdot \phi_i^*(t) + \varepsilon_{ij}(t) \quad (4.4.12)$$

Fig. 4.17 compares the empirical CDF of the adjusted R^2 distribution from eq. (4.4.11) (blue) with that of eq. (4.4.12) (green). The distributions average read, respectively, 0.5513 for the S&P500 (left panel) and 0.5659 for the EURONEXT 100 (right panel). Somewhat counter-intuitively, it appears that the performance of the lower dimensional regression featuring the CIV eigenfunction $\psi^*(t)$ first-order stochastically dominates, in both markets, that of the extended model with the individual eigenfunctions $\phi_i^*(t)$.

4.5. Concluding remarks

We investigate the intra-day volatility trajectories of the single constituents of two major international financial markets, namely the S&P500 and the EURONEXT 100 indices. The incumbent literature has so far been addressing one security (or market index) at a time. We therefore extend the theoretical

[7] Strictly speaking, our extended model exploits more information with respect to the baseline since its market volatility branch is based on actual data regarding the market index. However, in a separate experiment which we don't report here, we constructed the market volatility series M_{jt}^{Δ} using the first five principal components of the cross section of returns as done in Herskovic et al. (2016) obtaining substantially analogous results.

framework and develop the necessary methodology, based on recent advances in the functional data analysis field, in order to study an arbitrarily large collection of assets in a joint fashion. We devise a 2-step dimension-reduction procedure based on functional principal component analysis that aims at singling out a small set of curves, consisting of mean functions and eigenfunctions, that attempt to incorporate most of the information contained in the original large database of individual intra-day trajectories. We find that the loss of information linked to the consecutive application of principal component analysis is negligible and that our multi-asset upgrade substantially improves the prediction performance of the single-asset model. On top of this, we propose an alternative model based on a CAPM-inspired distinction between market volatility, i.e. the volatility of the index or of the underlying market portfolio, and the residual, asset-specific idiosyncratic volatility. We find that this extended model outperforms the original model in predicting the individual volatility trajectories, and that the contribution of the common idiosyncratic volatility curve, i.e. the curve that synthesise the principal mode of variation among all idiosyncratic volatilities, is not negligible thereof. In all our instances, we also find that the mean function alone is not an optimal predictor and that adding the first functional principal component to the equation, even discarding higher orders, brings a substantial improvement.

We believe that our framework could be fruitfully extended in a number of directions. First, the second step of our 2-step procedure only tackles the dimension reduction of a set of eigenfunctions, one for every underlying asset, to a single curve, while we make no attempt to likewise reduce the collection of mean functions resulting from the first step. Expanding our procedure to include such a reduction for the means would enable the prediction of an arbitrarily large set of original trajectories using a constant number of curves, at a cost in terms of prediction performance which is worth investigating. Second, in our study we neglect the possibility of serial dependence in the repeated observations of intra-day volatility trajectories. As recognised by Hörmann and Kokoszka (2010), while (weak) dependence does not prevent, in performing the functional principal component analysis, the consistent estimation of the eigendecomposition of a covariance operator, it fails to take into account the potentially very valuable information that is carried by the past values of the functional observations under study. State-of-the-art techniques such as dynamic functional principal component analysis (see e.g. Hörmann et al., 2014), based on a frequency domain approach, specifically target time series exhibiting some

serial dependence structure. Upgrading our 2-step procedure along these lines is not straightforward, although it would likely improve the prediction performance of our volatility model.

This page intentionally left blank.

Appendix

A

Mathematical proofs

A.1. Proof of Proposition 2.1

Recall, first of all, that a null dividend yield is ruled out by our phase space, and therefore the trivial fixed point $(0,0)$ is not applicable. By lagging and substituting the first equation into the second, it is easy to further reduce system (2.3.1) to a 1-dimensional map $\tilde{f}^{\mathfrak{C}} : \mathbb{R}_{++} \to \mathbb{R}_{++}$ solely in terms of the dividend yield:

$$e_t = \tilde{f}^{\mathfrak{C}}(e_{t-1}) = e_{t-1}\frac{1-\overline{x}}{1+\overline{x}(e_{t-1}-1)}(1+g). \tag{A.1.1}$$

Map (A.1.1) admits a unique non-trivial fixed point

$$e^* = g\frac{1-\overline{x}}{\overline{x}}. \tag{A.1.2}$$

Substituting e^* into the first equation of system (2.3.1) yields the equilibrium return $r^* = g$. It is also possible to show that:

- ◇ map $f^{\tilde{\mathfrak{C}}}(\cdot)$ is strictly increasing over $\mathbb{R}_{++}$ since

$$(f^{\tilde{\mathfrak{C}}})'(\cdot) = \frac{(\overline{x}-1)^2}{[1+\overline{x}(e_{t-1}-1)]^2}(1+g) > 0; \tag{A.1.3}$$

- ◇ the slope of $f^{\tilde{\mathfrak{C}}}$ at the fixed point lies within the unit circle since

$$(f^{\tilde{\mathfrak{C}}})'(e^*) = \frac{1}{1+g} \in (0,1). \tag{A.1.4}$$

Therefore, fixed point $\tilde{\mathfrak{C}} = \left(g\,,\, g\,\frac{1-\overline{x}}{\overline{x}}\right)$ is the unique, globally stable fixed point of the system.

□

A.2. Proof of Proposition 2.2

Recall, first of all, that a null dividend yield is ruled out by our phase space, and therefore the trivial (deterministic) fixed point $(0,0)$ is not applicable. We shall show that the 1-dimensional map (2.3.10) admits a unique globally stable random fixed point. In fact, defining

$$e^u = g\frac{1-x^u}{x^u} \quad \text{and} \quad e^d = g\frac{1-x^d}{x^d}, \tag{A.2.1}$$

as in (2.3.9), it holds

$$e^u = f^{\tilde{S}}_{du}(e^d) = f^{\tilde{S}}_{uu}(e^u), \tag{A.2.2}$$

$$e^d = f^{\tilde{S}}_{ud}(e^u) = f^{\tilde{S}}_{dd}(e^d). \tag{A.2.3}$$

Therefore, the state $e^{\tilde{S}}(x_{-1},x) = g\,\dfrac{1-x_{-1}}{x_{-1}}$ is a random fixed point of map (2.3.10). To show uniqueness and stability we shall use the following properties:

- ◇ for every couple $(x_{t-2},x_{t-1}) \in \bigtimes_2\{x^u,x^d\}$ map $f^{\tilde{S}}_{x_{t-2},x_{t-1}}$ is both strictly increasing,

$$(f^{\tilde{S}}_{x_{t-2},x_{t-1}})'(\cdot) = \frac{x_{t-2}(x_{t-2}-1)(x_{t-1}-1)}{x_{t-1}\left[1+x_{t-2}(e_{t-1}-1)\right]^2}(1+g) > 0, \tag{A.2.4}$$

  and concave over $\mathbb{R}_{++}$,

$$(f^{\tilde{S}}_{x_{t-2},x_{t-1}})''(\cdot) = -\frac{2(x_{t-2})^2(1-x_{t-2})(1-x_{t-1})}{x_{t-1}\left[1+x_{t-2}(e_{t-1}-1)\right]^3}(1+g) < 0. \tag{A.2.5}$$

- ◇ for every couple $(x_{t-2},x_{t-1}) \in \bigtimes_2\{x^u,x^d\}$

$$(f^{\tilde{S}}_{x_{t-2},x_{t-1}})'(\cdot)\Big|_{e=e^{x_{t-2}}} = \frac{1}{1+g} \in (0,1), \tag{A.2.6}$$

  where $e^{x_{t-2}} = e^u$ when $x_{t-2} = x^u$ and $e^{x_{t-2}} = e^d$ when $x_{t-2} = x^d$.

Uniqueness Given monotonicity and concavity, there exists only one $\bar{e}$ such that $f_{uu}(\bar{e}) = \bar{e}$. It follows that the random fixed point is unique.

Global stability For every t and for every realisation of the Markov process $\{x_t\}$ the composition of maps

$$f^{\tilde{\mathfrak{S}}}_{x_{t-2},x_{t-1}} \circ \cdots \circ f^{\tilde{\mathfrak{S}}}_{x_0,x_1} \tag{A.2.7}$$

is monotone and concave because it is the composition of monotone and concave maps. Moreover, due to (A.2.2), (A.2.3), and (A.2.6)

$$(f^{\tilde{\mathfrak{S}}}_{x_{t-2},x_{t-1}} \circ \ldots \circ f^{\tilde{\mathfrak{S}}}_{x_0,x_1})'\Big|_{e=e^{x_0}} = \left(\frac{1}{1+g}\right)^t. \tag{A.2.8}$$

It follows that for all initial dividend yields in the open interval $(0, e^d)$ the dynamics converges to $e^{\tilde{\mathfrak{S}}}$ from below, whereas for all initial dividend yields in the open interval (e^d, ∞), the dynamics converges to $e^{\tilde{\mathfrak{S}}}$ from above. Note that the convergence does not depend on $\mathbb{P}$.

Having proved global stability of $e^{\tilde{\mathfrak{S}}}$ for map $\tilde{f}^{\mathfrak{S}}$ in (2.3.10), global stability of $\tilde{\mathfrak{S}}$ for map $\tilde{\mathcal{F}}^{\mathfrak{S}}$ in (2.3.5) follows from the implications of the dividend yield dynamics on the dynamics of returns in the first equation of system (2.3.5).

□

A.3. Proof of Proposition 2.3

Fixed points $\mathfrak{C}$ and $\mathfrak{S}$ are straightforward generalisations of fixed points $\tilde{\mathfrak{C}}$ and $\tilde{\mathfrak{S}}$ derived in Proposition 2.1 and Proposition 2.2 when border conditions $\varphi^* = 0$ and $\varphi^* = 1$ are, respectively, imposed. It suffices to show that there exists no fixed point other than the aforementioned. Hence, consider the case $\varphi^* \in (0,1)$: the first equation of system (2.2.15) reduces to

$$x_{-1} = \varphi^* x_{-1} + (1-\varphi^*)\overline{x}. \tag{A.3.1}$$

Under Assumption 2.2 this condition is never satisfied since $x^d \neq x^u$. Therefore there exist no fixed points other than those found above.

□

A.4. Proof of Lemma 2.1

Consider first the Jacobian matrix $\mathcal{J}^{\mathfrak{C}}$ of system (2.2.15) computed at fixed point $\mathfrak{C}$

$$\mathcal{J}^{\mathfrak{C}} = \begin{bmatrix} \left.\frac{\partial \varphi}{\partial \varphi_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial \varphi}{\partial r_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial \varphi}{\partial e_{-1}}\right|_{\mathfrak{C}} \\ \left.\frac{\partial r}{\partial \varphi_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial r}{\partial r_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial r}{\partial e_{-1}}\right|_{\mathfrak{C}} \\ \left.\frac{\partial e}{\partial \varphi_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial e}{\partial r_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial e}{\partial e_{-1}}\right|_{\mathfrak{C}} \end{bmatrix} = \begin{bmatrix} \frac{\overline{x}+gx_{-1}}{\overline{x}(1+g)} & 0 & 0 \\ \frac{\overline{x}(x-x_{-1})-g(\overline{x}-x)}{\overline{x}^2(1-\overline{x})} & -\frac{g}{1+g} & \frac{\overline{x}}{1-\overline{x}} \\ 0 & -\frac{g}{1+g}\cdot\frac{1-\overline{x}}{\overline{x}} & 1 \end{bmatrix}. \tag{A.4.1}$$

The matrix is stochastic since it depends on two consecutive realisations x_{-1} and x of the stochastic trader's portfolio position. Moreover, it displays the following structure:

$$\mathcal{J}^{\mathfrak{C}} = \begin{bmatrix} \clubsuit & \mathbf{0} \\ \heartsuit & \spadesuit \end{bmatrix}, \tag{A.4.2}$$

where

$$\clubsuit = \left[\left.\frac{\partial \varphi}{\partial \varphi_{-1}}\right|_{\mathfrak{C}}\right], \qquad \mathbf{0} = \begin{bmatrix} 0 & 0 \end{bmatrix}, \qquad \heartsuit = \begin{bmatrix} \left.\frac{\partial r}{\partial \varphi_{-1}}\right|_{\mathfrak{C}} \\ 0 \end{bmatrix}, \qquad \spadesuit = \begin{bmatrix} \left.\frac{\partial r}{\partial r_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial r}{\partial e_{-1}}\right|_{\mathfrak{C}} \\ \left.\frac{\partial e}{\partial r_{-1}}\right|_{\mathfrak{C}} & \left.\frac{\partial e}{\partial e_{-1}}\right|_{\mathfrak{C}} \end{bmatrix} \tag{A.4.3}$$

Null block $\mathbf{0}$ prevents the $\heartsuit$ block from having any long-lasting effect on fixed point stability. Under Assumption 2.1, the rate of growth of paid dividends is strictly positive and the deterministic eigenvalues associated with the $\spadesuit$ block are always non-negative and less than unity:

$$\lambda_1^{\spadesuit} = \frac{1}{1+g} < 1 \quad \text{if and only if} \quad g > 0, \tag{A.4.4}$$

$$\lambda_2^{\spadesuit} = 0 < 1. \tag{A.4.5}$$

Since stochastic block $\clubsuit$ of Jacobian matrix (A.4.1) is 1×1 its eigenvalue equals the $\clubsuit$ element itself. The latter equals the ratio of the growth rate of wealth of the two groups, after a given realisation (x_{-1}, x). Local asymptotic stability follows by applying the same argument used in Theorem 4.3 Bottazzi and

Dindo (2014) that exploits Oseledets' multiplicative ergodic theorem and the local Hartman–Grobman theorem (Coayla-Teran and Ruffino, 2004, cf. Theorem 2.1 and 3.2, respectively). In particular, a sufficient condition for local stability can be given in terms of the geometric expected value of ♣ according to the invariant distribution $\tilde{\pi}$ of the Markov process $\{x_{-1},\ x\}$, as derived in Corollary 2.1. From the definition of $\rho^i\big|_{j'}$ the following equality holds:

$$\mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}\left[\clubsuit\right] = \mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}\left[\left.\frac{\partial\varphi}{\partial\varphi_{-1}}\right|_{\mathfrak{C}}\right] = \frac{\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}}{\rho^{\mathfrak{C}}\big|_{\mathfrak{C}}}. \tag{A.4.6}$$

Fixed point $\mathfrak{C}$ is asymptotically stable when $\mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}\left[\clubsuit\right] < 1$ and unstable when $\mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}\left[\clubsuit\right] > 1$. The latter inequalities can be therefore characterised in terms of the relative value of $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}$ and $\rho^{\mathfrak{C}}\big|_{\mathfrak{C}}$.

An analogous argument is applicable to the Jacobian matrix $\mathcal{J}^{\mathfrak{S}}$ of system (2.2.15) computed at fixed point $\mathfrak{S}$. The latter appears structurally similar to $\mathcal{J}^{\mathfrak{C}}$. We don't report its complete derivation for the sake of brevity.

□

A.5. Proof of Proposition 2.4

Following Lemma 2.1, a sufficient condition for local stability can be given in terms of the geometric expected value of ♣ in Jacobian $\mathcal{J}^{\mathfrak{C}}$ according to the invariant distribution $\tilde{\pi}$ of the Markov process $\{x_{-1},\ x\}$, as derived in Corollary 2.1:

$$\mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}\left[\clubsuit\right] = \frac{\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}}{\rho^{\mathfrak{C}}\big|_{\mathfrak{C}}} = \frac{\left(\overline{x}+gx^{u}\right)^{\frac{\pi^{u}}{\pi^{u}+\pi^{d}}}\left(\overline{x}+gx^{d}\right)^{\frac{\pi^{d}}{\pi^{u}+\pi^{d}}}}{\overline{x}(1+g)}. \tag{A.5.1}$$

Inequality (2.4.3) guarantees that the $\mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}\left[\clubsuit\right]$ lies inside the unit circle. Instability follows along the same lines.

□

A.6. Proof of Proposition 2.5

Due to the strict concavity of the logarithm function, for the growth rate $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}$ it holds

$$\log\left[\left(1+g\frac{x^u}{\overline{x}}\right)^{\frac{\pi^u}{\pi^u+\pi^d}}\left(1+g\frac{x^d}{\overline{x}}\right)^{\frac{\pi^d}{\pi^u+\pi^d}}\right] < \log\left[1+\frac{g}{\overline{x}}\left(\frac{\pi^u}{\pi^u+\pi^d}\cdot x^u+\frac{\pi^d}{\pi^u+\pi^d}\cdot x^d\right)\right]$$

$$= \log\left[1+g\frac{\mathbb{E}[x_t]}{\overline{x}}\right]. \tag{A.6.1}$$

It follows that if $\overline{x}\geq\mathbb{E}[x_t]$ then $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}} < \rho^{\mathfrak{C}}\big|_{\mathfrak{C}}$.

□

A.7. Proof of Proposition 2.6

$\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}(\cdot)$ is a continuous and strictly decreasing function of $\overline{x}$ on the open interval $(0,1)$ as it holds

$$\frac{\partial}{\partial\overline{x}}\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}(\cdot) = -g\,\frac{\sum_{i\in\{u,d\},j\in\{u,d\},i\neq j}\pi^i x^i(\overline{x}+gx^i)^{\frac{\pi^i}{\pi^i+\pi^j}}(\overline{x}+gx^j)^{\frac{\pi^i+2\pi^j}{\pi^i+\pi^j}}}{\overline{x}^2(\pi^u+\pi^d)(\overline{x}+gx^u)(\overline{x}+gx^d)} < 0. \tag{A.7.1}$$

Moreover, its limiting behaviour at the extrema of the support is characterised by

$$\lim_{\overline{x}\to 0^+}\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}(\cdot) = +\infty, \tag{A.7.2}$$

$$\lim_{\overline{x}\to 1^-}\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}(\cdot) = \mathbb{E}^{\tilde{\pi}}_{\mathcal{G}}[1+gx_{-1}] < 1+g. \tag{A.7.3}$$

Applying the intermediate value theorem yields the desired result.

$\boxed{x^d < x'}$

If $\overline{x} = x^d$, then condition (2.4.3) is violated since

$$\left[\frac{x^d+gx^u}{x^d+gx^d}\right]^{\frac{\pi^u}{\pi^u+\pi^d}} > 1. \tag{A.7.4}$$

Since by Proposition 2.6 $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}(\cdot)$ is monotone decreasing in $\overline{x}$, it must be that $x' > x^d$.

$\boxed{x' < x^u}$

If $\overline{x} = x^u$, then LHS of eq. (2.4.3) reads

$$\left[\frac{x^u + g x^d}{x^u + g x^u}\right]^{\frac{\pi^d}{\pi^u + \pi^d}} < 1. \tag{A.7.5}$$

Since by Proposition 2.6 $\rho^{\mathcal{S}}\big|_{\mathcal{C}}(\cdot)$ is monotone decreasing in $\overline{x}$, by continuity $\exists\, \varepsilon > 0$ such that $\rho^{\mathcal{S}}\big|_{\mathcal{C}}(x^u - \varepsilon) < 1 + g$. Therefore, it must be that $x' < x^u$.

Moreover, due to Proposition 2.5 it holds that $x' < \mathbb{E}[x_t]$.

□

A.8. Proof of Proposition 2.7

Following Lemma 2.1, a sufficient condition for local stability can be given in terms of the geometric expected value of ♣ in Jacobian $\mathcal{J}^{\mathcal{S}}$ according to the invariant distribution $\tilde{\pi}$ of the Markov process $\{x_{-1},\ x\}$, as derived in Corollary 2.1:

$$\begin{aligned}\mathbb{E}_{\mathcal{G}}^{\tilde{\pi}}[\clubsuit] = \frac{\rho^{\mathcal{C}}\big|_{\mathcal{S}}}{\rho^{\mathcal{S}}\big|_{\mathcal{S}}} = &\left[1 + \frac{g\overline{x}}{x^u}\right]^{\frac{\pi^u(1-\pi^d)}{\pi^u+\pi^d}} \cdot \left[1 + g\frac{\overline{x}}{x^u} + \frac{(x^u - \overline{x})(x^u - x^d)}{x^u(1-x^u)}\right]^{\frac{\pi^u\pi^d}{\pi^u+\pi^d}} \\ &\cdot \left[1 + g\frac{\overline{x}}{x^d} + \frac{(\overline{x} - x^d)(x^u - x^d)}{x^d(1-x^d)}\right]^{\frac{\pi^u\pi^d}{\pi^u+\pi^d}} \cdot \left[1 + \frac{g\overline{x}}{x^d}\right]^{\frac{\pi^d(1-\pi^u)}{\pi^u+\pi^d}} \cdot (1+g)^{-1}.\end{aligned} \tag{A.8.1}$$

Inequality (2.4.3) guarantees that the $\mathbb{E}_{\mathcal{G}}^{\tilde{\pi}}[\clubsuit]$ lies inside the unit circle. Instability follows along the same lines.

□

A.9. Proof of Proposition 2.8

Start by decomposing eq. (2.4.6) in its multiplicative components. Let us define $\rho_{ij}^{\mathcal{C}}\big|_{\mathcal{S}}$ as the value of $\rho^{\mathcal{C}}\big|_{\mathcal{S}}$ when $x_{-1} = x^i$ and $x = x^j$, with $x^i, x^j \in \{x^u, x^d\}$. Note that we can safely get rid of the probability exponent since it plays no role in the

current analysis. It is straightforward to check that

$$\rho^{\mathfrak{C}}_{uu}\big|_{\mathfrak{S}} \lesseqgtr 1+g \quad \Longleftrightarrow \quad \overline{x} \lesseqgtr x^u, \tag{A.9.1}$$

$$\rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}} = 1+g \quad \Longleftrightarrow \quad \overline{x} = x^u \quad \vee \quad g = \tilde{g} := \frac{x^u - x^d}{1-x^u}, \tag{A.9.2}$$

$$\rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}} \lesseqgtr 1+g \quad \Longleftrightarrow \quad \left(\overline{x} \lesseqgtr x^u \quad \wedge \quad g \gtreqless \tilde{g}\right) \quad \vee \quad \left(\overline{x} \gtreqless x^u \quad \wedge \quad g \lesseqgtr \tilde{g}\right), \tag{A.9.3}$$

$$\rho^{\mathfrak{C}}_{du}\big|_{\mathfrak{S}} \lesseqgtr 1+g \quad \Longleftrightarrow \quad \overline{x} \lesseqgtr x^d, \tag{A.9.4}$$

$$\rho^{\mathfrak{C}}_{dd}\big|_{\mathfrak{S}} \lesseqgtr 1+g \quad \Longleftrightarrow \quad \overline{x} \lesseqgtr x^d, \tag{A.9.5}$$

Condition (2.4.7) and follows from the joint satisfaction of eqs. (A.9.1) and (A.9.3) to (A.9.5) with strict inequality signs '$<$'.

Moreover, it is possible to show that $\rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}} \cdot \rho^{\mathfrak{C}}_{du}\big|_{\mathfrak{S}} > 1+g$, $\forall g > 0$ that is, independently of the threshold $\tilde{g}$ (note that $\rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}}$ and $\rho^{\mathfrak{C}}_{du}\big|_{\mathfrak{S}}$ are powered by the same exponent in eq. (2.4.6), as per the invariant distribution $\tilde{\boldsymbol{\pi}}$ in eq. (2.3.14)). As a consequence, the sufficient condition (2.4.8) for the joint satisfaction of eqs. (A.9.1) and (A.9.3) to (A.9.5) with strict inequality signs '$>$' only requires $\overline{x} \geq x^u$.

□

A.10. Hint of Conjecture 2.1

$\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(\cdot)$ is a continuous and differentiable function of $\overline{x}$ over the open interval $(0,1)$. At the extrema of the support it holds

$$\lim_{\overline{x}\to 0^+} \rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(\cdot) \;=\; 1 \;<\; 1+g, \tag{A.10.1}$$

$$\lim_{\overline{x}\to 1^-} \rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(\cdot) \;=\; \mathbb{E}^{\tilde{\boldsymbol{\pi}}}_{\mathcal{G}}\left[1+\frac{g}{x_{-1}}\right] \;>\; 1+g. \tag{A.10.2}$$

Moreover, it is possible to numerically show that $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(\cdot)$ exhibit no inflection points over $\overline{x} \in (0,1)$. We look for the roots of $\frac{\partial^2}{\partial \overline{x}^2}\, \rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(\cdot) = 0$ using the Anderson-Björck method $\forall x^u, x^d, \pi^u, \pi^d \in \{0.01, 0.02, \ldots, 0.99\}$ such that $x^u > x^d$, $\forall g \in 10^k, k \in \{-4,-3,\ldots,+1\}$. We then check that none of them belong to the interval $(0,1)$, finding no exceptions. Applying the intermediate value theorem yields the desired result.

A.11. Proof of Proposition 2.9

Start by decomposing eq. (2.4.6) in its multiplicative components. Let us define $\rho^{\mathfrak{C}}_{ij}\big|_{\mathfrak{S}}$ as the value of $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}$ when $x_{-1} = x^i$ and $x = x^j$, with $x^i, x^j \in \{x^u, x^d\}$. It is straightforward to check that if $\overline{x} = x^d$ then $\rho^{\mathfrak{C}}_{du}\big|_{\mathfrak{S}}(x^d) = \rho^{\mathfrak{C}}_{dd}\big|_{\mathfrak{S}}(x^d) = 1 + g$ and $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}$ reduces to

$$\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(x^d) = \rho^{\mathfrak{C}}_{uu}\big|_{\mathfrak{S}}(x^d)^{\frac{\pi^u(1-\pi^d)}{\pi^u+\pi^d}} \cdot \rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}}(x^d)^{\frac{\pi^u\pi^d}{\pi^u+\pi^d}} \cdot (1+g)^{\frac{\pi^u\pi^d}{\pi^u+\pi^d}+\frac{\pi^d(1-\pi^u)}{\pi^u+\pi^d}}. \tag{A.11.1}$$

From Proposition 2.8 it follows that

$$\rho^{\mathfrak{C}}_{uu}\big|_{\mathfrak{S}}(x^d) < 1 + g, \tag{A.11.2}$$

$$\rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}}(x^d) \gtreqless 1 + g \quad \Longleftrightarrow \quad g \lesseqgtr \frac{x^u - x^d}{1 - x^u}. \tag{A.11.3}$$

Moreover,

$$\lim_{g\to 0^+} \rho^{\mathfrak{C}}_{uu}\big|_{\mathfrak{S}}(x^d) = 1, \tag{A.11.4}$$

$$\lim_{g\to 0^+} \rho^{\mathfrak{C}}_{ud}\big|_{\mathfrak{S}}(x^d) = \frac{x^u + x^d(2x^u - x^d)}{x^u(1 - x^u)} > 1. \tag{A.11.5}$$

Therefore, there exists a $\hat{g} > 0$ such that $\forall\, g < \hat{g}$ condition (2.4.6) is strictly violated, i.e. $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(x^d) > 1 + g$ and it must be that $x'' < x^d$. Finally, from Proposition 2.6 it holds that $x^d < x'$.

□

A.12. Hint of Conjecture 2.2

The unique solution x' of $\rho^{\mathfrak{S}}\big|_{\mathfrak{C}}(\overline{x}) = 1 + g$ is computed using the Anderson-Björck method and then the condition $\rho^{\mathfrak{C}}\big|_{\mathfrak{S}}(x') > 1 + g$ is tested $\forall x^u, x^d, \pi^u, \pi^d \in \{0.01, 0.02, \ldots, 0.99\}$ such that $x^u > x^d$, and $\forall g = \hat{g} \cdot 10^k$, $k \in \mathbb{N}_+$ under the restriction $g \leq 10$. No exception is found.

This page intentionally left blank.

References

Aït-Sahalia, Yacine, Jianqing Fan and Yingying Li (2013). 'The leverage effect puzzle: Disentangling sources of bias at high frequency'. *Journal of Financial Economics* 109(1), pp. 224–249. DOI: 10.1016/j.jfineco.2013.02.018.

Aït-Sahalia, Yacine, Per A. Mykland and Zhang Lan (2005). 'How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise'. *Review of Financial Studies* 18(2), pp. 351–416. DOI: 10.1007/0-387-28359-5_1.

Aït-Sahalia, Yacine and Dacheng Xiu (2018). 'Principal Component Analysis of High-Frequency Data'. *Journal of the American Statistical Association*. DOI: 10.1080/01621459.2017.1401542.

Aldridge, Irene (2013). *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. Wiley. DOI: 10.1002/9781119203803.

Allen, Franklin, Stephen Morris and Hyun Song Shin (2006). 'Beauty Contests and Iterated Expectations in Asset Markets'. *Review of Financial Studies* 19(3), pp. 719–752. DOI: 10.1093/rfs/hhj036.

Allen, Helen and Mark P. Taylor (1990). 'Charts, Noise and Fundamentals in the London Foreign Exchange Market'. *Economic Journal* 100(400), pp. 49–59. DOI: 10.2307/2234183.

Almazan, Andres, Keith C. Brown, Murray Carlson and David A. Chapman (2004). 'Why constrain your mutual fund manager?' *Journal of Financial Economics* 73(2), pp. 289–321. DOI: 10.1016/j.jfineco.2003.05.007.

Aloud, M., E. Tsang and R. Olsen (2013). 'Modeling the FX market traders' behavior: An agent-based approach'. In: *Simulation in Computational Finance and*

Economics: Tools and Emerging Applications. Ed. by B. Alexandrova-Kabadjova, S. Martinez-Jaramillo, A. Garcia-Almanza and E. Tsang. Hershey, PA: Business Science Reference. DOI: 10.4018/978-1-4666-6268-1.ch018.

Andersen, Torben G. and Tim Bollerslev (1997). 'Intraday periodicity and volatility persistence in financial markets'. *Journal of Empirical Finance* 4(2-3), pp. 115–158. DOI: 10.1016/s0927-5398(97)00004-2.

Anufriev, Mikhail and Giulio Bottazzi (2010). 'Market Equilibria under Procedural Rationality'. *Journal of Mathematical Economics* 46(6), pp. 1140–1172. DOI: 10.1016/j.jmateco.2010.09.005.

Anufriev, Mikhail, Giulio Bottazzi, Matteo Marsili and Paolo Pin (2012). 'Excess covariance and dynamic instability in a multi-asset model'. *Journal of Economic Dynamics and Control* 36(8), pp. 1142–1161. DOI: 10.1016/j.jedc.2012.03.015.

Anufriev, Mikhail, Giulio Bottazzi and Francesca Pancotto (2006). 'Equilibria, stability and asymptotic dominance in a speculative market with heterogeneous traders'. *Journal of Economic Dynamics and Control* 30, pp. 1787–1835. DOI: 10.1016/j.jedc.2005.10.012.

Anufriev, Mikhail and Pietro Dindo (2010). 'Wealth-driven selection in a financial market with heterogeneous agents'. *Journal of Economic Behavior & Organization* 73, pp. 327–358. DOI: 10.1016/j.jebo.2009.11.006.

Arthur, W. Brian (1999). 'Complexity and the Economy'. *Science* 284(5411), pp. 107–109. DOI: 10.1126/science.284.5411.107.

Arthur, W. Brian (2014). *Complexity and the Economy*. Oxford University Press.

Arthur, W. Brian, John H. Hilland, Blake LeBaron, Richard Palmer and Paul Tayler (1997). 'Asset Pricing Under Endogenous Expectations in an Artificial Stock Market'. In: *The Economy as an Evolving Complex System II*. Ed. by W. Brian Arthur, Steven N. Durlauf and David A. Lane. Addison-Wesley.

Bachelier, Louis (1900). 'Théorie de la spéculation'. *Annales Scientifiques de l'École Normale Supérieure* 3(17), pp. 21–86. DOI: 10.24033/asens.476.

Bandi, Federico M. and Peter C. B. Phillips (2003). 'Fully Nonparametric Estimation of Scalar Diffusion Models'. *Econometrica* 71(1), pp. 241–283. DOI: 10.1111/1468-0262.00395.

Barberis, Nicholas C. (2013). 'Thirty Years of Prospect Theory in Economics: A Review and Assessment'. *Journal of Economic Perspectives* 27(1), pp. 173–196. DOI: 10.3386/w18621.

Barberis, Nicholas C. and Richard H. Thaler (2003). 'A survey of behavioral finance'. In: *Handbook of the Economics of Finance*. Vol. 1B. Elsevier. Chap. 18, pp. 1053–1128. DOI: 10.1016/S1574-0102(03)01027-6.

Barigozzi, Matteo and Marc Hallin (2016). 'Generalized dynamic factor models and volatilities: recovering the market volatility shocks'. *Econometrics Journal* 19(1), pp. C33–C60. DOI: 10.1111/ectj.12047.

Basu, S. (1977). 'Investment performance of common stocks in relation to their price-earnings ratios: a test of the efficient market hypothesis'. *Journal of Finance* 32(3), pp. 663–682. DOI: 10.1111/j.1540-6261.1977.tb01979.x.

Benartzi, Shlomo and Richard H. Thaler (1995). 'Myopic Loss Aversion and the Equity Premium Puzzle'. *Quarterly Journal of Economics* 110(1), pp. 73–92. DOI: 10.3386/w4369.

Bhargava, Alok (2014). 'Firms' fundamentals, macroeconomic variables and quarterly stock prices in the US'. *Journal of Econometrics* 183(2), pp. 241–250. DOI: 10.1016/j.jeconom.2014.05.014.

Biais, Bruno, Pierre Hillion and Chester Spatt (1995). 'An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse'. *Journal of Finance* 50(5), pp. 1655–1689. DOI: 10.1111/j.1540-6261.1995.tb05192.x.

Black, Fischer (1976). 'Studies of Stock Price Volatility Changes'. *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section*, pp. 177–181.

Black, Fischer (1986). 'Noise'. *Journal of Finance* 41(3), pp. 529–543. DOI: 10.1111/j.1540-6261.1986.tb04513.x.

Bloomfield, Robert and Alyssa Anderson (2010). 'Experimental Finance'. In: *Behavioral Finance: Investors, Corporations, and Markets*. Ed. by H. Kent Baker and John R. Nofsinger. John Wiley & Sons. DOI: 10.1002/9781118258415.ch7.

Blume, Lawrence and David Easley (1992). 'Evolution and Market Behavior'. *Journal of Economic Theory* 58, pp. 9–40. DOI: 10.1016/0022-0531(92)90099-4.

Blume, Lawrence and David Easley (2006). 'If You're so Smart, why Aren't You Rich? Belief Selection in Complete and Incomplete Markets'. *Econometrica* 74(4), pp. 929–966. DOI: 10.1111/j.1468-0262.2006.00691.x.

Bollerslev, Tim, Julia Litvinova and George Tauchen (2006). 'Leverage and Volatility Feedback Effects in High-Frequency Data'. *Journal of Financial Econometrics* 4(3), pp. 353–384. DOI: 10.1093/jjfinec/nbj014.

Bonanno, Giovanni, Fabrizio Lillo and Rosario Mantegna (2000). 'Dynamics of the number of trades of financial securities'. *Physica A: Statistical Mechanics and its Applications* 280(1-2), pp. 136–141. DOI: 10.1016/S0378-4371(99)00629-9.

Bottazzi, Giulio and Pietro Dindo (2014). 'Evolution and market behavior with endogenous investment rules'. *Journal of Economic Dynamics and Control* 48, pp. 121–146. DOI: 10.1016/j.jedc.2014.08.012.

Bouchaud, Jean-Philippe (2008). 'Economics needs a scientific revolution'. *Nature* 455, p. 1181. DOI: 10.1038/4551181a.

Bouchaud, Jean-Philippe, Andrew Matacz and Mark Potters (2001). 'Leverage Effect in Financial Markets: The Retarded Volatility Model'. *Physical Review Letters* 87(22), p. 228701. DOI: 10.1103/PhysRevLett.87.228701.

Breen, Willam and James Savage (1968). 'Portfolio distributions and tests of security selection models'. *Journal of Finance* 23(5), pp. 805–819. DOI: 10.1111/j.1540-6261.1968.tb00318.x.

Breen, William (1968). 'Low Price-Earnings Ratios and Industry Relatives'. *Financial Analysts Journal* 24(4), pp. 125–127. DOI: 10.2469/faj.v24.n4.125.

Brock, William A. and Cars H. Hommes (1998). 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model'. *Journal of Economic Dynamics and Control* 22(8-9), pp. 1235–1274. DOI: 10.1016/S0165-1889(98)00011-6.

Buffett, Warren E. (1984). 'The superinvestors of Graham-and-Doddsville'. *Hermes* Fall(4–15).

Campbell, John Y. (2000). 'Asset Pricing at the Millennium'. *Journal of Finance* 55(4), pp. 1515–1567. DOI: 10.3386/w7589.

Campbell, John Y., Sanford J. Grossman and Wang Jiang (1993). 'Trading Volume and Serial Correlation in Stock Returns'. *Quarterly Journal of Economics* 108(4), pp. 905–939. DOI: 10.2307/2118454.

Chiarella, Carl (1992). 'The dynamics of speculative behaviour'. *Annals of Operations Research* 37(1), pp. 101–123. DOI: 10.1007/BF02071051.

Chiarella, Carl, Roberto Dieci and Laura Gardini (2006). 'Asset price and wealth dynamics in a financial market with heterogeneous agents'. *Journal of Economic Dynamics and Control* 30(9-10), pp. 1755–1786. DOI: 10.1016/j.jedc.2005.10.011.

Chiarella, Carl and Xue-Zhong He (2001). 'Asset Price and Wealth Dynamics Under Heterogeneous Expectations'. *Quantitative Finance* 1(5), pp. 509–526. DOI: 10.1088/1469-7688/1/5/303.

Chiarella, Carl, Xue-Zhong He and Min Zheng (2011). 'An analysis of the effect of noise in a heterogeneous agent financial market model'. *Journal of Economic Dynamics and Control* 35(1), pp. 148–162. DOI: 10.1016/j.jedc.2010.09.006.

Chiarella, Carl and Giulia Iori (2002). 'A simulation analysis of the microstructure of double auction markets'. *Quantitative Finance* 2, pp. 346–353. DOI: 10.1088/1469-7688/2/5/303.

Cioffi-Revilla, Claudio (2002). 'Invariance and Universality in Social Agent-Based Simulations'. *Proceedings of the National Academy of Sciences of the United States of America* 99(10), pp. 7314–7316. DOI: 10.1073/pnas.082081499.

Coates, J. M. and J. Herbert (2008). 'Endogenous steroids and financial risk taking on a London trading floor'. *Proceedings of the National Academy of Sciences* 105(16), pp. 6167–6172. DOI: 10.1073/pnas.0704025105.

Coayla-Teran, Edson A. and Paulo R. Ruffino (2004). 'Stochastic versions of Hartman-Grobman theorems'. *Stochastics and Dynamics* 4(4), pp. 571–593. DOI: 10.1142/S0219493704001206.

Conlisk, John (1996). 'Why Bounded Rationality?' *Journal of Economic Literature* 34(2), pp. 669–700.

Cont, Rama (2001). 'Empirical properties of asset returns: stylized facts and statistical issues'. *Quantitative Finance* 1, pp. 223–236. DOI: 10.1080/713665670.

Cont, Rama (2007). 'Volatility Clustering in Financial Markets: Empirical Facts and Agent–Based Models'. In: *Long Memory in Economics*. Ed. by Gilles Teyssière and Alan P. Kirman. Springer, pp. 289–309. DOI: 10.1007/978-3-540-34625-8_10.

Cont, Rama (2011). 'Statistical Modeling of High-Frequency Financial Data'. *IEEE Signal Processing Magazine* 28(5), pp. 16–25. DOI: 10.1109/MSP.2011.941548.

Cutler, David M., James M. Poterba and Lawrence H. Summers (1989). 'What moves stock prices?' *Journal of Portfolio Management* 15(3), pp. 4–12. DOI: 10.3905/jpm.1989.409212.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldmann (1990). 'Noise Trader Risk in Financial Markets'. *Journal of Political Economy* 98(4), pp. 703–738. DOI: 10.1086/261703.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldmann (1991). 'The Survival of Noise Traders in Financial Markets'. *Journal of Business* 64(1), pp. 1–19. DOI: 10.3386/w2715.

DeMiguel, Victor, Lorenzo Garlappi and Raman Uppal (2009). 'Optimal Versus Naive Diversification: How Inefficient is the $1/N$ Portfolio Strategy?' *Review of Financial Studies* 22(5), pp. 1915–1953. DOI: 10.1093/rfs/hhm075.

Dreman, David N. and Michael A. Berry (1995). 'Overreaction, Underreaction, and the Low-P/E Effect'. *Financial Analysts Journal* 51(4), pp. 21–30. DOI: 10.2469/faj.v51.n4.1917.

Easley, David, Marcos M. López DE Prado and Maureen O'Hara (2012). 'The Volume Clock: Insights into the High-Frequency Paradigm'. *Journal of Portfolio Management* 39(1), pp. 19–29. DOI: 10.3905/jpm.2012.39.1.019.

Engle, Robert F. (2000). 'The Econometrics of Ultra-High-Frequency Data'. *Econometrica* 68(1), pp. 1–22. DOI: 10.1111/1468-0262.00091.

Euronext (2017). *Trading manual for the Universal Trading Platform*. URL: https://www.euronext.com/sites/www.euronext.com/files/trading_manual_en_2018_01_03.pdf.

Evstigneev, Igor, Thorsten Hens and Klaus Reiner Schenk-Hoppé (2011). 'Local stability analysis of a stochastic evolutionary financial market model with a risk-free asset'. *Mathematics and Financial Economics* 5(3), pp. 185–202. DOI: 10.1007/s11579-011-0056-z.

Fama, Eugene F. (1965). 'The Behavior of Stock-Market Prices'. *Journal of Business* 38(1), pp. 34–105. DOI: 10.1086/294743.

Fama, Eugene F. (1970). 'Efficient Capital Markets: A Review of Theory and Empirical Work'. *Journal of Finance* 25(2), pp. 383–417. DOI: 10.2307/2325486.

Farmer, J. Doyne (2002). 'Market force, ecology and evolution'. *Industrial and Corporate Change* 11(5), pp. 895–953. DOI: 10.1093/icc/11.5.895.

Farmer, J. Doyne and Andrew W. Lo (1999). 'Frontiers of finance: Evolution and efficient markets'. *Proceedings of the National Academy of Sciences* 96(18), pp. 9991–9992. DOI: 10.1073/pnas.96.18.9991.

Florens-Zmirou, Danielle (1993). 'On estimating the diffusion coefficient from discrete observations'. *Journal of Applied Probability* 30(4), pp. 790–804. DOI: 10.1017/s0021900200044570.

Foster, Andrew J. (1995). 'Volume-volatility relationships for crude oil futures markets'. *Journal of Futures Markets* 15(8), pp. 929–951. DOI: 10.1002/fut.3990150805.

Frankel, Jeffrey A. and Kenneth A. Froot (1990). 'Chartists, Fundamentalists, and Trading in the Foreign Exchange Market'. *American Economic Review* 80(2), pp. 181–185.

Gallant, A. Ronald, Peter E. Rossi and George Tauchen (1992). 'Stock Prices and Volume'. *Review of Financial Studies* 5(2), pp. 199–242. DOI: 10.1093/rfs/5.2.199.

Gallo, Giampiero M. and Christian T. Brownlees (2006). 'Financial econometric analysis at ultra-high frequency: Data handling concerns'. *Computational Statistics & Data Analysis* 51(4), pp. 2232–2245. DOI: 10.1016/j.csda.2006.09.030.

Ghoulmie, François, Rama Cont and Jean-Pierre Nadal (2005). 'Heterogeneity and feedback in an agent-based market model'. *Journal of Physics: Condensed Matter* 17(14), S1259–S1268. DOI: 10.1088/0953-8984/17/14/015.

Gigerenzer, Gerd (2008). 'Why Heuristics Work'. *Perspectives on Psychological Science* 3(1), pp. 20–29. DOI: 10.1111/j.1745-6916.2008.00058.x.

Gigerenzer, Gerd and Wolfgang Gaissmaier (2011). 'Heuristic Decision Making'. *Annual Review of Psychology* 62(1), pp. 451–482. DOI: 10.1146/annurev-psych-120709-145346.

Goldstein, Daniel G. and Gerd Gigerenzer (2002). 'Models of Ecological Rationality: The Recognition Heuristic'. *Psychological Review* 109(1), pp. 75–90. DOI: 10.1037//0033-295X.109.1.75.

Hays, Spencer, Haipeng Shen and Jianhua Z. Huang (2012). 'Functional dynamic factor models qith application to yield curve forecasting'. *Annals of Applied Statistics* 6(3), pp. 870–894. DOI: 10.1214/12-aoas551.

Hens, Thorsten and Klaus Reiner Schenk-Hoppé (2005). 'Evolutionary finance: introduction to the special issue'. *Journal of Mathematical Economics* 41(1–2), pp. 1–5. DOI: 10.1016/j.jmateco.2004.09.001.

Herskovic, Bernard, Bryan Kelly, Hanno Lustig and Stijn Van Nieuwerburgh (2016). 'The common factor in idiosyncratic volatility: Quantitative asset pricing implications'. *Journal of Financial Economics* 119(2), pp. 249–283. DOI: 10.1016/j.jfineco.2015.09.010.

Hirshleifer, David (2001). 'Investor Psychology and Asset Pricing'. *Journal of Finance* 56(4), pp. 1533–1597. DOI: 10.1111/0022-1082.00379.

Hirshleifer, David and Guo Ying Luo (2001). 'On the survival of overconfident traders in a competitive securities market'. *Journal of Financial Markets* 4(1), pp. 73–84. DOI: 10.1016/s1386-4181(00)00014-8.

Hirshleifer, David and Tyler Shumway (2003). 'Good Day Sunshine: Stock Returns and the Weather'. *Journal of Finance* 58(3), pp. 1009–1032. DOI: 10.1111/1540-6261.00556.

Hirshleifer, Jack (1977). 'Economics from a Biological Viewpoint'. *Journal of Law and Economics* 20(1), pp. 1–52. DOI: 10.1086/466891.

Hommes, Cars H. (2006). 'Heterogeneous Agent Models in Economics and Finance'. In: *Handbook of Computational Economics, volume 2*. Ed. by Leigh Tesfatsion and Kenneth L. Judd. North-Holland, pp. 1109–1186. DOI: 10.1016/S1574-0021(05)02023-X.

Hörmann, Siegfried, Łukasz Kidziński and Marc Hallin (2014). 'Dynamic functional principal components'. *Journal of the Royal Statistical Society, Statistical Methodology Series B* 77(2), pp. 319–348. DOI: 10.1111/rssb.12076.

Hörmann, Siegfried and Piotr Kokoszka (2010). 'Weakly dependent functional data'. *Annals of Statistics* 38(3), pp. 1845–1884. DOI: 10.1214/09-aos768.

Hull, J.H. (2017). *Options, Futures, and Other Derivatives*. 9th ed. Pearson.

Jacob Leal, Sandrine, Mauro Napoletano, Andrea Roventini and Giorgio Fagiolo (2016). 'Rock around the clock: An agent-based model of low- and high-

frequency trading'. *Journal of Evolutionary Economics* 26(1), pp. 49–76. DOI: 10.1007/s00191-015-0418-4.

Jain, Prem C. and Gun Ho Joh (1988). 'The Dependence between Hourly Prices and Trading Volume'. *Journal of Financial and Quantitative Analysis* 23(3), pp. 269–283. DOI: 10.2307/2331067.

Jensen, Michael C. (1978). 'Some anomalous evidence regarding market efficiency'. *Journal of Financial Economics* 6(2–3), pp. 95–101. DOI: 10.1016/0304-405x(78)90025-9.

Kahneman, Daniel (2003). 'Maps of bounded rationality: psychology for behavioral economics'. *American Economic Review* 93(5), pp. 1449–1475. DOI: 10.1257/000282803322655392.

Kahneman, Daniel and Amos Tversky (1979). 'Prospect Theory: An Analysis of Decision under Risk'. *Econometrica* 47(2), pp. 263–292. DOI: 10.2307/1914185.

Keynes, John M. (1936). *The General Theory of Unemployment, Interest and Money*. Harcourt, Brace and World, New York.

Kirman, Alan P. (1992). 'Whom or What Does the Representative Individual Represent?' *Journal of Economic Perspectives* 6(2), pp. 117–136. DOI: 10.1257/jep.6.2.117.

Kirman, Alan P. (2011). *Complex Economics: Individual and Collective Rationality*. Routledge.

Kirman, Alan P. and Gilles Teyssière (2002). 'Microeconomic Models for Long Memory in the Volatility of Financial Time Series'. *Studies in Nonlinear Dynamics & Econometrics* 5(4), pp. 281–302. DOI: doi.org/10.2202/1558-3708.1083.

Kluger, Brian D. and Mark E. McBride (2011). 'Intraday trading patterns in an intelligent autonomous agent-based stock market'. *Journal of Economic Behavior & Organization* 79(3), pp. 226–245. DOI: 10.1016/j.jebo.2011.01.032.

Kokoszka, Piotr, Hong Miao and Xi Zhang (2014). 'Functional Dynamic Factor Model for Intraday Price Curves'. *Journal of Financial Econometrics* 13(2), pp. 456–477. DOI: 10.1093/jjfinec/nbu004.

Kon, Stanley J. (1984). 'Models of Stock Returns—A Comparison'. *Journal of Finance* 39(1), pp. 147–165. DOI: 10.1111/j.1540-6261.1984.tb03865.x.

Lakonishok, Josef and Seymour Smidt (1988). 'Are Seasonal Anomalies Real? A Ninety-Year Perspective'. *Review of Financial Studies* 1(4), pp. 403–425. DOI: 10.1093/rfs/1.4.403.

LeBaron, Blake (2006). 'Agent-based Computational Finance'. In: *Handbook of Computational Economics, volume 2*. Ed. by Leigh Tesfatsion and Kenneth L. Judd. North-Holland. DOI: 10.1016/S1574-0021(05)02024-1.

Levy, Moshe, Haim Levy and Sorin Solomon (1994). 'A microscopic model of the stock market: Cycles, booms, and crashes'. *Economics Letters* 45(1), pp. 103–111. DOI: 10.1016/0165-1765(94)90065-5.

Lo, Andrew W. (2004). 'The Adaptive Markets Hypothesis: Market efficiency from an evolutionary perspective.' *Journal of Portfolio Management* 30(5), pp. 15–29. DOI: 10.3905/jpm.2004.442611.

Lo, Andrew W. and A. Craig MacKinley (2002). *A Non-Random Walk Down Wall Street*. Princeton University Press.

Lo, Andrew W., Dmitry V. Repin and Brett N. Steenbarger (2005). 'Fear and Greed in Financial Markets: A Clinical Study of Day-Traders'. *American Economic Review* 95(2), pp. 352–359. DOI: 10.1257/000282805774670095.

Lockwood, Larry J. and Scott C. Linn (1990). 'An Examination of Stock Market Return Volatility During Overnight and Intraday'. *Journal of Finance* 45(2), pp. 591–601. DOI: 10.1111/j.1540-6261.1990.tb03705.x.

Lux, Thomas (1995). 'Herd Behaviour, Bubbles and Crashes'. *Economic Journal* 105(431), pp. 881–896. DOI: 10.2307/2235156.

Lux, Thomas (1998). 'The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions'. *Journal of Economic Behavior & Organization* 33(2), pp. 143–165. DOI: 10.1016/S0167-2681(97)00088-7.

Lux, Thomas and Michele Marchesi (2000). 'Volatility clustering in financial markets: a microsimulation of interacting agents'. *International Journal of Theoretical and Applied Finance* 3(4), pp. 675–702. DOI: 10.1142/S0219024900000826.

Malkiel, Burton G. (2003). 'The Efficient Market Hypothesis and Its Critics'. *Journal of Economic Perspectives* 17(1), pp. 59–82. DOI: 10.1257/089533003321164958.

Malkiel, Burton G. (2016). *A Random Walk down Wall Street: The Time-tested Strategy for Successful Investing*. 11th ed. W. W. Norton & Company.

Mandelbrot, Benoit B. (1963). 'The Variation of Certain Speculative Prices'. *Journal of Business* 36(4), pp. 394–419. DOI: 10.1086/294632.

Mandelbrot, Benoit B. (1971). 'When Can Price be Arbitraged Efficiently? A Limit to the Validity of the Random Walk and Martingale Models'. *Review of Economics and Statistics* 53(3), pp. 225–236. DOI: 10.2307/1937966.

Markowitz, Harry (1952). 'Portfolio Selection'. *Journal of Finance* 7(1), pp. 77–91. DOI: 10.1111/j.1540-6261.1952.tb01525.x.

McInish, T. H. and R. A. Wood (1991). 'Hourly returns, volume, trade size, and number of trades'. *Journal of Financial Research* 14(4), pp. 303–315. DOI: 10.1111/j.1475-6803.1991.tb00668.x.

McWilliams, James D. (1966). 'Prices, Earnings and P-E Ratios'. *Financial Analysts Journal* 22(3), pp. 137–142. DOI: 10.2469/faj.v22.n3.137.

Menkhoff, Lukas (2010). 'The use of technical analysis by fund managers: International evidence'. *Journal of Banking & Finance* 34(11), pp. 2573–2586. DOI: 10.1016/j.jbankfin.2010.04.014.

Miller, John H. and Scott E. Page (2007). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press.

Müller, Hans-Georg, Rituparna Sen and Ulrich Stadtmüller (2011). 'Functional data analysis for volatility'. *Journal of Econometrics* 165(2), pp. 233–245. DOI: 10.1016/j.jeconom.2011.08.002.

Mulvey, John M. and Woo Chang Kim (2010). 'Fixed Mix Strategy'. In: *Encyclopedia of Quantitative Finance*. Ed. by Rama Cont. Wiley. DOI: 10.1002/9780470061602.eqf14010.

Nelson, Richard R. and Sydney G. Winter (1982). *An Evolutionary Theory of Economic Change*. Belknap Press.

Nicholson, S. Francis (1968). 'Price Ratios in Relation to Investment Results'. *Financial Analysts Journal* 24(1), pp. 105–109. DOI: 10.2469/faj.v24.n1.105.

Palczewski, Jan, Klaus Reiner Schenk-Hoppé and Tongya Wang (2016). 'Itchy feet vs cool heads: Flow of funds in an agent-based financial market'. *Journal*

of Economic Dynamics and Control 63, pp. 53–68. DOI: 10.1016/j.jedc.2015.12.002.

Pellizzari, Paolo and Frank Westerhoff (2009). 'Some effects of transaction taxes under different microstructures'. *Journal of Economic Behavior & Organization* 72(3), pp. 850–863. DOI: 10.1016/j.jebo.2009.08.010.

Raberto, Marco, Enrico Scalas and Francesco Mainardi (2002). 'Waiting-times and returns in high-frequency financial data: an empirical study'. *Physica A: Statistical Mechanics and its Applications* 314(1-4), pp. 749–755. DOI: 10.1016/S0378-4371(02)01048-8.

Ramsay, James, Giles Hooker and Spencer Graves (2009). *Functional Data Analysis with R and MATLAB*. Springer. DOI: 10.1007/978-0-387-98185-7.

Ramsay, James and B. W. Silverman (2005). *Functional Data Analysis*. 2nd ed. Springer. DOI: 10.1016/b978-0-08-097086-8.42046-5.

Renò, Roberto (2008). 'Nonparametric estimation of the diffusion coefficient of stochastic volatility models'. *Econometric Theory* 24(5), pp. 1174–1206. DOI: 10.1017/s026646660808047x.

Rosser, J. Barkley (1999). 'On the Complexities of Complex Economic Dynamics'. *Journal of Economic Perspectives* 13(4), pp. 169–192. DOI: 10.1257/jep.13.4.169.

Samuelson, Paul A. (1965). 'Proof that properly anticipated prices fluctuate randomly'. *Industrial management review* 6(2), pp. 41–49.

Sandroni, Alvaro (2000). 'Do Markets Favor Agents able to Make Accurate Predictions?' *Econometrica* 68(6), pp. 1303–1341. DOI: 10.1111/1468-0262.00163.

SEC (2014). *Equity Market Structure Literature Review. Part II: High Frequency Trading*. White Paper. March 18, 2014. URL: https://www.sec.gov/marketstructure/research/hft_lit_review_march_2014.pdf.

Shiller, Rober J. (2016). *Irrational Exuberance*. 3rd ed. Princeton University Press.

Shiller, Robert J. (1981). 'Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?' *American Economic Review* 71(3), pp. 421–436. DOI: 10.3386/w0456.

Simon, Herbert A. (1956). 'Rational choice and the structure of the environment.' *Psychological Review* 63(2), pp. 129–138. DOI: 10.1037/h0042769.

Simon, Herbert A. (1979). 'Rational Decision Making in Business Organizations'. *American Economic Review* 69(4), pp. 493–513.

Simon, Herbert A. (1997). *Administrative behavior: A Study of Decision Making Processes in Administrative Organizations*. 4th ed. The Free Press.

Tauchen, George E. and Mark Pitts (1983). 'The Price Variability-Volume Relationship on Speculative Markets'. *Econometrica* 51(2), pp. 485–505. DOI: 10.2307/1912002.

Taylor, Mark P. and Helen Allen (1992). 'The use of technical analysis in the foreign exchange market'. *Journal of International Money and Finance* 11(3), pp. 304–314. DOI: 10.1016/0261-5606(92)90048-3.

Thaler, Richard H. (1987). 'Anomalies: Weekend, Holiday, Turn of the Month, and Intraday Effects'. *Journal of Economic Perspectives* 1(2), pp. 169–177. DOI: 10.1257/jep.1.2.169.

Timmermann, Allan and Clive W. J. Granger (2004). 'Efficient market hypothesis and forecasting'. *International Journal of Forecasting* 20(1), pp. 15–27. DOI: 10.1016/s0169-2070(03)00012-8.

ISTITUTO
DI ECONOMIA

Sant'Anna
Scuola Universitaria Superiore Pisa