DISS. ETH NO. 20439

On the emergence of volatility, return autocorrelation and bubbles in Equity markets

A dissertation submitted to
ETH ZURICH

for the degree of
DOCTOR OF SCIENCES

presented by
GEORGES EMMANUEL HARRAS
Msc. in Physics, University of Fribourg
born on December 24^{th} 1980
Citizen of Luxembourg

accepted on the recommendation of
Prof. Dr. Didier Sornette, examiner
Prof. Dr. Thorsten Hens, co-examiner

2012

Acknowledgements

This thesis was carried out under the supervision of Didier Sornette to whom I owe my sincerest gratitude for his academic guidance, for the liberty that was given to me to enjoy the work on this thesis and the willingness to share his time for interesting discussions and insightful feedback. Among others, his great enthusiasm and relentless optimism had a great impact on me and my research. Also, I wish to thank Thorsten Hens for being my co-examiner and Gudela Grote for chairing my defense.

I am particularly grateful to Claudio J. Tessone, with whom I had the pleasure to work on my second paper and who not only contributed to deepening my knowledge of statistical physics but also showed me the power and beauty of the GNU/Linux world and how to properly code in various languages. Here special thanks are also due to Ryan Woodard, who introduced me and the whole Entrepreneurial Risks group to python and helped with many computational related problems, as well as critical feedback on many manuscripts.

For a nice office atmosphere and stimulating discussions, I want to thank Vladimir Filimonov, Moritz Hetzer, Andres Hüsler, Thomas Maillart, Moritz Müller and Shengsui Hu. For her enormous help with administrative tasks, I am indebted to Heidi Demuth, who did so without hesitation. I am grateful to Gilles Daniel and Riley Crane, for introducing me to life as a graduate student and all their advice, and to Peter Cauwels for his constructive feedback on major parts of this thesis.

Not on an academic, but on an equally important gastronomic level, I am grateful to Stocker for their quality food, their friendly staff and their six o'clock policy, without which many long hours would not have been enjoyable. I also want to thank my friends and my family for moral support and joyous moments outside of the office. Finally, I want to thank the reader for his or her interest in this thesis, in the hope that he or she will find its content insightful and of practical use for the questions at hand.

Contents

Summary

This thesis addresses three important aspects of price dynamics of publicly traded assets. For one, the emergence of a trending price is studied, resulting from the myopic optimization of socially influenceable investors, leading to the destabilization of prices and the growth of bubbles. My second contribution focuses on the volatility of price dynamics and, more generally, on the volatility of dynamic macroscopic observables, governed by a large number of interconnected units under the influence of a rapidly varying external signal. Whereas the first two contributions explore the first and second moment of collective/price dynamics via theoretical studies, my third contribution is an empirical study investigating the autocorrelation of daily price returns and its dependencies on other macroscopic variables such as volatility, long-term price movements and illiquidity.

My first scientific contribution consists of a financial market model, where the behavior of trading agents, interconnected by their social network, and the resulting price dynamics are investigated. Agents invest according to their opinion on future price movements, which is based on three sources of information, (i) public information, i.e. news, (ii) information from their social network and (iii) private information. In order to form the best predictor of future price movements, agents are continuously adapting their trading strategy to the current market regime by weighting the news and information from their peers according to their recent predicting performance. Paradoxically, it is their myopic adaptation to the current market regime which leads to a dramatic amplification of the price volatility and the occurrence of a bubble, followed by a crash. The model offers a simple reconciliation of the two opposite (herding versus fundamental) explanations for the origin of crashes within a single framework and shows that a crash is not a reaction to an extreme negative news event, but a sudden correction of an unsustainable high price. More general, this model shows that even with bounded rational and adapting agents, bubbles and crashes emerge naturally.

By reducing the complexity of the previous model, but keeping the same three basic influences, it is possible to apply this model to a very wide range of systems, generalizing the interpretation of the individual agent from an investor to any bistable entity, susceptible to its surrounding, a common and varying driving force and independent noise sources. This model, which is based on the kinetic Ising model, is a priori a physical model but can easily be related to social systems via the derived equivalence between the Ising model and a discrete choice model with social interactions. It is found that, independently of the shape of the driving force, increased levels of fluctuations in the macroscopic dynamics are observed for an intermediate noise strength (or coupling strength,

depending on the setup). Whereas for periodic forcing, the peak in the fluctuation amplitude corresponds to a pronounced amplification of the signal, with a strong correlation between the macroscopic dynamics and the driving force, this correlation is completely destroyed if the system is driven by a stochastic signal. This shows that even though these fluctuations are induced by the common forcing, the macroscopic dynamics have an endogenous origin. As an example of a system where this phenomenon can be observed, the social system of stock markets is proposed, explaining not only the excess of volatility observed in stock prices, but also the apparent absence of correlation between news and price changes and the persistence of volatility during times of crises.

The last part of this thesis contains an empirical study, motivated by the question of whether investors behave differently in different market regimes. For individual stocks traded on the New York Stock Exchange, I investigate the dynamics of the cross-sectional average of the first order autocorrelation of their daily returns and show that changes in the average autocorrelation of returns strongly correlate with prior changes in the cross-sectional volatility and market trends. It is found that return autocorrelation relates negatively to past volatility changes and positively to past market trends. This observation, which is a market-wide phenomenon, is persistent for over 20 years of data and also present in individual stocks. In contrast to the existing literature on return autocorrelation, illiquidity and bid-ask bounce can be rejected as driving forces behind the return autocorrelation dynamics. A behavioral origin of the phenomenon is proposed, where high volatility and bear markets lead to uncertainty and panic, reflected in overreacted behavior on a daily scale, whereas low volatility and bull markets lead to overconfidence, identified by price momentum.

Kurzfassung

Diese Dissertation behandelt drei wichtige Aspekte der Preisdynamik von öffentlich gehandelten Wertpapieren. Zum einen wird das Erscheinen von Preistrends behandelt, die aus einer kurzsichtigen Optimierung von sozial beeinflussbaren Investoren hervorgehen und zur Destabilisation von Preisen und zur Bildung wirtschaftlicher Blasen führen können. Meine zweiter Beitrag richtet sich auf die Volatilität von Preisen oder, im Allgemeinen, die Volatilität der makroskopischen Dynamik von Systemen die aus einer Sammlung von vielen untereinander verbundenen Einheiten bestehen und beeinflusst durch ein schnell variierendes externes Signal sind. Im Gegensatz zu den beiden ersten Beiträge, die dem ersten und zweiten Moment der Dynamik von Renditen oder kollektiven Systemen gewidmet sind, besteht mein dritter Beitrag aus einer empirischen Studie, welche die Abhängigkeit der Autokorrelation täglicher Renditen von gehandelter Wertpapier gegenüber andern makroskopischen Variablen untersucht, wie zum Beispiel Volatilität, langfristige Preisänderungen oder Illiquidität.

Um die Preisdynamik und das Verhalten von Agenten, die in ein soziales Netzwerk eingebunden sind, zu erforschen, wird ein Börsenmodell vorgestellt indem Investoren ein Wertpapier handeln und der Preis durch Angebot und Nachfrage verändert wird. Die Agenten investieren bezüglich ihrer Meinung zu bevorstehenden Preisänderungen. Diese Meinung basiert auf drei verschiedenen Informationen, (i) öffentliche Informationen, i.e., Neuigkeiten, (ii) Informationen aus ihrem sozialen Netzwerk und (iii) private Informationen. Um die Preisbewegungen möglichst gut voraussagen zu können, adaptieren die Agenten kontinuierlich ihre Handelsstrategien indem sie die verschiedene Informationsquellen bezüglich ihrer rezenten Leistung gewichten. Paradoxerweise ist es ihre kurzsichtige Anpassung an das herrschenden Marktregime das zu einer dramatischen Verstärkung der Preisfluktuationen und dem Aufkommen von Blasen führt, die von einem Börsenkrach beendet werden. Unser Modell vereinigt auf eine einfacher Art und Weise die zwei widersprüchlichen Erklärungen zum Ursprung von Börsenkrachen (Herdenverhalten und Nachrichten bezüglich des Fundamentalpreises) und zeigt, dass ein Krach nicht eine Reaktion zu extrem schlechten Nachrichten ist, sondern eine plötzliche Korrektur eines übermässig aufgeblähten Preises. Im Allgemeinen zeigt dieses Modell, dass auch mit rationalen und optimierenden Investoren, Börsenblasen und Krache ganz natürlich entstehen können.

Durch das Vereinfachen des vorherigen Modells, indem nur die drei Grundeinflüsse erhalten bleiben, ist es möglich das Modell auf eine breite Spannweite von Systemen anzuwenden. So kann man die Interpretation vom einzelnen Agent als ein Investor, auf jede beliebige bistabile Einheit verallgemeinern,

welche von ihrem Umfeld, von einem gemeinen wechselnden Treiben und einem unabhängigen Rauschen beeinflusst wird. Dieses Modell, welches auf dem dynamische Ising Modell basiert ist, ist also a priori ein physikalisches Modell, kann jedoch einfach auch für soziale Systeme verwendet werden durch die Äquivalenz zwischen dem Ising Modell und diskrete Wahl Modell mit sozialen Beeinflussung. Wir dokumentieren, unabhängig von der Natur der treibenden dynamischen Kraft, dass für intermediär Intensitäten von Rauschen (oder Wechselwirkung) ein starkes Ansteigen der Fluktuationen der makroskopischen Dynamik. Für eine periodisch treibende Kraft entspricht dieses Maximum an Fluktuationen dem Verstärken des Signal, mit einer ausgeprägten Korrelation zwischen dem Signal und der makroskopischen Dynamik. Diese Korrelation wird jedoch stark verringert wenn das System einem aperiodischen Treiben unterliegt. Dies zeigt, dass obwohl die Fluktuationen durch das gemeine Treiben ausgelöst wird, sie ein endogenen Ursprung haben. Als Beispiel für ein System wo dieses Phänomen beobachtet wird, schlagen wir den Finanzmarkt vor, womit nicht nur die übermässig Volatilität von Preisen erklärt werden kann, jedoch auch die Abwesenheit der Korrelation zwischen Neuigkeiten und Preisänderungen, sowie die anhaltende Volatilität in Krisenzeiten.

Der letzte Teil dieser Dissertation enthält eine empirische Studie, angeregt durch die Frage, wie und ob Investoren ihr Verhalten in verschiedenen Marktphasen verändern. Für individuelle Wertpapiere die auf dem New York Stock Exchange gehandelt werden, untersuche ich die Änderungen des querschnittlichen Mittels der Autokorrelation ersten Grades täglicher Renditen und zeige, dass diese Änderungen stark korreliert sind mit vorhergehenden Änderungen der querschnittlichen Volatilität oder des Preises. Die Renditenautokorrelation Änderungen sind negative proportional zur vergangenen Änderungen der Volatilität, und positive proportional zu vergangenen Preisänderungen. Diese Beobachtung, welche ein marktweites Phänomen ist, besteht anhaltend seit über 20 Jahren und wird auch für individuelle Papiere gefunden. Im Gegensatz zu der bestehenden Literatur zum Thema von Renditenautokorrelation, kann Illiquidität und Geld-Brief-Sprung als Ursache verworfen werden. Ein Erklärung aus der Verhaltensökonomie wird vorgeschlagen, in der eine hohle Volatilität und fallende Preise zu Unsicherheiten und Panik führen, was in Überreaktioen während dem täglichen Handeln wiedergespiegelt wird, wogegen niedrige Volatilität und steigende Preise zu übermässigem Selbstvertrauen und scheinbarer Sicherheit führen, was sich durch einen Trend im Preis zeigt.

Chapter 1

Introduction

Financial markets serve a multitude of purposes and are of paramount importance in open economies. Among others, they facilitate the raising of capital for companies, allow for the transfer of risk and liquidity, enable international trade and give private persons the possibility to expose themselves to the dynamics of national and international economies. It is, however, the emergence of a price and its dynamics due to the collective acting of large numbers of individuals, which is investigated in this work. As such, financial markets are global polling instruments, which give researchers the opportunity to study human behavior in their quest to make profitable investments. Even though this work concentrates on equity markets, the general concepts derived and studied here, also apply to markets in general.

As this manuscript is a cumulative thesis, i.e., a collection of research papers, the major part of my scientific contribution is concentrated in three self-contained papers, of which two are already published in peer-reviewed journals. In order to put my work into context, the papers are preceded by an extended literature review. The outline of the thesis is as follows:

- The objective of Chapter 2 is to derive the basic concepts on which the models of Chapter 4 and 5 are based. Thematically, Chapter 2 is divided in two parts. In the first part, from Section 2.1 to Section 2.3, the Boltzmann framework is derived from first principles and applied to solve the Ising model under the mean-field approximation. In the second part of Chapter 2, the origins of the discrete choice models are introduced and the relation between discrete choice models with social interactions, and the Ising model is shown. The Ising model was originally developed to explain the magnetic properties of ferromagnets, but can also, by reinterpreting its components in the context of decision makers exposed to a binary choice, be considered as a simple model of describing the competition between the ordering force of imitation or contagion and the disordering impact of private information or idiosyncratic opinions that promotes heterogeneous decisions.

- Chapter 3 contains an extended review on the scientific literature on bubbles and crashes, and related subjects such as momentum, the overreaction-underreaction phenomenon and imitation among analysts and institutional investors. Both theoretical models, as well as empirical studies are summarized, revealing the rich and diverse research that was performed on this subject. The reviewed empirical studies validate the basic assumptions upon which the models in Chapter 4 and 5 are build.

- A simple model of bounded rational agents, is presented in Chapter 4, which focuses on the emergence of bubbles and consequential crashes, and investigates how their proximate triggering factor might relate to their fundamental origin, and vice versa. Agents invest according to their opinion on future price movements, which is based on three sources of information, (i) public information, i.e. news, (ii) information from their friendship network and (iii) private information. Agents continuously adapt their trading strategy to the current market regime by weighting each of these sources of information according to its recent predicting performance. It is found that bubbles originate from a random lucky streak of positive news, which, due to a feedback mechanism of these news on the agents strategies develop into a transient collective herding regime. After this self- amplified exuberance, prices reach an unsustainable high value, which is corrected by a crash. These ingredients provide a simple mechanism for the excess volatility documented in financial markets. Paradoxically, it is the attempt for investors to adapt to the current market regime which leads to a dramatic amplification of the price volatility. A positive feedback loop is created by the two dominating mechanisms (adaptation and imitation) which, by reinforcing each other, result in bubbles and crashes. The model offers a simple reconciliation of the two opposite (herding versus fundamental) proposals for the origin of crashes within a single framework and shows that even with rational and adapting agents, bubbles and crashes can naturally emerge.

- In Chapter 5, a novel phenomenon of an increased level of fluctuations is presented, which is found for the collective dynamics of a system composed of many bistable units in the presence of a rapidly varying external signal, and intermediate noise levels. The archetypical signature of this phenomenon is that –beyond the increase in the level of fluctuations– the response of the system becomes uncorrelated with the external driving force. Numerical simulations and an analytical theory of a stochastic dynamical version of the Ising model on regular and random networks demonstrate the ubiquity and robustness of this phenomenon, which is argued to be a possible cause of excess volatility in financial markets, of enhanced effective temperatures in a variety of out-of-equilibrium systems, and of strong selective responses of immune systems of complex biological organisms.

- An empirical study, inquiring investors behavior in different market regimes, is presented in Chapter 6. For individual stocks traded on the New York Stock Exchange, the dynamics of the cross-sectional average of the first order autocorrelation of their daily returns is investigated and it is shown that changes in the average autocorrelation of daily returns strongly correlate with prior changes in the cross-sectional volatility and market trends. It is found that return autocorrelation relates negatively to past volatility changes and positively to past market trends. This observation, which is a market-wide phenomenon, is persistent for over 20 years of data and also present in individual stocks. In contrast to the existing literature on return autocorrelation, illiquidity and bid-ask bounce can be rejected as driving forces behind the return autocorrelation dynamics. A behavioral origin of the phenomenon is proposed, where high volatility and bear markets lead to uncertainty and panic, reflected in overreacted behavior on a daily scale, whereas low volatility and bull markets lead to overconfidence, identified by price momentum. In order to address the non-stationarity of some of the analyzed time-series, a very powerful and yet intuitive method had been developed and is used to compute meaningful correlations between time-series with various memories.

- An overall conclusion is given in Chapter 7, summarizing my contributions to the fields of finance and physics.

Chapter 2

The Ising model and random utility

In the field of thermodynamics and statistical mechanics, the objective is to analyze systems composed of a large number of simple microscopic units (particles, magnetic moments, ...) and to make statements about the system's macroscopic properties and dependencies. The familiarity of physicists with studying the macroscopic properties of many-body systems despite the many unknowns on the micro-level is one of the main reasons for their interest in social sciences, especially finance where there is an abundance of data to test models against.

This chapter is divided into two parts. In the first part the foundations of the Ising model will be reviewed. In order for the chapter to be self-consistent, the Boltzmann formalism will first be introduced, which constitutes the basics of statistical mechanics. This introduction will be kept as concise as possible, as statistical mechanics is not the subject of this thesis. For detailed treatment of the subject, see the book by Greiner et al. (1995) or Reif (1965).

In the second part of this chapter, the concepts of random utility models and discrete choice with social interactions are introduced and their relation to the Boltzmann formalism and the Ising model is established. The framework reviewed in the chapter constitute the basis of the models studied in Chapter 4 and 5, and of a rich and long list of models of the social and economic sciences, investigating the behavior of interacting agents (Schelling, 1971; Föllmer, 1974; Galam and Moscovici, 1991; Blume, 1993; Brock, 1993; Kirman, 1993; Lux, 1995; Orléan, 1995; Durlauf, 1999; Brock and Durlauf, 2001; Michard and Bouchaud, 2005).

2.1 Boltzmann statistics and the Canonical Ensemble

The exact state of any system, which specifies all its properties, is characterized by its position in phase space. A 3-dimensional (3D) system composed of N particles has a phase space of $6N$ dimensions, as every particle is fully described by its position (3D-vector) and its momentum (3D-vector). For a grid of spins (magnetic moments), which can take the values of ± 1, the phase space is a set of 2^N micro-states.

In order to compute the probabilities to find the system in a certain position in phase space, different kinds of systems can be considered, grouped by their level of openness. They can be completely closed, i.e. isolated, such that the phase space is constrained by the available energy. The ensemble containing all micro-states with that particular energy is called the *micro-canonical ensemble*. The next level of openness is achieved by submerging the system in a heat bath, keeping it at a constant temperature and allowing for the exchange of energy between the heat bath and the system. The ensemble of micro-states, unconstrained by their energy, is called the *canonical ensemble*. By allowing the system to not only exchange energy, but also particles with its surrounding, the set of micro-states is called the *grand canonical ensemble*. As the Ising model is defined with a fixed temperature and constant number of spins, we will focus on the canonical ensemble and derive the probabilities for its different micro-states.

Via the exchange with a heat bath[1], a system can take any amount of energy. A situation where the system has given all its energy to the heat bath is possible, as well as the situation where the system has absorbed all of the heat bath's energy, although the likelihood of these configurations is, as we will see, very small. A priori every single micro-state, a particular position in phase space, is equally likely to occur. We will however find that, given some constraints, certain macroscopic characteristics are more likely to occur as they can be obtained by a larger number of micro-states.

Let us consider $\mathcal{N}$ identical systems (replicas) composed of N units, whose phase space is either discrete or, if continuous, divided into small same-sized and numbered cells. Let n_i be the number of replicas that are in phase space state (or cell) i and which have the energy e_i. Then

$$\mathcal{N} = \sum_i n_i, \tag{2.1}$$

where the summation is performed over the entire phase space. As $\mathcal{N}$ is considered large, $p_i = n_i/\mathcal{N}$, is the probability for finding a system in state i. Even though in general, due to the heat bath, the system can take any amount of energy, in equilibrium however, the system will have an average energy, given by

$$U = \langle e_i \rangle = \sum_i p_i e_i = \frac{1}{\mathcal{N}} \sum_i n_i e_i. \tag{2.2}$$

For a certain arrangement of the systems $\{n_i\} = \{n_1, n_2, \ldots\}$, we can enumerate the systems and compute the number of different possible configurations. As there are exactly $\mathcal{N}!$ permutations of all the systems and $n_i!$ for every state, the number of rearrangements of $\{n_i\}$ are

$$W\{n_i\} = \frac{\mathcal{N}!}{\prod_i n_i!}. \tag{2.3}$$

It is said that a system, described by Eq. (2.3), obeys *Boltzmann statistics*. Given that every elementary phase space state (or cell) has the same probability, the arrangement that maximizes W corresponds to the most probable

[1] The heat bath is considered large compared with the size of the system.

distribution. For large integer n, Stirling's approximation,

$$\ln n! = n \ln n - n, \tag{2.4}$$

can be used to simplify the factorials of Eq. (2.3), leading to

$$\ln W\{n_i\} = \mathcal{N} \ln \mathcal{N} - \mathcal{N} - \sum_i [n_i \ln n_i - n_i]. \tag{2.5}$$

To find the extrema under constraints of Eq. (2.1) and (2.2), the method of Lagrange multipliers is used, such that

$$\Lambda\{n_i\} = \ln W\{n_i\} + \lambda(\sum_i n_i - \mathcal{N}) - \beta(\sum_i n_i e_i - \mathcal{N}U) \tag{2.6}$$

has to be maximized with respect to $\{n_i\}$. Here the minus sign in front of β is arbitrary, but it will be advantageous in the following section. Differentiating Eq. (2.6) and equating it to zero returns

$$\begin{aligned} \frac{d\Lambda\{n_i\}}{d\,n_i} &= -\ln n_i + \lambda - \beta e_i = 0 \\ \Leftrightarrow \;\; n_i &= e^{\lambda}\, e^{-\beta e_i} \end{aligned} \tag{2.7}$$

Using Eq. (2.1), we find that $e^{\lambda} = \mathcal{N} / \sum_i e^{-\beta e_i}$, leading to

$$p_i = \frac{n_i}{\mathcal{N}} = \frac{e^{-\beta e_i}}{\sum_i e^{-\beta e_i}}, \tag{2.8}$$

where $e^{-\beta e_i}$ is called the *Boltzmann factor*, and

$$Z = \sum_i e^{-\beta e_i} \tag{2.9}$$

is the *canonical partition function*, where the letter Z comes from the German word Zustandssumme. The Lagrange multiplier β relates to the average energy in the system (via Eq. (2.2)), which depends on the temperature for thermodynamic systems. As constant temperature is assumed due to the heat bath of the canonical system, as such the average energy of the system will also be constant. The connection between the Boltzmann framework and thermodynamics is realized by identifying that $\beta = 1/kT$, which is the subject of the next section. The intuition behind Eq. (2.8) is that the probability of a system having a certain amount of energy decreases exponentially with the energy, with a rate inversely proportional to the given temperature.

2.2 Relation to thermodynamics

From the second law of thermodynamics, we know that for constant volume V and number of particles N

$$\left(\frac{\partial S}{\partial U}\right)_{V,N} = \frac{1}{T}, \tag{2.10}$$

where S is the entropy of the system, defined by Boltzmann as

$$S_B = k \ln W, \tag{2.11}$$

with k the Boltzmann's constant and W the number of micro-states corresponding to a given macro-state. Boltzmann's entropy is defined only for isolated systems in equilibrium, i.e., system in the micro-canonical ensemble, where every micro-state is equally probable to be occupied by the system. However for system where exchanges are allowed, not all possible micro-states are equally probable, as seen in the previous section for the case of a system able to exchange energy with a heat bath. Gibbs introduced a generalized formulation of entropy, given by

$$S = -k \sum_i p_i \ln p_i, \tag{2.12}$$

which is valid for any ensemble and where p_i is the probability of the system being in the micro-state i of the $6N$-dimensional phase space and having the energy e_i. Boltzmann's formulation, for the micro-canonical ensemble, can be recovered by using uniform probability across states ($p_i = 1/W, \forall i$) and that the summation is performed over all W micro-states.

Now substituting the probabilities computed in Section 2.1, the entropy can be rewritten as

$$\begin{aligned} S &= -k \sum_i p_i \ln p_i \\ &= -k \sum_i \frac{e^{-\beta e_i}}{Z}\big(-\beta e_i - \ln(Z)\big) \\ &= k\beta \sum_i \frac{e_i e^{-\beta e_i}}{Z} + k \ln Z \sum_i \frac{e^{-\beta e_i}}{Z} \\ &= k\beta U + k \ln Z \end{aligned} \tag{2.13}$$

It can be shown that

$$\frac{\partial \ln Z}{\partial \beta} = -U \tag{2.14}$$

and combining Eq. (2.10), Eq. (2.13) and Eq. (2.14) yields

$$\begin{aligned} \frac{1}{T} &= \frac{\partial S}{\partial U} \\ &= k\beta + kU\frac{\partial \beta}{\partial U} + k\frac{\ln Z}{\partial \beta}\frac{\partial \beta}{\partial U} \\ &= k\beta, \end{aligned} \tag{2.15}$$

relating the results from the ensemble theory with thermodynamics.

2.3 Ising model

2.3.1 Introduction

Ernest Ising, together with his adviser Wilhelm Lenz, proposed the Ising model in 1925 as a simple model for ferromagnetic behavior (Brush, 1967). Real ferromagnets have a complicated structure and have to be studied via band theory. The Ising model, on the other hand, offers a strongly simplified approach, which still embodies the main qualitative features of a ferromagnet. The model consists of a large number of magnetic moments connected by a regular grid in d

dimensions, i.e. a hypercube. The magnetic moments, called spins, can only take two values (± 1), which represent the direction in which they point (up or down). Each spin interacts with its direct neighbors and with an external magnetic field, leading to a Hamiltonian given by

$$H = -J \sum_{i,j \in <i,j>} s_i s_j - h \sum_{i=i}^{N} s_i, \tag{2.16}$$

where s_i are the different spins, h is the external magnetic field, $J > 0$ is the interaction strength for a ferromagnet and $< i, j >$ represents the nearest neighbors. The number of nearest neighbors depends on the dimension of the system and will be referred to by $z = 2d$.

Not only was the Ising model used to study the behavior of ferromagnets, but since its introduction it was increasingly used as a toy model of phase transitions, with the average magnetization being the order parameter. In one dimension, the model does however not experience a phase transition at finite temperatures, but must be cooled down to zero temperature ($T = 0$) for the average magnetization to be different from zero. The intuition behind this result is that for any finite temperature it is possible for a single spin to divide the infinite long chain into two regions of opposite magnetization, destroying any finite magnetization. This fact demotivated Ising to continue his research on the model. Later it was however shown that in two dimensions the model experiences a transition at finite temperatures, which was confirmed by the seminal paper of Onsager (1944), who was able to exactly solve the model in 2D. There exists no closed form solution for three dimensions, which is thought to be impossible to derive. The behavior under 3D is, however, also well known as many partial or approximate solutions can be developped.

Knowing the Hamiltonian and that the Ising model is kept at a constant temperature, the partition function of the canonical ensemble of Eq. (2.9) is to be used, yielding

$$Z = \sum_{s_1=\pm 1} \sum_{s_2=\pm 1} \dots \sum_{s_N=\pm 1} e^{-\beta H}, \tag{2.17}$$

with $\beta = 1/kT$ and the energy of the system given by the Hamiltonian. Due to the nearest neighbor interaction, a lot of work is required to further use Eq. (2.17) without any approximations.

On simple approximation is to neglect spin interaction, i.e. $J = 0$. In this situation the spins are independent from each other and the system can be described by a single particle Hamiltonian,

$$H = -hs_i. \tag{2.18}$$

As the probabilities for the spin values are given by Eq. (2.8), the average magnetization is given by

$$\langle s \rangle = \frac{e^{+\beta h} - e^{-\beta h}}{e^{+\beta h} + e^{-\beta h}} = \tanh(\beta h), \tag{2.19}$$

where no phase transition can be observed.

2.3.2 Mean-field approximation

An alternative method to study the Ising model is the mean-field approximation, where the influence of the neighbor spins is simplified by replacing their value by an average magnetization. The here derived mean-field Ising model is inspired by Kadanoff (2000).

The approximation is performed by neglecting feedback effects between the spins and separating one particular spin i from the rest, which is performed by rearranging the Hamiltonian of Eq. (2.16) into

$$\begin{aligned} H &= -Js_i \sum_{j\in<i,j>} s_j - hs_i + c \\ &= -s_i\Big(\frac{K}{z}\sum_{j\in<i,j>} s_j + h\Big) + c, \end{aligned} \tag{2.20}$$

where j sums over the neighbors of s_i and c absorbs the contributions of all other spins to the Hamiltonian, which are independent of and do not ifluence spin i. The coupling strength J was replaced by K/z, such that the dimensionality of the system does not change the impact of the neighbor-interaction relative to the magnetic field. Identifying the influence of spin i as an effective field

$$h_{\text{eff}} = \frac{K}{z}\sum_{j\in<i,j>} s_j + h \tag{2.21}$$

Eq. (2.20) can be rewritten as

$$H = -s_i\, h_{\text{eff}} + c. \tag{2.22}$$

As spin i is only controlled by its effective field, h_{eff}, the probability for finding the spin i in either direction can be computed via Eq. (2.8), yielding

$$p(s_i) = \frac{e^{-\beta s_i h_{\text{eff}}}}{\sum_{s_i=\pm 1} e^{-\beta s_i h_{\text{eff}}}}, \tag{2.23}$$

such that spin i's average value over different realizations is given by

$$\langle s_i\rangle = \tanh(\beta\, h_{\text{eff}}). \tag{2.24}$$

By replacing the s_j in Eq. (2.21) by their average $\langle s_j\rangle$, which assumes that their fluctuations are well behaved, and by assuming that the system is translationally invariant (i.e., every spin is equivalent) such that $\langle s_i\rangle$ can be replaced by the ensemble average, the famous Ising model mean-field solution is obtained,

$$\begin{aligned} m = \frac{1}{N}\sum_i^N s_i &= \tanh(\beta\, h_{\text{eff}}) \\ &= \tanh(\beta\, Km + \beta\, h). \end{aligned} \tag{2.25}$$

The accuracy of Eq. (2.25) will increase with z, as that means that the average will be computed based on more spins, decreasing its fluctuations and increasing the likelihood of being a good representation of the entire system. In the case of infinite-range interaction, where every spin is connected to every other spin, the mean-field approximation is exact.

In the following sections, discrete choice models and the concept of random utility models will be introduced. As such, this paragraph indicates the split of the chapter in physics and social sciences. As will become clear in the last section of this chapter, these two subjects, which a priori seem distinct from one another, have some common intersection opening the possibility to apply the wealth of knowledge of physical models to social and economic models.

2.4 Discrete choice models

Discrete choice models are used to model a situation, where a decision maker has to select one choice out of a set of n alternatives. Compared to the standard economic setup, where agents can choose an amount from a continuous value (e.g., the amount of money invested in the stock market, given a set of explanatory variables), a different framework is needed for a discrete set of choices (e.g., taking the car, bus or bike for commuting, given a set of explanatory variables). For an in-depth treatment of the subject of discrete choice models,see to the book of Train (2003).

First the concept of random utility will be introduced, which is used to derive the most prominent dicrete choice models, the Probit and the Logit model. A stronge resemblence between the Logist model and the Boltzmann statistics will be observed. Later, a binary choice model of socially interacting agents will be introduced, and its equivalence with the Ising model will be shown, creating a connection between studies on Ising-like systems and collective behavior of social decision makers.

2.4.1 Random Utility

A Random Utility Model (RUM) is the standard framework used for the modeling of discrete choice scenarios. The decision maker has to choose one alternative out of a set X of n possible ones. For each alternative, $x \in X$, the decision maker obtains the utility (payoff) $U(x)$. The decision maker will choose the alternative which maximizes his/her utility.

On the other side, the researcher who wishes to model the decision maker's behavior does not know the exact utility attached to the various alternatives. Instead he/she possesses a set of attributes and explanatory variables describing the decision maker and its surrounding. These variables are used to compute the representative utility, $V(x)$, which obviously is not identical to the utility perceived by the decision maker, $V(x) \neq U(x)$. The utility can be decomposed as

$$U(x) = V(x) + \epsilon(x), \tag{2.26}$$

where $\epsilon(x)$ captures all factors which are not included in $V(x)$, i.e., the unobserved utility. This decomposition is fully general as $\epsilon(x)$ is defined as the difference between the modeler's representative utility and the decision maker's real utility. As $\epsilon = \{\epsilon(x) \mid x \in X\}$ is unknown to the researcher, it will be assumed random, hence the name, random utility model.

The probability for the decision maker to chose x over all other alternatives

$Y = X - \{x\}$ is given by

$$P(x) = \text{Prob}\big(U(x) > U(y)\,,\ \forall y \in Y\big) \tag{2.27}$$
$$= \text{Prob}\big(V(x) - V(y) > \epsilon(y) - \epsilon(x)\,,\ \forall y \in Y\big)$$
$$= \int_\epsilon I\Big(V(x) - V(y) > \epsilon(y) - \epsilon(x)\,,\ \forall y \in Y\Big) f(\epsilon) d\epsilon, \tag{2.28}$$

where $f(\epsilon)$ is the multivariate distribution of the unknown factors and $I(\cdot)$ the indicator function. The various discrete choice models are obtained by different assumptions of $f(\epsilon)$. It is assumed that the ϵ are i.i.d., such that their distribution sets the discrete choice model.

2.4.2 Probit model

While studying the relation between physical stimuli and induced psychological sensation, Thurstone (1927) was the first to introduce a discrete choice model. The aim of his study was to investigate the answers of participants to binary questions (e.g., which of the two weights is heavier?), leading to the *law of comparative judgment*. Participants are exposed to two stimuli x_1 and x_2, whose intensity is given by $V(x_i)$. Due to neuronal or mechanical irregularities, the sensation perceived by the participants amounts to

$$U(x_i) = V(x_i) + \epsilon_i,\ i \in \{1, 2\}, \tag{2.29}$$

where ϵ_i is a noise term, which was assumed to follow a normal distribution. The assumption of the normal distribution is not take out of some ulterior motive nor is it based on any knowledge of the "measurement error". This choice seems however intuitively valid as, after the central limit theorem, the normal distribution results from the repeated summation of a random variable from any distribution whose variance is defined.

Participants were asked to select the stimulus with the highest intensity. Given Eq. (2.29), the probability of selecting stimulus x_1 is

$$p(x_1) = \text{Prob}(\max_{i\in\{1,2\}} U(x_i))$$
$$= \text{Prob}(U(x_1) > U(x_2))$$
$$= \text{Prob}(V(x_1) - V(x_2) > \epsilon_2 - \epsilon_1)$$
$$= F_n(V(x_1) - V(x_2)), \tag{2.30}$$

where $F_n(x)$ is the cumulative distribution function (CDF) of the random variable $\epsilon_2 - \epsilon_1$, which is also normal by construction. A similar formalism can be derived for more than two choices, resulting in a probability given by Eq. (2.28), where $f(\epsilon)$ represents the normal distribution. In case all ϵ_i are i.i.d., this discrete choice model is referred to as the probit model.

2.4.3 Logit model

The most prominent discrete choice model, the logit model, is based on Luce's choice axiom and was introduced by Luce (1959). Suppose that X represents the complete set of possible choices and $S \subseteq X$, a subset of these choices. If for any element $x \in X$ there is a finite probability of being chosen, $p_X(x) \in]0,1[$, then Luce's choice axiom is defined as

$$p_X(x) = p_S(x) p_X(S), \tag{2.31}$$

with $p_X(S)$ the probability of choosing any element in S from the set of X. The axiom states that the probability of choosing one possibility over another from a set of alternatives is not affected by the addition or removal of other alternatives, leading to the name of "independence from irrelevant alternatives" (IIA), which is how the axiom is referred to in more recent literature. That this relation follows immediately from Eq. (2.31), can be seen by rewriting Eq. (2.31) for any other element $y \in X$ and equating the last factor in the RHS, leading to

$$\frac{p_S(x)}{p_S(y)} = \frac{p_X(x)}{p_X(y)}. \tag{2.32}$$

The assumption of IIA is valid in many scenarios, but inconsistencies arise when alternatives are added to the pool, which are very similar to choices already available, a situation which is of no concern in this work[2].

Another important formulation of the axiom can be found by rewriting Eq. (2.31) as

$$p_S(x) = \frac{p_X(x)}{\sum_{y \in S} p_X(y)}, \tag{2.33}$$

where $P_X(S)$ is replaced by the sum of the probabilities of all elements in S. As proven in Luce (1959), a set of probabilities satisfies the Choice Axiom if and only if there exists a set of numbers $\{v(x)\}$, which attach a weight to every alternative, such that

$$p_S(x) = \frac{v(x)}{\sum_{y \in S} v(y)}, \tag{2.34}$$

for every $x \in S \subseteq X$. The weights $\{v(x)\}$ are uniquely determined by the set of probabilities $\{p_X(x)\}$ up to a multiplication by a constant.

Marschak (1959) showed that probabilities, specified by IIA, are consistent with random utility models, i.e. that the weights in Eq. (2.34) can be obtained from a function of the representative utility $V(x)$, such that the probabilities of Eq. (2.34) can be expressed in the form of a random utily model as given by Eq. (2.27). One way of obtaining the weights of Eq. (2.34) is by setting $v(x) = e^{V(x)}$. As an example of the simple case of $S = \{x, y\}$, the probability for choosing x is

$$\begin{aligned} p_{\{x,y\}}(x) &= \frac{e^{V(x)}}{e^{V(x)} + e^{V(y)}} \\ &= \frac{1}{1 + e^{-(V(x) - V(y))}} \\ &= F_l(V(x) - V(y)), \end{aligned} \tag{2.35}$$

where $F_l(x)$ is the CDF of the logistic distribution. Here the similarity between Eq. (2.35) and Eq. (2.30) it to be noted, indicating that if the difference between

[2]The problem was fist mentioned by Debreu (1960) and is now known under the name of the "red bus, blue bus" problem.

the random terms of the utilities $U(x)$ and $U(y)$ is distributed according to a logistic distribution, the resulting probabilities are consistent with IIA.

More generally, it was Holman and Marley (as cited in Luce and Suppes (1965)) who showed that if the unknown utility, $\epsilon(x)$, is distributed according to the double exponential distribution, also called the Gumbel distribution, which has a CDF given by

$$F_G(x) = e^{-e^{-(x-\mu)/\gamma}} \tag{2.36}$$

with μ and $\gamma > 0$ constants, then the difference of the unknown utility will be logistically distributed and the resulting model will be equivalent to IIA for any choice scenario, not only pair comparison.

The other direction was proven by McFadden (1974), who showed that if the probability satisfies the IIA condition, then the unknown utility has to be distributed according to the Gumbel distribution.

Next, the relation between the random utility model and IIA via Gumbel distributed unobserved utilities will be derived. Eq. (2.27) can be rewritten as

$$\begin{aligned} p(x) &= \text{Prob}\big(U(x) > U(y)\,,\ \forall y \in Y\big) \\ &= \text{Prob}\big(V(x) - V(y) > \epsilon(y) - \epsilon(x)\,,\ \forall y \in Y\big). \end{aligned} \tag{2.37}$$

From the RHS of Eq. (2.37) together with their i.i.d. character, it follows that it is the difference of the unobserved utilities that sets the probability. As such the mean of the unobserved utilities, is of no importance and the parameter μ, which controls the mean of the Gumbel distribution can be set to zero without loss of generality. With $\mu = 0$, the distribution of random utility is left with one tunable parameter γ, which controls its variance, given by $(\pi\gamma)^2/6$. By using the definition of the CDF, Eq. (2.37) can be written as

$$\begin{aligned} p(x) &= \text{Prob}\big(V(x) - V(y) + \epsilon(x) > \epsilon(y)\,,\ \forall y \in Y\big) \\ &= \int_{-\infty}^{\infty} \left(\prod_{y \in Y}^{n} e^{-e^{-(V(x)-V(y)+\epsilon(x))/\gamma}} \right) f_G\big(\epsilon(x)\big) d\epsilon(x), \end{aligned} \tag{2.38}$$

where $f_G(x) = \frac{1}{\gamma} e^{-x/\gamma} e^{-e^{-x/\gamma}}$, the probability density function of the Gumbel distribution with $\mu = 0$. Performing a change of variable, $u = e^{-\epsilon(x)/\gamma}$, together with the corresponding adaptation of the integration borders, Eq. (2.38) can be rewritten as

$$\begin{aligned} p(x) &= \int_{-\infty}^{\infty} \left(\prod_{y \in Y}^{n} e^{-e^{-(V(x)-V(y)+\epsilon(x))/\gamma}} \right) \frac{1}{\gamma} e^{-\epsilon(x)/\gamma} e^{-e^{-\epsilon(x)/\gamma}} d\epsilon(x) \\ &= \int_{0}^{\infty} \left(\prod_{y \in Y} e^{-e^{-(V(x)-V(y))/\gamma} u} \right) e^{-u} du \\ &= \int_{0}^{\infty} e^{-\left(u \sum_{y \in Y} e^{-(V(x)-V(y))/\gamma}\right)} e^{-u} du \\ &= \int_{0}^{\infty} e^{-u\left(1+\sum_{y \in Y} e^{-(V(x)-V(y))/\gamma}\right)} du \\ &= \frac{1}{1+\sum_{y \in Y} e^{-(V(x)-V(y))/\gamma}} \\ &= \frac{1}{1+e^{-V(x)/\gamma} \sum_{y \in Y} e^{V(y)/\gamma}}. \end{aligned} \tag{2.39}$$

Now, by multiplying both sides of the ratio by $e^{V(x)/\gamma}$, keeping in mind that $Y = X - x$, the well known logit formulation can be recovered,

$$p(x) = \frac{e^{V(x)/\gamma}}{\sum_{y \in X} e^{V(y)/\gamma}}, \tag{2.40}$$

which obviously fulfills IIA and defines the probabilities for the logit model. Interestingly Eq. (2.40) bears a strong resemblance to Eq. (2.8), derived in Section 2.1, which specifies the probability of finding the system in a state i with energy e_i at a given temperature. As such, the maximization of entropy together with the constraint on the average energy is equivalent to maximization of the representative utility in a discrete choice system, where the "temperature" is proportional to the standard deviation of the unobserved utility, which tunes the scale of the unboserved utility.

2.4.4 Discrete choice with social interaction

The different variables which influence the utility of the decision maker have not been considered up to this point. As will be discussed in Section 3.5.2, there exist many scenarios in which individuals are positively influenced by the people around them, a behavior for which there are strategically rational as well as psychological reasons. Even if the intrinsic utility for one alternative is strongly heterogeneous among the entire population, the strong desire for conformity in humans (as well as other social animals) can result in the majority choosing that one alternative. In case of social influence, the total utility of the decision maker's alternatives will depend on the choices of the decision maker's group of reference.

Consider a population of N decision makers who are interacting among each other and have to take an individual choice out of a set of n alternatives. As introduced in Section 2.4.1, decision makers will select the alternative which maximizes their utility. Under social influence, the simplest form for the utility of choice x for individual i, denoted x_i, will consist of three terms,

$$U(x_i) = v(x_i) + S(x_i, \mu_i^e(x_{-i})) + \epsilon(x_i), \tag{2.41}$$

where $\mu_i^e(x_{-i})$ represents the set of agent i's expectation, at the time of his decision making, on the choices taken by his surrounding (excluding i). There are three terms contributing to agent i's total utility of choice x: the private utility, $v(x_i)$, which can be estimated by explanatory variables; the social utility, given by $S(x, i, \mu_i^e(x_{-i}))$; and the unobserved utility, represented by $\epsilon(x_i)$ as introduced in Eq. (2.26). The formulation of the total utility of socially interacting decision makers as given in Eq. (2.41) is in agreement with the proposed utility in a prominent theoretical study of discrete choice with social interaction by Brock and Durlauf (2001), which also influenced the general structure of this section.

The simplest, and also the most common, discrete choice scenario is the one where the decision maker has to choose between two alternatives (going to university: yes/no, having a child: yes/no, US presidential elections: Democrats/Republican, ...), resulting in a binary choice model with $x \in \{-1, +1\}$. For simplicity, it is assumed that every acquaintance is equally important and that all decisions makers are acquaintances of each other such that $\mu_i^e(x_{-i})$ will be approximated by

$$\bar{x}_i^e = \frac{1}{N-1} \sum_{j \neq i}^{N} x_{i,j}^e, \tag{2.42}$$

where $x_{i,j}^e$ represents agent i's expectation of agent j's choice. As such, $\bar{x}_i^e$ denotes the average choice of the population as expected by agent i.

On the basis that the decision makers long for conformity, the simplest form that their social utility term can take is

$$S(x, i, \bar{x}_i^e) = J x_i \bar{x}_i^e, \tag{2.43}$$

such that they get positive social utility when x_i and $\bar{x}_i^e$ have the same sign. The next simplest alternative is

$$\begin{aligned} S(x, i, \bar{x}_i^e) &= -\frac{J}{2}(x_i - \bar{x}_i^e)^2 \\ &= J x_i \bar{x}_i^e - \frac{J}{2}(1 + (\bar{x}_i^e)^2, \end{aligned} \tag{2.44}$$

where they want to minimize their distance to the concensus and where the fact that $x_i^2 = 1$ was used. By comparing Eq. (2.43) and Eq. (2.44), it is noted that even though they differ in level, the dependence of x_i coincides and as such, similar behavior is expected. Due to this similarity only Eq. (2.43) will be considered for this work, see Brock and Durlauf (2001) for further detail on Eq. (2.44).

For the private utility, $v(x_i)$, a first choice is a linear relation, such that $v(x_i) = x_i h + k$. The parameters h and k are assumed to be the same for every agent and to be measurable by the modeler. The external incentive or influence of one alternative (like news or advertisement campaign) is represented by h, whereas k is a constant utility. As agents compare their alternatives and choose the alternatives that maximize their utilities, a constant utility will have no impact on their decisions and can be disregarded, leading to a measurable private utility of

$$v(x_i) = x_i h. \tag{2.45}$$

Combining the different approximations, the decision maker's utility, given by Eq. (2.41), yields

$$\begin{aligned} U(x_i) &= x_i h + J x_i \bar{x}_i^e + \epsilon(x_i) \\ &= x_i(h + J\bar{x}_i^e) + \epsilon(x_i), \end{aligned} \tag{2.46}$$

which can be identified as a random utility model (Eq. (2.26)) where the first term in Eq. (2.46) represents the representative utility. Together with the assumption of IIA, such that the unobserved utility is distributed according to a Gumbel distribution, the probability of choosing x_i, computed via Eq. (2.40), is

$$p(x_i) = \frac{e^{\frac{1}{\gamma}x_i(h+J\bar{x}_i^e)}}{\sum_{x_i=\pm1} e^{\frac{1}{\gamma}x_i(h+J\bar{x}_i^e)}}. \tag{2.47}$$

Comparing Eq. (2.47) to Eq. (2.23), the probability of a spin of the Ising model, an equivalence between the binary choice model with social interactions and the Ising model becomes apparent. With the assumption of global interactions and rational expectations of decision makers ($\bar{x}_i^e = \sum_i x_i/N$), the behavior of the average choice is described by exactly the same framework as the mean-field Ising model of Eq. (2.25). Due to this equivalence, the results obtained in Chapter 5, which are framed for an audience of physicists, could equally well be targeted to the social science community by using the framework presented here.

2.4.5 Probit vs. Logit model

The probit model and the logit model are the two most prominent models to study discrete choice scenarios. As was documented above, both are random utility models (i.e. can be described by Eq. (2.26)) where the unobserved utility for the probit model is distributed according to a normal law, emerging from the central limit theorem, while the stochastic term in the utility of the logit model follows a Gumbel distribution, a consequence of Luce's choice axiom (IIA). The different distributions for the unobserbed utility is their only distinction. As the difference between the logistic distribution (resulting from the difference between two Gumbel distributed random variables) and the normal distribution is very small, with slighty more tail events in the former case, it is empirically very hard to distinguish between the models and there are no a priori theoretical reasons to prioritize one before the other in most cases.

Among these two models, the logit model enjoys a stronger popularity among researchers. The reason for this is found in the closed form solution of the resulting probabilities (for the probit function the error function has to be used), which also simplifies the parameter estimation via log-likelihood.

As both binary choice models in Chapter 4 and 5 use a stochastic utility term which is normally distributed, they are probit models. However as stated above, similar results would be found for the logit model.

Chapter 3

Economic Bubbles and Related Literature

Economic bubbles and crashes are frequently discussed, but still controversial topics in the financial world and literature. Bubbles, being defined as intermittent regimes –persistent on the time-scale of years– of increasing over-valuation of assets, and crashes, the rapid –on the scale of days to months– deflation of high prices, bringing the price closer to its fair value, are strange beasts. According to the mainstream financial models and theories of the second half of the 20th century, they are sheer impossible, ignoring, among others, the events occurring in the Octobers of 1929 and 1987, as well as the beginning of 2000, and the phenomenal increase of stock prices that preceded these events.

The basic assumption of these models is the efficient markets hypothesis (EMH), which is one of the pilars of modern economics and finance, and constitutes the Null Hypothesis for any newly proposed model or empirical analysis. The EMH states that "prices fully reflect all available information", an idea, which was independently developed by Paul A. Samuelson (1965) and Eugene F. Fama (1965a,b). It is a beautiful, elegant and bottom up approach, which is closely related to the wisdom of the crowd. In this framework, markets are interpreted as platforms, where all investors can submit their opinion on the real and fundamental value of the traded asset, resulting in a price, which is a combination of all the information pooled together. Unsophisticated traders, who are willing to exchange assets at a different price, would either not find a counter-party or would be exploited by more sophisticated investors and would die out in the long run, leaving free markets populated by only fully rational investors. As such, according to the EMH, a free market is self-cleaning, reflects all available information and prices only change as novel information about the asset's value is revealed.

Unfortunately, detailed studies of financial market crashes concluded that they could not find any revelation of information that could explain the price movements in the previously mentioned October crashes of 1929 and 1987, nor is there evidence for such news events concerning the price drop of internet related companies in 2000, or the devaluation of financial institutions at the end of 2007. The general explanation behind the dramatic price drops is that they are corrections of a previously grossly overpriced asset class, which experienced a buildup in the preceding years, i.e., a economic bubble. Moreover, bubbles are not an phenomenon exclusively reserved for financial markets, some famous historical examples are the Dutch Tulip Mania (1634-1637), the Mississippi Bubble (1719-1730), the South Sea Bubble (1720), the British Railway Mania (1840s). Details about these bubble are elaborated by Sornette (2003). Besides the fact that the EMH is unable to explain bubbles and crashes, there are other statistical regularities, such as "excess volatility"(Shiller, 1981), momentum, overreaction and reversal (Jegadeesh, 1990; Bondt and Thaler, 1985; Jegadeesh and Titman, 1993) which are inconsistent with this framework (West, 1988). Building on anecdotal evidence, popular books like Galbraith (1954) and Kindleberger (1978), acknowledge bubbles and crashes and see their origin in human irrationalities, leading to overoptimism, fads and manias.

Next, I will review the academic literature and present a non-exhaustive list of models and empirical studies, which aims at explaining the existence of bubbles and their underlying origins. The structure of the review was partly influenced by two other reviews on this subject, namely Brunnermeier (2008) and Kaizoji and Sornette (2010).

3.1 Rational bubbles

Under the hypothesis of symmetric information, it is impossible for bubbles to exist in a financial market where finite maturity assets are traded by agents with rational expectations (Tirole, 1982). However, rational bubbles can exist in models with infinite maturity assets. Blanchard and Watson (1982) show that rational agents may be invested in a bubble if its growth is fast enough to remunerate the risk of the bubble bursting. The setting for such bubble is however very restrictive, as the permuted inception of a rational bubble can only be at the first date of trading (Diba and Grossman, 1988b) and the price must not have an upper bound.

Empirical testing for such bubbles is a problematic task. Diba and Grossman (1988a) find no evidence that stock prices contain rational bubbles, while West (1987) can reject the null hypothesis of no bubbles. The inconclusiveness of empirical evidence is confirmed by Flood and Hodrick (1990) and Evans (1991), who acknowledge the difficulty to detect explosive patterns of bubbles due to poorly specified models and the linearity of economic methods not adapted to this issue. Another criticism of such rational bubble models is its reliance on the backward induction of investors to compute the value of the asset, a propriety, which is not found in experiments (McKelvey and Palfrey, 1992) where participants play the centipede game (Rosenthal, 1981), an iterative two-person game. A formal criticism is reported by Lux and Sornette (2002), who show that models of rational bubbles, as introduced by Blanchard and Watson (1982), exhibit returns whose tail-distribution is described by a power-law with an exponent smaller than 1, which is at odds with empirical evidence.

Dropping the symmetric information hypothesis allows for models with heterogeneity in agent's knowledge. In this setup, asset prices, for one, reveal the scarcity of the asset, and second, constitute the aggregated information of all active agents. Under asymmetric information, it might be that not every investor is aware of the bubble if assets are persistently overvalued. Another possibility is that every agent knows that prices exceed fundamental values but not everyone knows that all other agents are aware of it. Such models are based on the "greater fools" phenomenon, (Kindleberger, 1978; Kindleberger and Aliber, 2005) where investors hold overpriced assets in hopes of selling them to greater fools before the bubble bursts. As an example of the existence of a rational bubble under the assumption of asymmetric information and short sale constraints, the model of Allen et al. (1993) is referred to.

3.2 Heterogeneous Beliefs and Limited arbitrage

In order for a market to be efficient, every new piece of information has to be immediately be incorporated into the market. This is, however, only possible if investors can react to positive, as well as negative news. If short-selling possibilities are limited or non-existent, the hands of rational investors are bound, leaving them unable to fulfill their "duty" and correct the mis-prices created by unsophisticated market actors. In this scenario, bubbles can persist.

A description of the various limits of arbitrage is given by Shleifer and Vishny (1997). The principal tool for rational investors to act against an overpriced asset is to short-sell the asset. As this practice is more strongly regulated than for long positions, a majority of private investors are not able to to express their views during a bubble, leaving the price-correcting trading to institutional investors. However, as reported by Almazan et al. (2004), a large majority of mutual funds are also not permitted to sell short. Furthermore, 79% of US equity mutual funds do not use derivatives, according to Koski and Pontiff (1999), suggesting that the practice of synthetically creating short position is not used by such funds either.

Under the reasonable assumption of limited arbitrage and investors with heterogeneous beliefs, there exists a large number of models generating an inflation of prices (Lintner, 1969; Miller, 1977; Jarrow, 1980; Chen et al., 2002; Duffie et al., 2002; Hong et al., 2006). The basic mechanism behind these models is that investors have very diverse expectations of future cash flows of the traded assets, but because of the short sales restrictions, not every investor is able to incorporate his information into the price. With the short sales restrictions, the conditions for a free market are no longer fulfilled, corrupting the wisdom of the crowd effect of aggregating information by sampling the opinion of a large number of investors. It is the exclusively sampling of the optimistic population that pushes the price above its fundamental value.

This mechanism is considered among the most convincing to explain the irrational exuberance behind the rise of the Internet related stock prices. The regulation of holding newly IPOed shares, which the majority of Internet company shares were, for a minimum period of 6 months, keeping the owners from closing their positions and stoping the increasing prices of these stocks. For a large number of companies this lockup expired in the spring and latter half of 2000, which can be interpreted as a removal of short sales restrictions and which, according to Cochrane (2002) and Ofek and Richardson (2003) resulted in the beginning of the correction of the overpriced industry.

A related mechanism for the persistence of overvalued assets is risk aversion. Rational, but risk averse, investors could restrain from heavily arbitraging inflated prices as shorting is riskier and costlier then long positions. It may be that it takes a long time till the price pressure is big enough to bring the price back to fundamental levels, or it may be that, by luck, the irrational expectations become right and then the arbitrageurs would take major losses.

3.3 Consciously Ridding the Bubble

Another type of model to investigate the persistence of bubbles and their termination was introduced by Abreu and Brunnermeier (2003). They assume rational investors, who, sequentially realize that the price is in a bubble, which will eventually collapse. They know that one single arbitrageur will not be able to bring prices down to fundamental values, but they do not know how many other investors are aware of the bubble. Based on his knowledge, the investor has two ways of reacting to this situation. He can either immediately bet on falling prices and not profit from the subsequent run-up, or decide to temporarily ride the bubble, while taking the risk of not getting out in time before the bubble bursts. As the bubble will only collapse if the number of arbitrageurs exceeds a threshold, investors face a synchronization problem and they try to forecast when other investors will attack the bubble. In this setup, a bubble can persist even when a large number of market actors are aware of it and a small and insignificant news event can trigger a crash, as it helps traders to synchronously engage their arbitrage strategies.

The strategy of riding a bubble, instead of acting against it, as is proposed by this model, was confirmed by the empirical studies of Brunnermeier and Nagel (2004) and Griffin et al. (2011). They found that, during the Internet-bubble from 1997 to March 2000, hedge funds were not only heavily invested in technology stock, but that they were also aware of their over-valuation as they reduced their potions as the bubble's end was in sight. Such evidence questions the stabilizing effect of rational investors on prices, and shows that sophisticated traders can profit from deviating, as well as returning prices to fundamental values.

Similar findings are reported by a survey from Shiller (1987) immediately after the 1987 crash. It is reported that investors were aware that the market was overvalued, which however did not stop them from being exposed to it. Many investors showed a strong overconfidence in their capacities, as they were convinced of being able to predict the market, i.e. getting out in time.

3.4 Positive feedback trading

In contrast to previously mentioned models, where bubbles emerge because investors had very high hopes for future cash flows or where they know prices were inflated but predicted further growth, this section is concerned with a class of models, where positive feedback traders, momentum traders, or "noise traders" as they were introduced by Kyle (1985) and Black (1986), are responsible for the persistent deviation of the fundamental value. These agents are not aware of a bubble, nor do they estimate future cash flows, instead they engage into trend-following investments, decoupled from any fundamental value[1]. This kind

[1] Generally, noise traders are traders, which base their trading decision on invalid information, i.e. which is not helpful to predict future cash flows, such as past price movements and technical analysis.

of strategy does not necessary have to be irrational. It may result from stop-loss orders, liquidation of positions to meet margin calls or investors' risk aversion rapidly declining with increasing wealth (Black, 1988; Leland and Rubinstein, 1988).

3.4.1 Experimental evidence and existence of momentum

On the other hand, evidence for the affinity of humans to engage in trend-following behavior was found in an experimental study by Andreassen and Kraus (1990). They asked their participants, who had some training in economics, to bet on future price movements of authentic stock price patterns. When, over some period, stock prices did not significantly change, subjects predicted mean-reverting patterns, where a positive return is followed by a negative one, and vice versa, such that the price remained at the same level. If, however, there was an apparent trend in the price time series, subjects intuitively changed their strategy and predicted a continuation of the trend. As these results were virtually universal among all participants, trend-following behavior appears to be a widespread phenomenon. Similar results were found by Offerman and Sonnemans (2004), who report that subjects mistook randomly appearing trends as a sign of positive autocorrelation.

Besides these technical and behavioral reasons for the existence of positive feedback trading, there is also statistical evidence that stocks exhibit momentum behavior at intermediate horizons. For companies traded on the New York Stock Exchange, Jegadeesh and Titman (1993) find that the strategy of shorting past losers and buying past winner stocks gives economically significant profits for six to twelve moths after the portfolio creation. Winners and losers are defined as the top and bottom decile of the last six-month returns. Similar results are found for stocks of European stock exchanges (Rouwenhorst, 1998) and emerging stock markets (Rouwenhorst, 1999). Beside momentum at intermediate horizons, stock price dynamics show a long-term reversal at time-scales of 3-5 years (Bondt and Thaler, 1985). Similar results are found by Cutler et al. (1990), who confirm positive return autocorrelation for returns up to one year for markets for stocks, bonds, foreign exchange, and various real assets, and negative return autocorrelation for returns on stocks, bonds, and foreign exchange on a time-scale of ~ 2 years. The subject of momentum in returns followed by price reversals will be covered in more detail in the introduction of the paper presented in Section 6.2.

3.4.2 Models with momentum traders

In the following, additional models are discussed, which are able to generate deviations from fundamentals, due to feedback trading, and subsequent price reversal. The aim of these model is, among others, to explain the observed positive mid-term and negative long-term return autocorrelation. Although bubbles and crashes, which constitute the extreme case of this stylized fact, are not explicitly mentioned in these studies, I still believe that they give some valuable insight into the underlying of these phenomena.

Investigating the effect of positive feedback traders, de Long et al. (1990) proposed a model studying the interaction between positive feedback traders and rational investors, resulting in a price inflation and a subsequent reversal. In this model, rational investors, who obtain information to estimate the fundamental value of the asset, are aware of the presence of positive feedback traders and their investment strategy. As such, they drive the price up, over its fundamental value, knowing that the noise traders, who are solely attracted by the price increase and are unaware of its real value, will buy the asset anyways. At the last period of this 3-period model, the asset's real value is revealed, bringing the price down to it fundamental value. This model shows that rational speculators, who are usually thought to stabilize prices, can have a destabilizing effect on prices when interacting with feedback traders. Another model which assumes feedback traders is proposed by Cutler et al. (1990). In that set-up, noise traders are able to make some profit, which is in contrast to de Long et al. (1990), where feedback traders are badly exploited.

Similar to the previous model, in that the deviation from the fundamental value is created by feedback trading, but different in its underlying origin, is the under- and overreaction model of Daniel et al. (1998). In their representative agent model, they investigate the price dynamics resulting from the trading of quasi-rational individuals, whose decision process is explicitly modeled and based on well documented psychological biases. These biases are overconfidence about the precision of private information, and biased self-attribution, the fact that people react asymmetrically to new information. They tend to credit themselves for past success and blame external factors for failures, or as put by Langer and Roth (1975): "Heads I win, tails its chance". Given the assumption of overconfidence, investors will overreact to their private information. If their private information is confirmed by public information, their confidence is further increased leading to an even stronger overreaction and generating momentum. Once public information starts contradicting their initial private information and signalling the deviation from the fundamental value, the confidence in their private information is only slowly decreased and they are underreacting to the public news. This underreaction results in a slowly reverting of the price back towards the fundamentals.[2]

Hong and Stein (1999) proposed a market model where two types of agents interact, "newswatchers" and "momentum traders". Both types are bounded rational, as every type relies exclusively on his type of information, which is private information and past price changes, respectively. The model is based on an initial underreaction of the newswatchers as information gradually diffuses across the population of newswatchers and is slowly integrated into the price. This slow integration creates a drift, which attracts the momentum traders. Their trading activity eventually leads to an overreaction of the price to the news and a deviation from the fundamentals. The reversal corresponds to the unwinding of momentum traders positions, which happens endogenously compared to de Long et al. (1990), where prices are just forced back to fundamentals on a terminal date.

A modern incarnation of market model with rational and noise traders is prosed by Kaizoji et al. (2011), where noise traders not only indulge in short term momentum trading, but are also subjected to social imitation, a mechanism that will be introduced in Section 3.5. Intermittened bubbles are witnessed for large enough populations of noise traders, who can push the price from from its

[2]Another model, which is also based on psychological evidence (conservatism bias) is studied by Barberis et al. (1998), where investors over- and under-react to different types of news and oscillate between two states, mean-reverting and trending.

fundamental value. During the bubble regimes, the momentum strategies are found to be protable.

3.4.3 Empirical evidence of positive feedback trading

In contrast to Daniel et al. (1998), where the momentum strategy is a result of two behavioral biases, de Long et al. (1990) and Hong and Stein (1999) assume, a priori, the presence of momentum traders, whose behavior is crucial for their results. Although not completely unchallenged (Lakonishok et al. (1992) only finds evidence for herding in small stocks), the majority of investor behavior studies finds evidence for such a trading strategy, confirming the validity of the momentum traders as an ingredient in financial market models.

Analyzing quarterly buying and selling imbalances of mutual funds, Klemkosky (1977) finds that after two months of abnormally positive stock returns, large buying imbalances followed and one month or more of abnormally negative returns engendered a large selling imbalances, i.e. a text-book example of momentum trading. Grinblatt et al. (1995) report that of their sample of 155 mutual funds over the period of 1975-1984, 77% of the mutual funds were momentum investors and that on average, funds that invested on momentum, realized significantly better performance than other funds cf. Wermers (1999). For individual security and NYSE portfolio, Sias and Starks (1997) find that daily return autocorrelations are an increasing function of the level of institutional ownership, evidence that institutional investors and funds engage into momentum trading. Also Nofsinger and Sias (1999) document intra-period momentum trading, and that this strategy is successful as “stocks institutional investors purchase subsequently outperform those they sell” in the following two years. Contrasting these intra-period momentum trading results, is the study of Sias et al. (2001), which states that the positive correlation between returns and changes in institutional holdings originates primarily from the price pressure exercised by the institutional trading.

Badrinath and Wahal (2002) observe significant differences in trading practices among different types of institutions, but report that generally institutions act as momentum traders when they enter stocks, i.e. exactly what is needed to fuel a bubble. Dennis and Strickland (2002) document that, on days when the overall market experiences large movements, the magnitude of a company’s stock return, with the same sign as the market, is positively related to institutional investors’ ownership in that company. Also the turnover of these stocks is positively related to institutional ownership, indicating that institutional investors sell (buy) more then private investors when the market drops (rises).

A nice clearcut result is reported by Griffin et al. (2003), who state that, “Based on the previous day’s stock return, the top performing decile of securities is 23.9% more likely to be bought in net by institutions (...) than those in the bottom performance decile.” Using quarterly data, Cai and Zheng (2004) confirm that stock returns Granger-cause institutional trading, and that this is especially true for purchases. The success of this trading strategy is also confirmed as “the stocks with heavy institutional buying (selling), experience positive (negative) excess returns over the previous 12 months”. Also in the Taiwan market, institutional investors are found to engage in momentum trading (Chen et al., 2008). Further evidence for momentum trading is reported by Sias (2004), where it is stated that institutional demand is only slightly related to past returns, but strongly related to lagged institutional demand. This finding suggests that there might be a different or additional mechanism behind the creation of mid-term return momentum and bubbles, namely that investors imitate each other. This mechanism will be discussed in the next section.

3.5 Herding

Herding, in the context of financial economics, refers to economic agents' behavior of imitating each others, i.e., the opinion and actions of an agent are not entirely motivated by her own private information and conviction, but mainly by the opinion and actions of agents in their surrounding. This mechanism is thought to be an important mechanism behind the growth and collapse of a bubble. As herding investors disregard some of their information, this information will not be incorporated into the price, making it possible for prices to deviate from fundamentals. Besides the price not reflecting all available information, unsubstantiated rumors and false news may impact the price and be amplified via the positive feedback mechanism of herding[3]. This behavior can lead to self-fulfilling prophecies, as supply and demand impact prices, confirming, on the short run, investors' opinions. On the long term however, prices will return to their fundamental value, when investors' irrational exuberance has to face economic reality.

The original meaning of herding, moving together, is found in nature as the default behavior of a large group of animals, ranging from sardines to grasshopper, and from buffaloes to lemmings[4]. But also the copying of specific actions is widely documented for a variety of tasks, such as foraging, diet choices, means of avoiding predators and selection of mates (Gibson and Höglund, 1992; Dugatkin, 1992; Dugatkin and Guy, 1992; Giraldeau, 1997; Pennisi, 2010). It is even suggested that, in the early evolution of humans, the ability to imitated innovative complex behavior was selected for (Blackmore, 2000), which may have promoted the development of our comparably large brain (Dunbar, 1998).

The tendency of humans to blend in and imitate their social environment is also well documented in psychology. Already early experiments showed that, within a group, individuals often abandon their own private signal to adopt the opinion shared by the majority. For instance Sherif (1937) asked participants to report the movements of a point of light in darkened room without any point of reference. In fact, the point was not moving at all, but when participants where in groups, they quickly, without any discussion, came to a consensus (differing across groups) on the amount of movement. Afterwards, interviews revealed that the participants were unaware of the group's influence on the final decision. A similar experiment was conducted by Asch (1952), where subjects had to compare the lengths of line segments. When they were alone most participants gave correct answers, but when in a group, where all other members were asked

[3]Due to its intrinsic sequential character, in order to copy an action, it has to be performed first, herding creates a positive feedback mechanism.

[4]Although the stories of Lemmings committing mass suicide are just urban legends, whose origin goes back to the 1958 Disney film "White Wilderness", where Lemmings are shown to jump in large numbers off a cliff into certain death. In this movie, the Lemmings "jumped" off the cliff, not because they were blindly following their fellow rodents, but because they were launched off the cliff using a turntable, as was reported by a documentary of the Canadian Broadcasting Corporation in 1983.

to give the same wrong answer, subjects frequently agreed with this wrong answer. Generally, subject were aware of the correct answer, as they reported in later interviews, but were afraid to contradict the group. The fact that social influence currupts the "wisdom of the crowd" effect was also shown in a recent experiment by Lorenz et al. (2011). Participants were asked factual questions, whose answers they could reconsider after having seen partial or full information about the responses of other subjects. Their initial "wise" answer gets distorted by this information, letting them converge towards wrong answers.

The reason for this change of opinion are manifold. In the situation of incomplete knowledge, it may be rational for an agent to change her opinion, as the group's opinion is potentially based on more information, increasing its likelihood of being correct. Another reason may be the human's innate wish for conformity and the urge of belonging to a group, decreasing the chances for discrimination and the distribution of a potential penalty over the whole group. Recent results from empirical psychology point to an additional reason. Zaki et al. (2011) find that observing the choices of others, changes the internal preferences of the observer, such that the change of opinion is not a conscience rational decision, but a subconscious change of preferences. Similar results are reported by Edelson et al. (2011).

The fact that herding in the context of financial markets is real and has important consequences was also acknowledged by Jean-Claude Trichet, President of the European Central Bank, who said about the incentives and behavior of fund managers that, "Some operators have come to the conclusion that it is better to be wrong along with everybody else, rather than take the risk of being right, or wrong, alone ... By its nature, trend following amplifies the imbalance that may at some point affect a market, potentially leading to vicious circles of price adjustments and liquidation of positions" (Trichet, 2001).

3.5.1 Theoretical models of rational herding

Contrarily to the intuition, of labeling herding as an irrational and counterproductive behavior, the following studies show that imitating one's peers may be the optimal choice, both in terms of profitability and sustainability (i.e., keeping your job and reputation). Here, these various mechanisms behind herding behavior will be elaborated and put in the context of financial markets.

Informational herding

One of the best-known models of information-based herding is the *informational cascades* model (Bikhchandani et al. (1992); Banerjee (1992); Welch (1992)). These occur in a setup, where every agent has only limited information and has to choose one of several alternatives. Agents choose sequentially and see the choices of the agents prior to them. As every agent has different information, it becomes rational to choose the alternative, which attracted the majority, even if this choice stands in contradiction with the agent's own information. "An informational cascade occurs when it is optimal for an individual, having observed the actions of those ahead of him, to follow the behavior of the preceding individual without regard to his own information", according to Bikhchandani et al. (1992).

This mechanism is best illustrated by the example of a hungry tourist, looking for a good restaurant. While strolling around in a city, unknown to him, he arrives at a place with many restaurants. One of these is packed with people, who look like locals to him, while the other restaurants are sparsely filled. Although his tourist-guide states that one of the little occupied restaurants is the best place in town, he will still go to the packed place, as, so he thinks, the locals surely know where to find the best food. Even though the tourist's choice is completely based on the choices of his surrounding, while ignoring all his private information, it is still a rational choice, which could lead to the optimal outcome. It might be that the tourist-guide is out-dated or that the reviewer had a bad day when he was rating the restaurant. If, however, every costumer in the restaurant took his decision based on its popularity at the time, the final population is entirely controlled by the first few guests and the occupancy of the restaurant is not necessarily related to the quality of its food.

A similar argument can be constructed for investors selecting a stock to invest in, where either the opinion of the investor's entourage or the latest changes of the stock price represents its popularity. The latter scenario is equivalent to momentum trading, which was discussed in the previous section, whereas the former constitutes the basis of a class of financial market models that will be introduced in Section 3.5.3.

There is, however, a major difference between choosing a restaurant and picking a stock. The tourist knows, as soon as he has finished his meal, whether his choice was the right one. For an investor on the other hand, it is not so clear-cut as the value of a stock equals the discounted future cash-flows, which is never known with certainty. What he however sees, is the dynamics of the price since his transaction. If the direction of the price change is in agreement with his prior opinion, he feels confirmed that his opinion was right. This mechanism, which will be investigated in the main section of the next chapter (Section 4.2), creates an herding-induced positive-feedback process and leads to self-fulfilling prophecies.

An important consequence of the sequential character of the setup is that once the informational cascade has occurred, any new choice is uninformative for later observers, as his private information is ignored, leading to an *information blockage*. Once a blockage occurs, any choice of newly arriving agents is on no benefit to others.

Reputational herding

In their seminal paper, Scharfstein and Stein (1990) propose a model of reputation based herding of fund managers. Outside observers infer the managers' ability from their investment choices and the resulting payoffs. They assume that competent managers should have similar choices, while for incompetent ones, independent noise should be observed. As the managers themselves are unsure about their own ability, it is rational for them to copy each other, even if it contradicts their private information. Collective bad payoffs will be attributed to bad economy or bad circumstances, while taking losses and the competition making profits out-weights the opposite situation if loss of reputation is to be avoided. Herding of risk averse managers is found to be rational in many models if information is costly (Gumbel, 2005), and if managers are compared relative to each other (Zwiebel, 1995; Maug and Naik, 2011).

An example of such a behavior can be found in the strong adpotion of CDOs and CDSs into the portfolios of many bank and funds during the rising retail pricees in the US from 2003-2006. During that period, these derivatives generated extraordinary profits, such that private investors went to funds, which offered these great rates. Even managers, who were aware of the potential dangers of a large exposure to such assets, were forced to invest in then, in order to compete with other funds who had less objections. As such, comparing fund performances with each other does not only lead to herding among managers, but also to trading strategies which are only myopically optimal and can result in a destabilized market, a subect which will be treated in more detail in Section 4.2. Besides, Prendergast (1993) shows that subordinate managers are incentivised to make recommendations consistent with the prior beliefs of their superiors.

The behavior of analysts is considered by Trueman (1994). He shows that, even without justifiable information, they are likely to make their current earning forecasts close to prior earnings expectations and to those previously announced by other analysts, as investors interpret this as a good analysis. The less skillful the analyst is, the stronger his bias. Similar behavior is predicted by the model of Graham (1999), which is validated by investigating analyst's reactions to a recommendation newsletter using thirteen years of data. Also for firm managers, it may be better to make investment decisions consistent with the market's "prejudices" (Brandenburger and Polak, 1996), even when they have superior information telling them otherwise. This behavior arises from the willingness to please investors and share-price maximization, which however, favors the short-term.

Investigative herding

Investigative herding refers to the situation where information is investigated only if the analyst believes that others also consider this information to be important, similar to the Beauty Contest of Keynes (1936).

In the seminal paper on investigative herding, Brennan (1990) proposed a model, where acquiring costly information only pays off if investors coordinate and move the price in the direction of the new information, giving an advantage to whose who respond fast. Otherwise, the costly information might end up useless and the investor will not be able to sell the asset at a profit. A similar model is investigated by Froot et al. (1992), where they show that for short-term investors, it can be profitable to learn what other informed investors know and to trade on this information, even if it is only spurious information. Hirshleifer et al. (1994) find a tendency to herd in a situation where some traders receive information before others. The choice of focusing on information regarding either long-term or short-term can also be a result of herding, according to Holden and Subrahmanyam (1996)

3.5.2 Empirical evidence of herding

Having presented a large variety of models containing agents that indulge in some sort of herding, here, a selection of studies is reviewed, which reports strong evidence for such a behavior among economic agents.

General herding evidence

As the empirical studies on momentum trading, reported in Section 3.4, investigate the changes in the portfolios of institutional investors given prior market movements, the data at the disposal of these researchers also enables them to search for herding evidence. Evidence of herding can take several forms, such as portfolio changes of one fund leading to similar changes in an other fund, or that funds perform the same changes at the same time, or that the dispersion between the portfolio of different funds is too small as to be explained by independent enterprises.

For instance Grinblatt et al. (1995), which find strong evidence for momentum investors, also report that funds tended to buy and sell the same stocks at the same time, even though their evidence is relatively weak. Wermers (1998) find stronger evidence of funds simultaneously buying the same stock in the 1975-84 period, especially stocks with high past returns. The fund's herding behavior is found to be profitable, as the stocks that were bought as a herd, significantly outperformed during the following quarters those, which were sold as a herd. Also sequential herding is documented, as he finds that, for popular stocks, some groups of funds imitate the future portfolio choices of other funds. In a different study on mutual funds from 1975 through 1994, Wermers (1999) reports little herding in average stock, but much higher levels of herding in small and growth stocks. This is in line with the predictions that herding is more likely to occur if fundamental value estimation is harder and when there is more uncertainty. Confirming his previous result he finds that, while herding, the purchases significantly outperform the sold assets.

Also other studies document copying behavior of funds (Sias and Starks, 1997) and acting in synchronicity (Nofsinger and Sias, 1999; Dennis and Strickland, 2002). Very clear evidence that institutions indulge in informational herding, i.e., inferring information from each other trades, is found by Sias (2004) for the period from 1983 to 1997. They state that, "institutional demand is more strongly related to lag institutional demand than lag returns". Similar results are found for the Taiwanease market, where Chen et al. (2008) report that institutional investors follow each other into and out of the same securities. Such findings are confirmed by Demirer et al. (2010), who in addition find that herding is stronger in bear markets, i.e., panic fueld herding.

Detailed herding evidence

To show that the previous studies are not a result of spurious herding, where the different funds had similar portfolios because they reacted to the same information and not because of any interaction, the following studies investigate the behavior of analysts and investors in more detail, such that spurious herding can be excluded.

By studying the relative tightness of analysts earning forecast distribution in the 1985-1987 period, Olsen (1996) finds that the majority (52 to 72 %) indulge in herding, with herding behavior even increasing with the unpredictablity of future earnings. Analysts' earnings prediction are shown to be biased towards the consensus forecast, seen in the decreased dispersion of their opinions. A tendency for optimistic predictions is also observed, as analysts are more likely to herd towards the consensus and disregard their private information, if their private information is pessimistic. Also De Bondt and Forbes (1999) find "excessive agreement" among analyst predictions. Analyzing U.K. companies between 1986 and 1997, their evidence supports overoptimism, overreaction and "a surprising degree of consensus relative to the predictability of corporate earnings". Chevalier and Ellison (1999) report herding among younger managers, as they avoid unsystematic risk and are invested into popular sectors. Evidence that herding is negatively related to experience is also found for analysts by Hong et al. (2000); Clement and Tse (2005). Minimizing the risk of a bad reputation at the start of a career, is thought to be the main reason for this herding. Herding also increases with the difficulty of the task, as is reported by Kim and Pantzalis (2003) for analysts-data of the 1980-98 period.

Sequential herding, i.e. cascading, of analysts is reported by Welch (2000), who shows that security analysts' recommendations are significantly influenced by previous recommendations. The influence is stronger the more recent the past recommendation is. It is also shown that their recommendations correlate with the prevailing consensus forecast and that the strength of the consensus' influence is independent of its accuracy. This leads to the conclusion that the herding towards the consensus is not based on fundamental information but on short-lived information, i.e., analysts herd after noise. In contrast to previous studies, Welch (2000) finds that consensus-herding is significantly stronger in bull markets and towards an optimistic consensus, consistent with the overall optimistic bias of analysts reported in other studies. Stronger herding during "good times" leads to a poorer information aggregation such that these times represent a fertile ground to grow bubbles, with a valuation based on thin air.

Further evidence for herding among analysts is provided by Guedj and Bouchaud (2005), who document over-optimism and very small variance of forecasts compared with forecasting errors. These effects were particularly strong during the early nineties and the Internet bubble.

Let us now turn our attention to investors and fund managers. In one of the earliest of such studies, Shiller and Pound (1989) asked institutional and individual investors, "How do investors develop interest in and receive important information, leading to decisions about investments?". Many investors reported that personal contacts had brought their attention to stocks they recently bought. Generally they found that direct interpersonal communications had a strong impact on investors' trading decisions. As for analysts, herding, as well as trend-following behavior, is more common in younger and inexperienced fund managers, as was found by Greenwood and Nagel (2009), who investigated managers behavior during the dot-com bubble.

Having provided evidence of herding of analysts above, Brown et al. (2009) show that mutual funds are sensible to analysts reports and herd into and out of the stocks that are up- and down-graded in analysis consensus predictions. Besides inducing changes in funds holdings, analysts' revisions also impact the market, as stocks traded by herds of mutual funds exhibit a price impact in the same quarter, together with a price reversal in the following quarter, showing the potential price destabilizing effect of analyst information when overused by fund managers.

Another great study on herding and word-of-mouth propagation of information, with similar results as the survey by Shiller and Pound (1989), is the study of Hong et al. (2005). They find that for fund managers living in the same city, changes in their portfolio are correlated, i.e. the probability of buying (selling) a certain stock significantly increases if an other manager in that city has bought (sold) the same stock. They can rule out local preferences, as the correlated portfolio changes also appear for companies far away from the managers' home.

3.5.3 Herding-based models of financial markets

In the context of the models presented in this section, herding refers to social interaction, the situation where the opinion and actions of individuals are positively influenced by those of their reference group. Social interactions can explain self-reinforcing behavior and may lead to multiple equilibria in the absence of external coordination. Besides, in certain situations, small exogenous influences on systems composed of interacting individuals may cause large changes at the aggregate level. With the exception of one model (Shiller et al., 1984), the models presented here are based on the Ising model (cf. Chap. 2), which arises from the binary choice setup (buy/sell), the peer-interaction among investors and the adjustable interaction strength or variance of the random term of their utility. For additional information on such and related models, Hommes (2006, 2008) reviews the literature of interacting agents in economics and finance.

One of the first models, incorporating social interaction, was introduced by Schelling (1971) and aimed at explaining the racial segregation in residential areas. He showed that even small racial preferences can lead to pronounced residual segregation. Social interaction was also found to be of crucial importance in the adoption of norms and habits (Schelling, 1973). For the majority of academics in the field of financial economics, it was Shiller et al. (1984) who first proposed a model of stock price dynamics which incorporates social interaction, leading decisions based on "animal spirits", as suggested by Keynes (1936), like "fads", where economic decisions are not correlated with economic reality. Among others, he affirms the importance of social interaction by pointing to its well documented effects in psychology and the impact of such behavior on stock prices, leading to excess volatility and extend deviations from the fundamentals. His affirmations were later confirmed by West (1988).

However, already ten years before Shiller et al. (1984), but not very known in the finance community, Föllmer (1974) analyzed the outcome of "Random economies with many interacting agents", although with less emphasis on the psychological origin or evidence of the contagious nature of opinions. His model of an economy, with agents governed by random preferences whose probability depends on the agent's environment, is inspired by contemporary research in physics and probability on interacting particle systems, with these final results being based on the Ising model. He shows that even short range interactions may propagate through the whole economy and give rise to price instabilities.

In a similar vein, Blume (1993) and Brock (1993) develop models with interacting agents facing discrete choices. Also the model by Kirman (1993), which was inspired by the puzzling behavior of foraging ants and aims at explaining epidemics and herding behavior in financial markets, follows the same scheme. If agents (or ants) are given the choice between two opinions (optimistic versus pessimistic or two different food sources), while being influenced by a randomly selected fellow agent (or fellow ant), a persistent and asymmetric distribution of opinions emerges, with the maximum suddenly switching between the two alternatives.

Other modes, directly building on the Ising model in the ferromagnetic state and investigating financial bubbles and crashes are proposed by Lux (1995); Kaizoji (2000); Kaizoji et al. (2002) among others. These are either static models (Kaizoji, 2000), relating to the multiple equilibria and hysteresis characteristic under external forcing of the Ising model, or dynamic models, including some kind of repulsion from the fixed points, such that the system can overcome a predominant opinion in one direction and switch between attractive fixed point (Lux, 1995; Kaizoji et al., 2002). These systems usually show periodic dynamics, with the majority switching between overly optimism and overly pessimistic.

An application of the Ising model to empirical data of opinion shifts is presented by Michard and Bouchaud (2005). Although not for financial data, they show how collective opinions can abruptly change (adoption of the birth control pill, cell phones, applause in concerts), when exposed to a slowly changing and weak signal, similar to the large response of the magnetization of the Ising model, relative to a weak change in the external magnetic field.

Similar to the previously presented models, in that it is inspired by the Ising model, Chapter 5 contains the study of a model which can be interpreted, among others, as a financial market model or opinion dynamics model. However, compared to all previous studies which either do not consider any external influence or a slowly varying driving force, the influence of a rapidly varying external signal will be considered, representing the constant flow of news, which economic agents are subjected to. In that setup a new phenomenon, which we call *noise-induced volatility*, will be identified, which is proposed as an explanation for the excess volatility found on financial markets (Shiller, 1981), together with the observed weak predictive power of news onto price dynamics (Parr, 1985; Cutler et al., 1989; Joulin et al., 2008).

3.6 Experimental evidence

Besides undoubtable evidence from general experiments of humans adopting the opinion of, and being influenced by, their surrounding, either consciously (Asch, 1952; Lorenz et al., 2011) or unconsciously (Sherif, 1937; Zaki et al., 2011; Edelson et al., 2011) as presented in the begining of Section 3.5, there exists also experimental studies specifically targeting the emergence of bubbles in the price of a good.

In their experimental exchange, Smith et al. (1988) study the price dynamics of a risky asset, traded via a double-auction setting by their participants. At every time-step of its finite life-time, the asset pays a random dividend, whose distribution and expected value are well known to the participants, such that the asset's fundamental value is monotonously decreasing in time. Despite the homogenous knowledge among the participants of its payoff-structure, there is intensive trading and prices initially rise, and surmount the asset's fair value. For the majority of their experimental price evolutions, they witness bubbly dynamics, which can be separated into three phases: (i) an initial rally, exceeding the fundamental price, (ii) a cooling off of the rally and (iii) a crash towards the end of the assets lifetime. Subsequent studies show that these bubbles persist even if short-sales are possible, if trading fees are introduced and if participants have an economic background. In other words, even if subjects are given the opportunity and incentives to fight the bubble, they still persist.

It was thought that the "greater fool" hypothesis and the lack of common knowledge are the reasons behind the intensive trading and the emergence of the bubble. In order to test this hypothesis, Lei et al. (2001) performed a similar experiment, where participants did not have the possibility of reselling a purchased share, rendering speculative trading impossible. Even in this setup bubbles and crashes are observed and are attributed to irrational elements in participants behavior as lack of common knowledge of pure speculation can be ruled out.

3.7 Concluding remarks

In the previous sections, a wealth of literature was presented, which either provide empirical evidence or theoretical models for the existence and underlying origins of bubbles, as well as their weakened manifestation, price momentum. To underline the impact of the preveiling consensus on an asset's value, in contrast to its fundamental value, I point to an interview with Eugene F. Fama (2007), conducted on November 2. 2007, i.e., at the top of the real eatate and credit bubble. In this intervew, Fama states that "The word *bubble* drives me nuts", and that he does not acknowledge the internet bubble, nor that the crashes of 1927 and 1987 contradict the EMH. When asked whether the housing markets are efficient, he replied that he does not know, but says that, when people are buying a house, "they look around very carefully and they compare prices", implying that his guess would be an affirmation of their efficiency. One year later, the S&P 500 had lost 36% of its value.

This statement of Prof. Fama shows the importance of the results documented in experiments such as Zaki et al. (2011), Edelson et al. (2011) or Lorenz et al. (2011), where social interaction makes people honestly believe that the majority is right, corrupting the wisdom of the crowd effect, and as such, also the efficiency of a free market.

Chapter 4

A model of myopic adapting agents

4.1 Introductory comments

My contribution to the literature of bubble formation and origin of financial crashes is presented in Section 4.2. It contains the full paper, as it is published in the Journal of Economic Behavior & Organization (Harras and Sornette, 2011). The paper is self-contained with introduction and conclusion, rendering an additional introduction superfluous. In Section 4.3, additional comparisons with previously discussed models are performed and further evidence for the validity of our assumptions and conclusions are provided.

4.2 The paper

Journal of Economic Behavior & Organization 80 (2011) 137–152

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

How to grow a bubble: A model of myopic adapting agents

Georges Harras*, Didier Sornette

Department of Management, Technology and Economics, ETH Zurich, CH-8001 Zurich, Switzerland

ARTICLE INFO

Article history:
Received 11 November 2010
Accepted 2 March 2011
Available online 21 March 2011

JEL classification:
G01
G14
D84

Keywords:
Stock market
Crash
Bubble
Herding
Adaptation
Agent-based model

ABSTRACT

We present a simple agent-based model to study the development of a bubble and the consequential crash and investigate how their proximate triggering factor might relate to their fundamental mechanism, and vice versa. Our agents invest according to their opinion on future price movements, which is based on three sources of information, (i) public information, i.e. news, (ii) information from their "friendship" network and (iii) private information. Our bounded rational agents continuously adapt their trading strategy to the current market regime by weighting each of these sources of information in their trading decision according to its recent predicting performance. We find that bubbles originate from a random lucky streak of positive news, which, due to a feedback mechanism of these news on the agents' strategies develop into a transient collective herding regime. After this self-amplified exuberance, the price has reached an unsustainable high value, being corrected by a crash, which brings the price even below its fundamental value. These ingredients provide a simple mechanism for the excess volatility documented in financial markets. Paradoxically, it is the attempt for investors to adapt to the current market regime which leads to a dramatic amplification of the price volatility. A positive feedback loop is created by the two dominating mechanisms (adaptation and imitation) which, by reinforcing each other, result in bubbles and crashes. The model offers a simple reconciliation of the two opposite (herding versus fundamental) proposals for the origin of crashes within a single framework and justifies the existence of two populations in the distribution of returns, exemplifying the concept that crashes are qualitatively different from the rest of the price moves.

1. Introduction

Bubbles and crashes in financial markets are events that are fascinating to academics and practitioners alike. According to the consecrated academic view that markets are efficient, bubbles, being temporally persistent, self-reenforcement deviations of the price from the fundamental value, are impossible. And crashes should only result from the revelation of a dramatic piece of information. Yet in reality, there is a large consensus both from professionals (Dudley, 2010; Trichet, 2010) and academia (Shiller, 2000; Abreu and Brunnermeier, 2003) that bubbles do exist, and even the most thorough post-mortem analyses are typically inconclusive as to what piece of information might have triggered the observed crash (Barro et al., 1989).

It is often observed that crashes occur soon after a long run-up of prices, referred to as a bubble. A crash is thus often the burst of the bubble. There is a vast amount of literature aiming at characterizing the underlying origin(s) and mechanism(s) of financial bubbles (Abreu and Brunnermeier, 2003; Kaufman, 2001; Sheffrin, 2005; Shiller, 2000; Sornette, 2003a) but

* Corresponding author.
E-mail address: gharras@ethz.ch (G. Harras).

0167-2681/$ – see front matter
doi:10.1016/j.jebo.2011.03.003

there is still no consensus in the academic community on what is really a bubble and what are its characteristic properties. Bubbles do not seem to be fully explained by bounded rationality (Levine and Zajac, 2007), speculation (Lei et al., 2001) or the uncertainty in the market (Smith et al., 1988). Finally, there is no really satisfactory theory of bubbles, which both encompasses its different possible mechanisms and adheres to reasonable economic principles (no arbitrage, equilibrium, bounded rationality, etc.).

Most approaches to explain crashes search for possible mechanism or effects that operate at very short time scales (hours, days, or weeks at most). Other mechanisms concentrate on learning an exogenously given crash rate (Sandroni, 1998). Here, we build on the radically different hypotheses summarized in (Sornette, 2003a) that the underlying cause of the crash should be found in the preceding months and years, in the progressively increasing build-up of a characteristic that we refer to as 'market cooperation', which expresses the growth of the correlation between investors' decisions leading to stronger effective interactions between them as a result of several positive feedback mechanisms. According to this point of view, the proximal triggering factor for price collapse should be clearly distinguished from the fundamental factor. A crash occurs because the market has entered an unstable phase towards the culmination of a bubble and any small disturbance or process may reveal the existence of the instability. Think of a ruler held up vertically on your finger: this very unstable position will lead eventually to its collapse, as a result of a small (or an absence of adequate) motion of your hand or due to any tiny whiff of air. This is the proximal cause of the collapse. But the fundamental cause should be attributed to the intrinsically unstable position.

What is then the origin of the maturing instability? Many studies have suggested that bubbles result from the over-optimistic expectation of future earnings and history provides a significant number of examples of bubbles driven by such unrealistic expectations (Kindleberger and Aliber, 2005; Sheffrin, 2005; Sornette, 2003a). These studies and many others show that bubbles are initially nucleated at times of burgeoning economic fundamentals in so-called "new economy" climates. This vocable refers to new opportunities and/or new technological innovations. But, because there are large uncertainties concerning present values of the economies that will result from the present innovations, investors are more prone to influences from their peers (Hong et al., 2005), the media, and other channels that combine to build a self-reflexive climate of (over-)optimism (Umpleby, 2007). In particular, these interactions may lead to significant imitation, herding and collective behaviors. Herding due to technical as well as behavioral mechanisms creates positive feedback mechanisms, which lead to self-organized cooperation and the development of possible instabilities or to the "building of castles in the air", to paraphrase Malkiel (1990). This idea is probably best exemplified in the context of the Internet bubble culminating in 2000 or the recent the CDO bubble in the USA peaking in 2007, where the new economies where the Internet or complex derivatives on sub-prime mortgages building on accelerating real-estate valuations.

Based on these ideas, the present paper adds to the literature by providing a detailed analysis of how the proximate triggering factor of a crash might relate to its fundamental mechanism in terms of a global cooperative herding mechanism. In particular, we rationalize the finding of Cutler et al. (1989) that exogenous news are responsible for no more than a third of the variance of the returns and that major financial crises are not preceded by any particular dramatic news.

In a nutshell, our multi-period many agent-based model is designed as follows. At each time step t, each investor forms an opinion on the next-period value of a single stock traded on the market. This opinion is shaped by weighting and combining three sources of information available at time t: (i) public information, i.e. news, (ii) information from their "friendship" network, promoting imitation and (iii) private information. In addition, we assume that the agents adapt their strategy, i.e., the relative importance of these different sources of information according to how well they performed in the past in predicting the next-time step valuation.

The a priori sensible qualities of our agents to gather all possible information and adapt to the recent past turn out to backfire. As their decisions are aggregated in the market, their collective impact leads to the nucleation of transient phases of herding with positive feedbacks. These nucleations occur as a result of random occurrences of short runs of same signed news. Our main findings can thus be summarized as follows: rallies and crashes occur due to random lucky or unlucky streaks of news that are amplified by the feedback of the news on the agents' strategies into collective transient herding regimes. In addition to providing a convincing mechanism for bubbles and crashes, our model also provides a simple explanation for the excess volatility puzzle (Shiller, 1981).

Before presenting the model and its results, it is useful to compare it with the relevant literature and related models. A related line of research aims at developing a theory of "convention" (Orléan, 1984, 1986, 1989a, 1989b, 1991, 1995), which emphasizes that even the concept of "fundamental value" may be a convention established by positive and negative feedbacks in a social system. A first notable implementation by Topol (1991) proposes a model with an additive learning process between an 'agent-efficient' price dynamics and a mimetic contagion dynamics. Similar to our own set-up, the agents of Topol (1991) adjust their bid-ask prices by combining the information from the other buyers' bid prices, the other sellers' ask prices and the agent's own efficient price corresponding to his knowledge of the economic fundamentals. Topol (1991) shows that mimetic contagion provides a mechanism for excess volatility. Another implementation of the concept of convention by Wyart and Bouchaud (2007) shows that agents who use strategies based on the past correlations between some news and returns may actually produce by their trading decisions the very correlation that they postulated, even when there is no a priori economic basis for such correlation. The fact that agents trade on the basis of how the information forecasts the return is reminiscent of our model, with however several important differences. The first important conceptual change is that Wyart and Bouchaud (2007) use a representative agent approach (in contrast with our heterogeneous agent framework), so that effect of imitation through the social network is neglected. The second difference is in the agent's calculation of the correlation to adapt their strategies. In Wyart and Bouchaud (2007), agents' strategies are controlled by the correlation between the news and the return resulting immediately from their aggregate action based on those news (taking into account the agents' own impact). Our agents' strategies are determined by the correlation between their information and the return one time step later, which embodies the more realistic situation, in which a postion first has to be open and then closed a time step later for the trade's payoff to be observed.

Another closely related line of research is known as "information cascades". According to (Bikhchandani et al., 1992), "an informational cascade occurs when it is optimal for an individual having observed the action of those ahead of him, to follow the behavior of the preceding individual without regard to his own information". In these models, agents know that they have only limited information and use their neighbors actions in order to complement their information set. Bikhchandani et al. (1992) showed that the fact that agents use the decisions of other agents to make their own decision will lead with probability 1 to an informational cascade under conditions where the decisions are sequential and irreversible. This model was later generalized by Orléan (1995) into a non-sequential version, where informational cascades were still found to be possible.

The concept of information cascades is not new in modeling bubbles. Chari and Kehoe (2003) developed a model where agents try to compensate their uncertainty about the a priori fixed payoff of an asset by observing all other agents' actions. In our model, agents are also using the opinions of their neighbors to determine how to act but the reason behind this is different. Our agents are not so much interested in the fundamental value of the stock, but more in its future directions. They try to buy the asset before its price rises and sell before it falls, making profit from the difference in the price. The true underlying equilibrium value is not the only important information to them, and they are more clever than purely fundamental value investors. They recognize that fundamental value is just one component among others that will set the market price. They include the possibility that the price may deviate from fundamental value, due to other behavioral factors. And they try to learn and adapt to determine what are the dominant factors. In principle, they should be able to discover the fundamental value and converge to its equilibrium. But it is a fact that they do not in some circumstances, due to the amplification of runs of positive or negative news in the presence of their collective behavior when sufficiently strong. In the "information cascade" set-up, one assume that the "truth" exists, that there is a true fundamental price or a correct choice to be made which is exogenously given, and agents have no influence on the outcome. In our model however the outcome, whether selling or buying a stock was the right choice, is endogenously emerging from the aggregated choices of all agents. There is no a priori right or wrong answer, it is decided during the process. Moreover, the strength of the influence of her neighbors onto a given agent is not constant in time. This influence by the social environment evolves in time according to its past relevance and success.

A model for the formation of a boom followed by a crash was also developed by Veldkamp (2005), where the price of an unknown company can rise only slowly due to infrequent news coverage. If the company performs well resulting in a slow boom, its susceptibility towards news increases as the media become more aware of the successful company so that, eventually, a single piece of bad news can induce a sudden crash. Although the subject of research is the same, we show how a boom can also be formed with news not being constantly positive and that a single piece of bad news does not necessarily lead to the burst a bubble.

The endogenization of the sources of information onto the decisions of the agents is inspired by the model of Zhou and Sornette (2007), which focuses on herding and on the role of "irrational" mis-attribution of price moves to generate most of the stylized facts observed in financial time series. Similarly to their model as well as many other artificial financial market models investigating the interaction between trading agents, our model is based on the Ising model, one of the simplest models describing the competition between the ordering force of imitation or contagion and the disordering impact of private information or idiosyncratic noise that promotes heterogeneous decisions (McCoy and Wu, 1973).

Our paper is organized into four sections. In Section 2, the detailed working of the model is presented. The results are shown and discussed in Section 3 and Section 4 concludes.

2. The model

2.1. General set-up

We consider a fixed universe of N agents who are trading (buying or selling) a single asset, which can be seen as a stock, the market portfolio or any other exchange traded asset. This asset is traded on an organized market, coordinated by a market maker. At each time step, agents have the possibility to either trade or to remain passive. The trading decision of a given agent is based on her opinion on the future price development.

To form their opinion, agents use information from three different sources: idiosyncratic opinion, global news and their network of acquaintances. In order to adapt their decision making process to the current market situation, they are weighting the different information sources by their respective past predicting performance. Limited to these sources of information, our agents act rationally, i.e., they use all information available to them to maximize their profits. Since they use backward looking adapting strategies with finite time horizons, our agents are boundedly rational, with limited competence, resources and available time.

A limitation of the model is to assume that agents do not have access to more liquidity than their initial wealth and that generated by their investments. Moreover, our universe has a fixed population, so that there is no flux of new "foreign" investors that may be attracted in the later stage of a bubble, and who could inflate it up further (Sornette and Zhou, 2004; Zhou and Sornette, 2006, 2008). We thus purposefully remove one of the mechanisms, namely the increasing credit availability and credit creation (Caginalp et al., 2001), which has often been reported as an important ingredient to inflate historical bubbles (Galbraith, 1997; Sornette, 2003b; Kindleberger and Aliber, 2005). This allows us to focus on the role of decision processes with conflicting pieces of information in the presence of local adaptation.

2.2. Three sources of information

At every time step, agents form anticipations concerning the future price movements based on three sources information.

A first source of information of a given agent is her private information, $\varepsilon_i(t)$, which may reflect the unique access to information not available publicly or the idiosyncratic, subjective view of the particular agent on how the stock will perform in the future. The private information is different for every agent, is taken uncorrelated across agents and time: the innovations $\varepsilon_i(t)$ are normally distributed ($\varepsilon_i(t) \sim N(0, 1)$) and i.i.d.

A second source is the public information, $n(t)$. Public information includes economic, financial and geopolitical news that may influence the future economic performance of the stock. To capture the idea that the public news, $n(t)$, is fully informational with no redundancy (Chaitin, 1987), we take $n(t)$ as a white Gaussian noise with unit variance, uncorrelated with the private information $\{\varepsilon_i(t), i = 1, \ldots, N\}$ of the agents. Although news are generated as a stationary process, we will see that their impact on the agents evolves because of the adaptive nature of the agents' strategies.

The third source of information for a given agent is provided by the expected decisions of other agents to whom she is connected in her social and professional network. With limited access to information and finite computing power (bounded rationality), it can be shown to be optimal to imitate others (Orléan, 1986; Roehner and Sornette, 2000). Moreover, there is clear empirical evidence that practitioners do imitate their colleagues (Hong et al., 2005). In our model, agents gather information on the opinions of their neighbors in their social network and incorporate it as an ingredient into their trading decision.

Incorporating agent interaction in the opinion formation process leads to dynamics described by models derived from the Ising model. Many earlier works have already borrowed concepts from the theory of the Ising models and of phase transitions to model social interactions and organization (e.g. Follmer, 1974; Callen and Shapero, 1974; Montroll and Badger, 1974). In particular, Orléan (1984, 1986, 1989a, 1989b, 1991, 1995) has captured the paradox of combining rational and imitative behavior under the name "mimetic rationality," by developing models of mimetic contagion of investors in the stock markets which are based on irreversible processes of opinion forming.

2.3. Opinion formation

Using the three sources of information described in the previous section, the opinion of agent i at time t, $\omega_i(t)$, consists of their weighted sum,

$$\omega_i(t) = c_{1i} \sum_{j=1}^{J} k_{ij}(t-1) E_i[s_j(t)] + c_{2i} u(t-1) n(t) + c_{3i} \epsilon_i(t), \tag{1}$$

where $\epsilon_i(t)$ represents the private information of agent i, $n(t)$ is the public information, J is the number of neighbors that agent i polls for their opinion and $E_i[s_j(t)]$ is the expected action of the neighbor j estimated by agent i at time t.[1] The functional form of expression (1) embodies our hypothesis that an agent forms her opinion based on a combination of different sources of information. This is a standard assumption in the social interaction literature (Bischi et al., 2006; Brock and Durlauf, 2001) and decision making theory (see for instance Körding, 2006).

To take into account the heterogeneity in trading style and preferences of traders, we assume that each agent i is characterized by a triplet of fixed traits, in the form of the weights (c_{1i}, c_{2i}, c_{3i}) she attributes to each of the three pieces of information (social network, news and idiosyncratic). The values (c_{1i}, c_{2i}, c_{3i}) for each agent are chosen randomly from three uniform distributions over the respective intervals $[0, C_1]$, $[0, C_2]$ and $[0, C_3]$, at the initialization of the system. In Section 2.4, we will extend this heterogeneity by allowing for different risk aversions.

In order to adapt to the recent market regime, each agent can modify the weights she attributes to the information from each of her neighbor j, via the factor $k_{ij}(t)$, and to the public news, via the factor $u(t)$. The factors $k_{ij}(t)$'s and $u(t)$ are updated such as to give more weight to an information source if it was a good predictor in the recent past, and to decrease its influence in the inverse case (more details in Section 2.7). The idiosyncratic term is not weighted and has a constant impact on agents actions.

Finally, for simplicity, our agents live on a virtual square lattice with $J = 4$ neighbors, with periodic boundary conditions. The reported results are not sensitive to this topology, and hold for random as well as complete graphs.

[1] We use a sequential updating mechanism with a random ordering. In this way, when agent i polls her neighbors, she has a mix of opinions coming from those who have already updated theirs and those have not yet. This procedure can be thought of as a device to account for the large distribution of reactions times of humans (Vazquez et al., 2006; Crane et al., 2010).

2.4. Trading decision

Until now, we have introduced heterogeneity between agents through their three personal traits (c_{1i}, c_{2i} for c_{3i}), unique to each agent, on how they combine information to form their opinion. Another important well-documented heterogeneity is that different people have different risk aversions. We capture this trait by assuming that each agent is characterized by a fixed threshold $\underline{\omega_i}$, controlling the triggering of an investment action, given her opinion level $\omega_i(t)$. An agent i decides to go long (buy a stock) if her conviction $\omega_i(t)$ is sufficiently positive so as to reach the threshold: $\omega_i(t) \geq \underline{\omega_i}$. Reversely, she decides to go short (sell a stock) if $\omega_i(t) \leq -\underline{\omega_i}$. Thus, we assume symmetric levels of conviction in order for a trade to occur either on the buy or sell sides. The parameter $\underline{\omega_i}$ captures one dimension of the agent's risk aversion: how much certitude she needs to break her hesitation and move into the market. The larger her threshold $\underline{\omega_i}$, the larger certitude about future price movements the agent requires in order to start trading. Each agent is characterized by a different $\underline{\omega_i}$, drawn randomly from a uniform distribution in the interval $[0, \underline{\Omega}]$.

As previously discussed in Section 2.1, our agents are liquidity constrained. The portfolio of an agent i at time t is the sum of her cash $cash_i(t)$ and of the number $stocks_i(t)$ of the single asset that is traded in our artificial market. When an agent decides to buy, she uses a fixed fraction g of her cash. When an agent decides to sell, she sells the same fixed fraction g of the value of her stocks. The fact that g is much smaller than 1 ensures time diversification. Our main results do not change significantly as long as g does not approach 1. Our agents are not allowed to borrow, because they can only buy a new stock, when they have the cash. Reciprocally, we impose short-sell constraints, in the sense that an agent can only sell a stock she owns. Thus, our model is related to the literature investigating the role of short-sale constraints (Miller, 1977; Chen et al., 2002; Ofek and Richardson, 2003).

These rules can be summarized in terms of the direction $s_i(t)$ of the trading decision and the volume $v_i(t)$ (in units of number of stock shares) of the agent i:

$$\begin{aligned}
\text{if} \quad \omega_i(t) > \underline{\omega_i}: \quad & s_i(t) = +1 \quad \text{(buying)} \\
& v_i(t) = g.\frac{cash_i(t)}{p(t-1)} \\
\text{if} \quad \omega_i(t) < -\underline{\omega_i}: \quad & s_i(t) = -1 \quad \text{(selling)} \\
& v_i(t) = g \cdot stocks_i(t),
\end{aligned}$$

where $p(t)$ is the price of the asset at time t. When an agent is buying assets, her order volume $v_i(t)$ is determined by her available cash and by the stock share price $p(t-1)$ at the previous time step (the main results remain unchanged if agents would use the expected $p(t)$ instead). Our agents are submitting market orders, such that the price to pay to realize an order is the new price $p(t)$ determined by the market maker. This new price is determined by the price clearing mechanism that aggregates the excess demand after all the traders have submitted theirs decisions.

2.5. Price clearing condition

Once all the agents have decided on their orders, the new price of the asset is determined by the following equations:

$$r(t) = \frac{1}{\lambda \cdot N} \sum_{i=1}^{N} s_i(t) \cdot v_i(t) \tag{2}$$

$$\log[p(t)] = \log[price(t-1)] + r(t), \tag{3}$$

where $r(t)$ is the return at time t and λ represents the relative impact of the excess demand upon the price, i.e. the market depth. Similar to Beja and Goldman (1980) and Wyart and Bouchaud (2007), we neglect all higher order contributions in expression (2) and use a linear market impact function, as a rough approximation at time scales significantly larger than the tick-per-tick time scales for which nonlinear impact functions are observed (Plerou et al., 2002).

Expressions (2) and (3) can be interpreted in two ways. One is that the trading is performed through a market maker, disposing of an unlimited amount of cash and stocks. Agents submit all their market orders to the market maker, who, after adapting the price to the excess demand, executes all the agents' trades. Because the market maker adapts the price before he executes the trades, he has a competitive advantage and gets on average a significant positive return for his service.

An alternative interpretation is that the trading style of our agents is midterm to longterm trading, excluding high-frequency traders like hedge-funds and such. Once our agents have absorbed their information and taken a trading decision, the price has already changed due to faster agents using similar trading information.

2.6. Cash and stock positions

We assume a frictionless market with no transaction fees. Once the return and the new price are determined by the market clearing Eqs. (2) and (3), the cash and number of stocks held by each agent i are updated according to

$$cash_i(t) = cash_i(t-1) - s_i(t)v_i(t)p(t) \tag{4}$$

$$stocks_i(t) = stocks_i(t-1) + s_i(t)v_i(t). \quad (5)$$

2.7. Adaptation

As described above, agents have pre-existing heterogeneous beliefs on the reliability of the three different sources of information, quantified by their three traits $c_{1/2/3i}$. In addition, we assume that agents adapt their belief concerning the credibility of the news $n(t)$ and their trust in the advice $E_i[s_j(t)]$ of their social contacts, according to time-dependent weights $u(t)$ and $k_{ij}(t)$, which take into account their recent past performance. Specifically, an agent estimates the value of a source of information by the correlation between the source's prediction and the realized return. For their strategy to be adapted to the current market regime, agents prioritize recent data in their calibration of the correlation. This prioritization of recent data is supported, first, by behavioral findings stating that individuals tend to overweight recent information and underweight prior data, second, by practitioners, who calibrate their trading strategies with recent data. The implementation of this prioritization is achieved by a standard auto-regressive update:

$$u(t) = \alpha u(t-1) + (1-\alpha)n(t-1)\frac{r(t)}{\sigma_r} \quad (6)$$

$$k_{ij}(t) = \alpha\, k_{ij}(t-1) + (1-\alpha)E_i[s_j(t-1)]\ \frac{r(t)}{\sigma_r} \quad (7)$$

Choosing $0 < \alpha < 1$ and with $0 < \sigma_r$,[2] the correct prediction of the sign of the realized stock return $r(t)$ from a given information source tends to reinforce the trust in that source of information, all the more so, the larger the return (scaled by its volatility σ_r) and the larger the strength of the signal. The time scale $1/|\ln(\alpha)|$ sets the memory duration over which past performance continues to impact the adaptive trust coefficients $u(t)$ and $k_{ij}(t)$. The update of u and k_{ij} via Eqs. (6) and (7) is performed at every time step.

3. Results of the model

3.1. General properties

Our model is an idealized "test tube" representation of a financial market and given the simplifications put into the model, we do not aim at reproducing faithful statistical characteristics of realistic price dynamics. Our objective is to obtain an understanding of how the interplay of news, herding and private information can lead to the formation of bubbles and crashes. We first point out a few properties of the model, that derive straightforwardly from our set-up.

Because we model a closed system, with no new influx of money or stocks after the initial endowment of $cash(0)$ and $stocks(0)$, there cannot be any money/wealth creation in the long run.[3] As a consequence, the price trajectory has an upper and lower bound.[4] The constraints on cash and stocks tend to push the price back to its initial value, $p(0) = 1$, such that the price performs a mean-reverting random walk[5] around its initial value, which will be refereed to in the following as the equilibrium price.

The adaptive process of our agents essentially consists in looking for persistent sources of information, which impact on the returns. In more detail: for a trade to be profitable, an agent has to first acquire a number of asset (at time t), then its value has then to increase in the following time step, explaining the offset of one time step between the information source and the realized return in Eqs. (6) and (7). The return $r(t+1)$ is however influenced by the information sources at time $t+1$, and not by those at time t, on which agents based their prediction of $r(t+1)$. This means that, for an information source to have some real predicting power, it must have some persistence (cf. Appendix A for a more detailed explanation of this mechanism).

3.2. First results

In our simulations, we fix the number of agents in the system to $N = 2500$, the market depth to $\lambda = 0.25$, the maximal individual conviction threshold to $\underline{\Omega} = 2.0$, the fraction of their cash or stocks that investors trade per action to $g = 2\,\%$, the initial amount of cash and stocks held by each agent to $cash_i(0) = 1$ and $stocks_i(0) = 1$, and the memory discount factor to $\alpha = 0.95$, corresponding to a characteristic time of $1/|\ln(\alpha)| \approx 20$ time steps. The news are modeled by i.i.d. Gaussian noise. Setting $C_1 = C_2 = C_3 = 1.0$, Fig. 1 shows a typical realization of the time evolution of the log-price $\log[p(t)]$, the one-time-step return $r(t)$, the prediction performance of the news, $u(t)$ and the ensemble average of the prediction performance of the neighbors, $\langle k_{ij} \rangle (t)$. The middle right panel shows the distribution of returns with clear evidence of a non-Gaussian fat tail structure. The lower right panel shows the absence of correlation between returns together with the presence of non-

[2] σ_r is in fact $\sigma_r(t)$, with $\sigma_r(t)^2 = \alpha \cdot \sigma_r(t-1)^2 + (1-\alpha)\cdot(r(t-1) - \langle r(t) \rangle)^2$ and $\langle r(t) \rangle = \alpha \cdot \langle r(t-1) \rangle + (1-\alpha)\cdot r(t-1)$.

[3] Strictly speaking, the model suffers however from a slight money destruction due to the price setting mechanism with the market maker in which the log-price change is linear in the excess demand. But the number of purchased stocks depends on the real price (=exp(log-price)). Therefore, a rapid increase

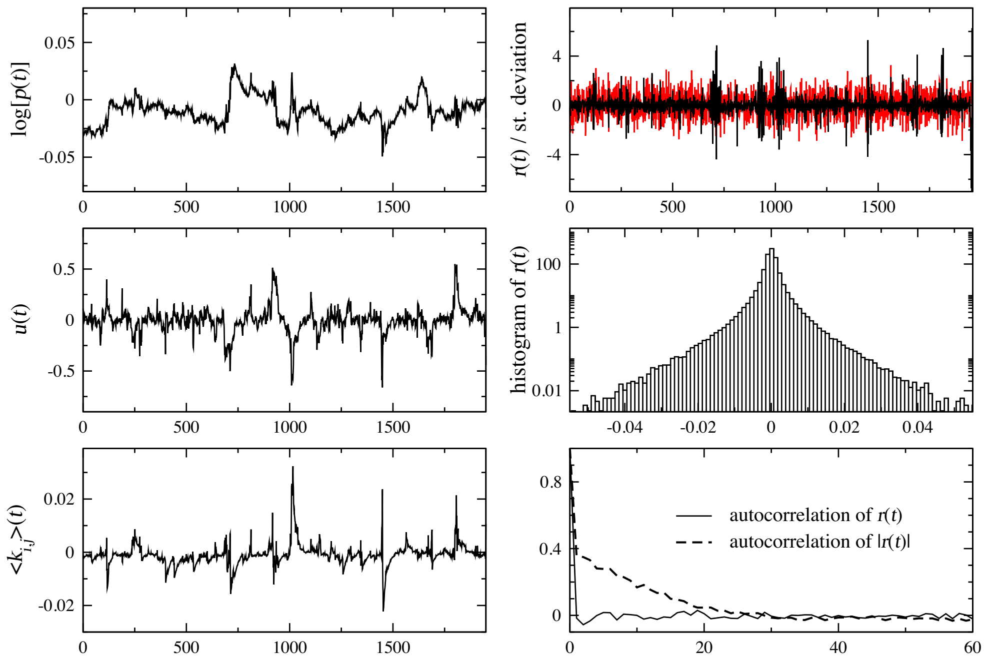

Fig. 1. This figure shows a typical realization of the major observables of the system. These observables are the time evolution of the price $p(t)$ (upper left panel), the one-time-step return $r(t)$ in black with clear evidence of clustered volatility (upper right panel) together with the news, $n(t)$, in the background in red, the news weight factor $u(t)$ (middle left panel) and propensity $\langle k_{ij} \rangle (t)$ to imitate (lower left panel). The middle right panel shows the distribution of returns: the linear-log scales would qualify a Gaussian distribution as an inverted parabola, a double-exponential as a double tent made of two straight lines; in contrast, one can observe a strong upward curvature in the tail of this distribution, qualifying a fat-tail property compatible with a stretched exponential or power law. The lower right panel shows the absence of correlation between returns together with the presence of non-negligible correlation of the volatility (here measured as the absolute value of the returns). Note the positive value of the correlation of the volatility up to a time about 25 time steps, followed by a small negative value up to 80 time steps. The time scale of the correlation of volatility is set by the memory factor $\alpha = 0.95$ corresponding to a characteristic time scale of 20 time steps. These results are obtained for $C_1 = C_2 = C_3 = 1.0$, and frozen weights attributed by the agents to the three information sources drawn out of a uniform distribution from 0 to C_1, C_2, C_3, respectively. The histogram and the correlation data are computed out of a realization with 6×10^4 time steps.

negligible correlation of the volatility (here measured as the absolute value of the returns), which confirms the clear evidence of clustered volatility in the time series of one-time-step returns.

While the perceived predicting power of the news, $u(t)$, fluctuates around its mean value of 0, it should be noted that it exhibits significant non-zero values, indicating that agents sometimes give a lot of importance to the news. If the agents were fully aware of the i.i.d. properties of the news, they would not use them.[6] But because of the adaptive nature of their strategy to the current market regime, agents do not use the complete price and news time series to update their trust into the news, but only recent data points.[7] Due to the use of a finite data set, the i.i.d. news may occasionally show persistence,[8] leeding to an increase of $u(t)$ as consequence. The statistical fluctuations associated with the random patterns that are always presents in genuine noise is misinterpreted by the agents as genuine predictability. It is the local optimization, that makes the agents see causality, where there is only randomness (Taleb, 2008).

The lower left panel of Fig. 1 shows the average propensity to imitate, which also fluctuates around 0. But, the amplitude of these fluctuations is much reduced compared to those of $u(t)$. This is because each agent updates individually her propensity to imitate her neighbors according to (7), so that the statistical average $\langle k_{ij} \rangle (t)$ is performed over the whole heterogenous population of agents, compared with no average for $u(t)$ which is common knowledge to all agents.

---

followed by a slow decrease of the price decreases the total wealth of the system, by the concavity of the logarithmic function. This effect is essentially negligible.

4 The upper bound is reached once agents have exhausted all their cash. The lower bound is then the agents are all in cash.

5 The increments of the walk are however not distributed according to a normal law, but to a distribution with fatter tails (cf. Fig. 1) due to the adaptive strategies of the agents.

6 Recall that the return $r(t+1)$ is influenced by the news at time $t+1$ on which agents based their prediction of $r(t+1)$, and not by those at time t. Because the news have no true persistence, they can not have true predictive power.

7 The weight of a data point in the update of $u(t)$ decreases exponentially with increasing age with a time scale $\sim 1/|\ln(\alpha)|$.

8 Our agents do not have a PhD in Econometrics and they do not perform proper statistical tests of their hypotheses.

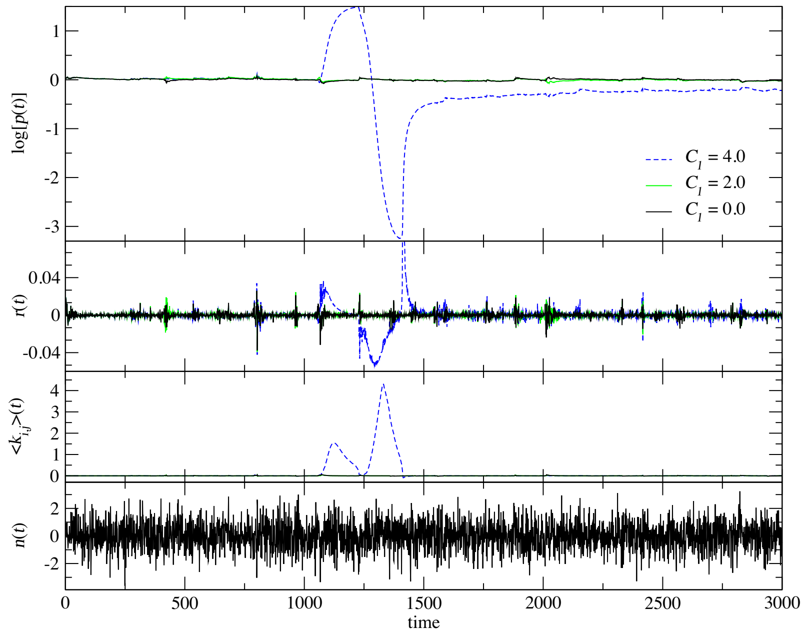

Fig. 2. Evolution of several variables for $C_1 = 0.0$, 2.0, 4.0 with the other parameters unchanged, including the random seed for the three realizations, resulting in the same realization of the news for the different runs (shown in the bottom panel). Other parameters are: $N = 2500$, $C_2 = C_3 = 1.0$, $\underline{\Omega} = 2.0$, $\alpha = 0.95$, $\lambda = 0.25$, $g = 0.02$.

The crucial parameters of our model are the parameters C_1, C_2, C_3, which control the level of heterogeneity and the a priori preference for the three different types of information. Changing these parameters changes the way the agents behave in ways that we now explore systematically.

3.3. C_1-dependence

Each agent is endowed with a fixed individual preference level, c_{1i}, controlling how much she takes into account the information stemming from the actions of their neighbors. This level is different from agent to agent, and is drawn from a uniform distribution in the interval $[0, C_1]$. Thus, the parameter C_1 sets the maximal and mean $(= C_1/2)$ innate weight, that agents give to their social influences.

Fig. 2 plots the evolution of several variables for three different values of C_1, all other parameters, including the seed of the random number generator, remaining the same. For vanishing propensity to imitate ($C_1 = 0$), some price spikes can be observed, which are generated by the news only, whose influence can be amplified by the positive feedback resulting from adaptation that tends to increase the relevance that investors attribute to news after a lucky run of news of the same signs. For $C_1 = 2.0$, one can observe that these peaks are amplified due to the imitation now also contributing to the agents' actions. For $C_1 = 4.0$, a qualitatively different price evolution appears. For such large values of the maximal susceptibility to their social environment, the price is driven to its extremes, its dynamics being only slowed down by the agents' finite cash and stock portfolio reaching their boundaries. We show below that this extreme behaviors results from a self-fulfilling prophecy, enabled through social interactions.

To better illustrate the effect of increasing C_1, the third panel in Fig. 2 shows the average weight factor $\langle k_{ij} \rangle (t)$[9] used by the agents to assess the relevance of the information stemming from their neighbors. By increasing C_1, agents are by definition more susceptible to their neighbors' opinions, making them more likely to act in the same way if they show some predictive power. Consequently, since the price dynamics is governed by the aggregate demand, herding in opinions leads to persistent returns, creating the very returns agents hoped for, which reinforce the prediction power of their neighbors in a positive feedback loop. With $\langle k_{ij} \rangle (t)$ and C_1 large, the opinions of the agents are completely shaped by their social component, while the news and their idiosyncratic term are essentially ignored. Due to this positive feedback loop, a small predictive success of some agents can trigger an avalanche of self-fulfilling prophecies, leading to price dynamics completely unrelated to the news and to large price deviations from the assets to its equilibrium value.

[9] $\langle k_{ij} \rangle (t) = 1/(N \cdot J) \sum_{i=1}^{N} \sum_{j=1}^{J} k_{ij}(t)$.

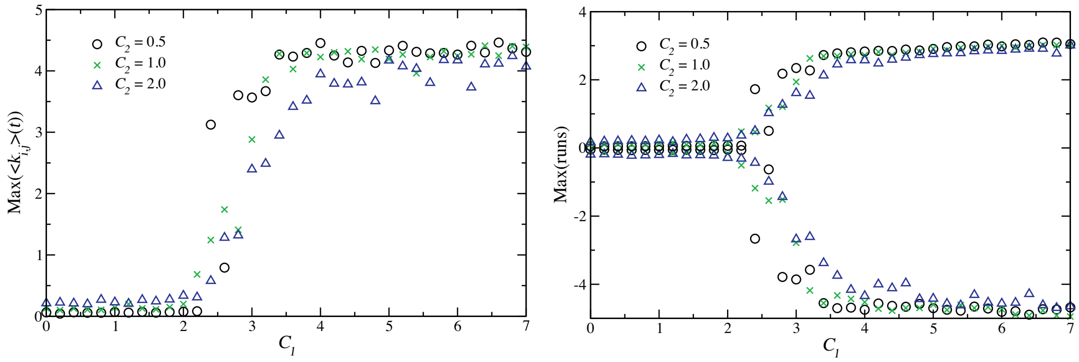

Fig. 3. Impact of C_2, the innate susceptibility to the news, onto the transition from the efficient to the excitable regime in function of C_1, the innate susceptibility to neighbors' actions. The transitionis measured by the maximal value of $\langle k_{ij} \rangle (t)$ (the recent prediction performance of agents' neighbors) (left) and the extremal draw-down and -ups (sum of consecutive same signed returns) (right), both averaged over many realizations. Both plots show that with increasing C_2, the critical C_1-value increases and the transition smoothens. Other parameters are: $N = 10^4$, $C_3 = 1.0$, $\underline{\Omega} = 2.0$, $\alpha = 0.95$, $\lambda = 0.25$, $g = 0.02$.

3.4. The existence of two regimes

The existence of a bifurcation beyond which a new regime appears is documented in Fig. 3 (left), where the maximum value of $\langle k_{ij} \rangle (t)$, averaged over many realizations and simulated with the same parameters $C_2 = C_3 = 1$, is plotted as a function of C_1. One can observe a rather abrupt transition occurring at around $C_1 = 3$. This transition is related to the phase transition of the Ising model, on which our model is based. Due to the presence of the adaption induced feedback loops, and with the dynamical character of $\langle k_{ij} \rangle (t)$, the precise nature of this transition can not be asserted. The existence of this change of regime explains the radical difference of properties shown in Fig. 2 for $C_1 = 0$, 2 to 4. The jump in $\langle k_{ij} \rangle (t)$ at $C_1 \approx 3$ is mirrored by a similar transition in the values of the maximal draw-downs (sum of consecutive negative returns) and draw-ups (sum of consecutive positive returns) as a function of C_1 in Fig. 3 (right). For $C_1 > 3$, a second regime is revealed where very large price moves occur. These market events are fundamentally different from the price fluctuations in the regime for $C_1 < 3$. These large price changes are reminiscent of the "outliers" documented by Johansen and Sornette (1998, 2001) and Sornette (2003a), and recently extended into the concept of "dragon-kings" (Sornette, 2009).

3.4.1. The efficient regime

For $C_1 < 3$, agents do not attribute sufficient importance to their neighbors in order to trigger the feedback loop that would lead to strong synchronized actions as occurs for large C_1s. For small C_1, the market is approximately efficient, in the sense that there is no autocorrelation of returns, as shown in Fig. 1, and the price fluctuates rather closely around its equilibrium value. While the major source of fluctuations are the news modeled as a Gaussian white noise process, the price fluctuations develop strong non-Gaussian features, as a result of the combined effect of the adaptive process that tends to amplify runs of same signed news and of the propensity to imitate that leads to small but non-negligible collective behaviors.

A first interesting conclusion can be drawn that our model provides a natural setting for rationalizing the excess volatility puzzle (Shiller, 1981), through the adaptive process of our agents. It could be argued that our setting is too simplified and unrealistic. But, how do real investors, traders, fund managers access the value their investment decisions? Necessarily by performing some kind of comparisons between the realized performance and some benchmarks, which can be a market portfolio, the results of competitors, the ex ante expectations, all the above or others. The adaptive process used by our agents is arguably a simple and straightforward embodiment of the tendency for investors to adjust their strategies on the basis of past recent performance, here on how well the news predicted the market returns. Because measurements are noisy, the resulting estimation leads unavoidably to an amplification of the intrinsic variability of the news into a much strong variability of the prices, i.e., to the excess volatility effect. Somewhat paradoxically, it is the attempt of industrious investors to continuously adapt to the current market situation, which leads to the dramatic amplification of the price volatility. This may be thought of as another embodiment of the "illusion of control" effect, found in the Minority and the Parrondo games (Satinover and Sornette, 2007a, 2007b, 2009), according to which sophisticated strategies are found to under-perform simple ones.

3.4.2. The excitable regime

A population of agents characterized by $C_1 > 3$, represents the situation in which many agents know that their idiosyncratic information and the news are incomplete. In order to compensate for this lack of information, agents tend to imitate the actions of successful acquaintances. Under these conditions, the average propensity to imitate, $\langle k_{ij} \rangle (t)$, exhibits extreme values, resulting in large price deviations from the equilibrium price and periods of persistent returns, as shown in Fig. 2. In this regime, the market is in an excitable state. By imitating the opinions of recent winners who profited from some

146 G. Harras, D. Sornette / Journal of Economic Behavior & Organization 80 (2011) 137–152

departure of the market price from its equilibrium value, our agents tend to amplify this anomaly, further strengthening the attraction of this strategy for other agents, eventually ending in a bubble and crash.

The triggering event responsible for the increasing weight that agents entrust to their neighbors' opinions is nothing but the random occurrence of a sequence of same signed news. As explained in Section 3.2, the weight $u(t)$ of the news in their opinions is increased when the agents perceive a pattern of persistence in the news, which also induces persistent returns. Then, the agents reassess their belief and give more importance to the news. Because the pattern of persistence of the news is common knowledge, this tends to align the decisions and actions of the agents. As a consequence, their aggregate impact makes happen the very belief that initially led to their actions, thus increasing the prediction power of the agents' opinions. As a whole, the agents see that the opinion of their friends is accurate, thus tending to increase their trust. This increase in the propensity $\langle k_{ij} \rangle (t)$ to imitate can lead to a cascade of trading activity, resulting in the rise of a bubble, as described in details in the appendix. This scenario is a cartoon representation of the well-documented fact that many bubbles start initially with a change of economic fundamentals. Translated in our agent-based model, this change of economic fundamentals is nothing but the streak of same signed news that tells the story of an increasing market (for positive news). This small positive bias can be sufficient to nucleate a process that eventually blossom into a full-fledged bubble. The corresponding amplification of the news put the price on an unsustainable trajectory. This occurs especially when the system lives in the excitable state, in which the price can easily overshoot the values implied by the good/bad news.

Once such a cascade has begun and the best strategy is to follow the herd, agents are, in the case of a bubble, buying stocks at every time step and pushing the price up till they have no money left to further increase the price. At this point, their predicting power decreases due to they decreasing impact on the returns and the cascade ends. As financial bubbles feed on new money pouring in the market, the lack of new liquidity is a well-known factor of instability for financial bubbles (Kindleberger and Aliber, 2005; Hussam et al., 2008). Following this buying phase, the portfolio of agents consists mainly of stocks, biasing their actions towards selling. Now, some randomly occurring negative news are sufficient to trigger a reverse cascade, the crash, leading to an overshoot of the price below its equilibrium value. This scenario provides a clear distinction between the fundamental cause of the crash (the unstable high position of the price that has dried up all liquidity available) and the triggering proximal factor (a random occurrence of a sequence of negative news). In line with many observations, crashes in our model do not need a dramatic piece of negative information. Only a trickle can trigger a flood once the market as a whole has evolved into an unsustainable unstable position.

Due to the symmetry between buying and selling in our model, the price can, starting from its equilibrium value, depart in either direction, creating either a bubble (over-valuation of the asset over an extended period of time) or an negative bubble (under-valuation of the asset). This deviation will then be ended by either a crash (fast drop in price after a bubble) or a rally (fast appreciation after a negative bubble). In our analysis, we concentrate on the case of bubbles followed by a crash, because this is the more common scenario. The reason for this is twofold. First, in real markets, short-selling can occur, but is not equally available to all market players. The second reason has behavioral origins. In a bullish regime, people are progressively attracted to invest in financial markets, tending to push the price upward. Once invested, their attention is more focused on the financial markets. Fear and greed often lead to over-reactions and possible panics when the sentiments become negative, triggering herd selling which self-fulfills the very fears at their origin (Veldkamp, 2005).

Another interesting characteristic of the herding regime occurring for $C_1 > 3$ is that it is very difficult to diagnose this regime from the properties of the price recorded outside those transient episodes of booms or crashes. Indeed, outside these special moments of "exuberance", the market behaves as if in the regime $C_1 < 3$. Bubbles and crashes do not belong to the normal regular dynamics of the model. They are only experienced when certain conditions are fulfilled, as explained above, that combine to create these transient instabilities. They can thus be considered as "outliers" in the sense of Johansen and Sornette (1998, 2001), or using a better more colorful terminology, they are "dragon-kings" (Sornette, 2009). The statistical analysis of the distribution of f $\langle k_{ij} \rangle (t)$ confirms this claim. Fig. 5 shows the appearance of an extremely fat tail in the distribution of $\langle k_{ij} \rangle (t)$ over the ensemble of different realizations as a function of time for $C_1 = 4$, while its bulk remains approximately identical to the distribution obtained for the smaller values of C_1 below the critical threshold $\simeq 3$. This confirms the existence of a class of transient regimes, the booms and crashes, which coexisting with the normal dynamics of the prices.

The hidden nature of the regime associated with $C_1 > 3$ and the random occurrence of triggering news lead to the prediction that advanced diagnostics of bubbles and crashes should lead to numerous false alarms. Consider the study of Kaminsky (1998), who has compiled a large list of indicators of financial crises, suggested by the fundamentalist literature on the period from 1970 to 1995 for 20 countries. Out of the 102 financial crises in her database, she finds that the specificity of the indicators is quite low: only 39% of the ex ante diagnostics coincided with a crisis, suggesting that fundamental reasons should be expanded by behavioral ones to explain the emergence of crises.

We thus come to our second important conclusion: the present model provides a simple mechanism for the existence of two populations in the distribution of prices, exemplifying the concept that booms and crashes are qualitatively different from the rest of the price moves. The second population of boom-crash (dragon-kings) appears when the innate propensity to herd reaches a threshold above which a self-reinforcing positive feedback loop starts to operate intermittently.

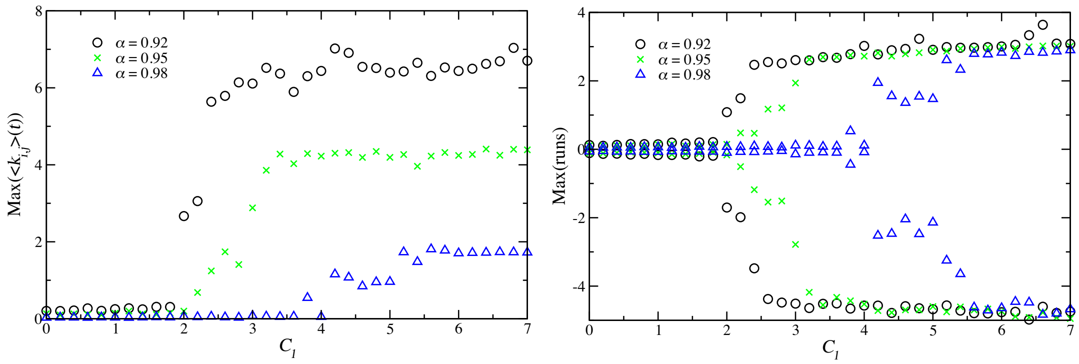

Fig. 4. Same as Fig. 3, except that C_2 is fixed at 1.0 and that the impact of α, which fixes the length of the time-span over which agents measure the predicting power of the different sources of information, is investigated. For larger α, the critical C_1-value increases, the transition smoothens and the largest possible $\langle k_{ij} \rangle(t)$ values is decreased. Other parameters are: $N = 10^4$, $C_3 = 1.0$, $\underline{\Omega} = 2.0$, $\lambda = 0.25$, $g = 0.02$.

3.5. C_2-dependence

Fig. 3 displays the impact of C_2, the a priori importance of the news, onto the transition from the efficient to the excitable regime. Both panels show that the more the agents trust the news, the stronger C_1 has to be for the system to become excitable. With increasing C_2-values, we also observe that the transition becomes smoother and the maximum $\langle k_{ij} \rangle(t)$ is decreased. This corresponds well to the intuition that, if traders are well informed and believe that the news correctly describe the economy, such drastic over- and under-valuations are less likely to happen and a higher level of panic is needed for a crash to happen.

3.6. α-Dependence

Although the presence of the α-parameter is crucial, its specific value (within a certain range below 1) has only a minor importance. Recall that α sets the time-scale of the market regime, since it controls the length of the time series that is used to estimate the predicting power of the different information sources. The fact that $\alpha < 1$, i.e. that agents' strategies are designed to identify local market regimes is the reason that makes them possible in the first place. It is the local adaptation, which is the true origin of a bubble and a subsequent crash.

Fig. 4 shows the impact of α on the transition from the efficient to the excitable regime. The closer α is to 1, the larger the critical C_1 for which the system becomes excitable. With a larger memory, the growth of the propensity $\langle k_{ij} \rangle(t)$ to imitate is more limited because the agents see now much better the bigger picture and are less easily carried away by a temporal coordination of their neighbors. Changing α leaves the maximal drawn-downs and -ups unchanged because, once an coordination of the agents starts, the only way to stop it is via the drying up of their cash/stock-reservoir.

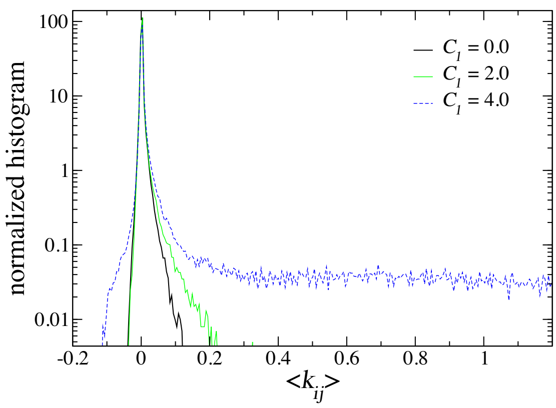

Fig. 5. Normalized histogram of $\langle k_{ij} \rangle(t)$ for different values of C_1.

3.7. Alternative clearing condition

We have played with variations of the implementations of the clearing conditions with different market maker' strategies. For instance, when the market make is adapting the price after (rather than before) the exchange of assets with the agents, we find the same qualitative results and bubbles and crashes occur by the same mechanism. With such an price clearing condition however, news bear real predictive power, destroying the efficiency of the market.

In general, the present model is very robust with respect to changes in the different ingredients. As long as agents can interact and are locally optimizing their strategies, bubbles and crashes do appear.

4. Conclusion

In this paper, we have addressed two major questions:

- Why do bubbles and crashes exist?
- How to they emerge?

We approached these questions by constructing a model of bounded rational, locally optimizing agents, trading a single asset with a very parsimonious strategy. The actions of the agents are determined by their anticipation of the future price changes, which is based on three different sources of information: private information, public information (news) and information from their neighbors in their network of professional acquaintance. Given these information, they try to maximize their usefulness by constantly scanning the market and adjusting the weight of the different sources to their opinion by the recent predicting performance of these sources. In this way, they are always adapting their strategy to the current market regime, such that they can profit from an opportunity if it arises.

We find that two regimes appear, depending on how strong the agents are influenced by their neighbors (controlled by the parameter C_1). In the regime of small C_1s, the low herding/efficient regime, agents are sometimes more influenced by the news and sometimes more by their neighbors, but due to the small level of trust they put into their neighbors by default, they do not get carried away in over-imitating their neighbors if the latter, for a short time interval, seem to be good predictors. The returns are mostly driven by the global and idiosyncratic news. The resulting market is approximately efficient, with the price not deviating much from its equilibrium value.

We find that the return distribution is however quite different from that describing the exogenous news. Our simple agents are able to transform the string of independent normally distributed news (both for the global and idiosyncratic news) into a return distribution with fatter than exponential tails, showing a clear sign of excess volatility. Also clustered volatility and a non-zero autocorrelation in the volatility of the returns are observed while the returns themselves remain uncorrelated, in agreement with the absence of arbitrage opportunities (at least at the linear correlation level). These different properties show that our simple model can reproduce some important stylized facts of the stock market, and can motivate the possibility to test its prediction in other market regimes.

By increasing C_1 above a certain critical value, the system enters a second regime where the agents give on average more importance to their neighbors' actions than to the other pieces of information. By increasing the awareness of their neighbors' actions, agents are more likely to coordinate their actions, which increases the probability that the direction of the return results in the predicted direction, which then again increases their trust in these successful predicting neighbors. Due to this positive feedback loop, the average coefficients $\langle k_{ij} \rangle$ (the dynamic trust of agent i in agent j) can surpass a critical value and the agents' opinions are dominated by only this information term, resulting in series of consistently large same-signed returns. Because the agents are always trying to maximize their returns, it is rational for the agents to follow the majority and to "surf" the bubble or the crash. This regime is characterized by large deviations from the equilibrium price resulting from a coordination of the agents' actions due to their local adaptation of their strategy to the mood of the market. Not only is it rational to follow the herd, we have also showed (in the appendix) that the agents who are early imitators of their successful neighbors in the early stage of a bubble/crash are those who will accumulate the largest wealth among all the agents after the market has returned to its normal regime.

We showed that the origin of these large deviations from the equilibrium price nucleate from the news. A random occurrence of a sequence of same signed news pushes the price in one direction and starts the coordination process of the agents. This situation is reminiscent from the mechanism for the initiation of real world bubbles, where an innovation leads to a period with a majority of positive news, which also move the market. Because the reason of the positive market move is an innovation, the agents are not entirely sure of its intrinsic value and seek advice from some of their professional collegues. If those colleagues tell them that they made large profits with this asset and they trust these colleagues, they will follow their advice, resulting in the same kind of behavior as produced by our model.

By following each other's actions, the agents push the price up, beyond its equilibrium value, up to an unsustainable level. Once the hype has cooled of and the agents have invested all their cash into the stock, just a little push by negative news can cause the price to collapse, resulting in a crash, without any apparent reason.

By increasing the prior propensity C_1 to imitate to a high value, the average behavior and properties of the dynamics of the model is unchanged. Outside of these large price variations occurring during rare bubbles and crashes, the dynamics looks similar to that documented in the low C_1 regime, i.e., appears to function like an efficient market. Therefore, attempting to

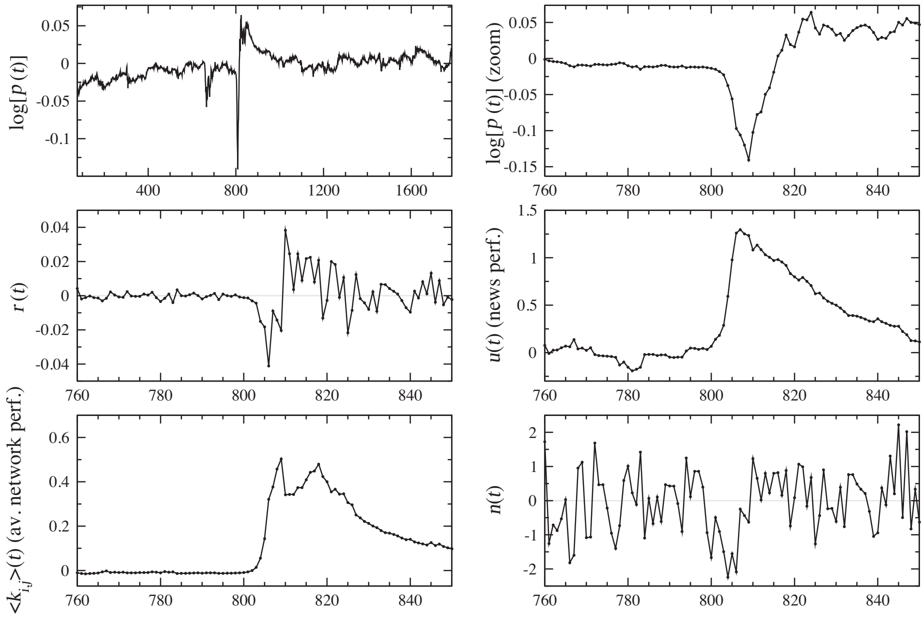

Fig. A.1. Time series of several key variables showing the response to the news for $C_1 = 1$ (efficient regime) and $C_2 = C_3 = 1.0$. Upper left panel: a portion of the price time series with a drop and rebound. Upper right panel: a magnification of the upper left panel around the increased volatility. Middle left panel: the time series of the returns. Middle right panel: the weight $u(t)$ of the news showing a fast growth over the time interval in which the news are all negative, followed by a decay over a time scale given by $1/\ln(\alpha) \approx 20$ time steps. Lower left panel: The average weight $\langle k_{ij} \rangle (t)$ of the propensity to imitate also exhibits a fast acceleration followed by a slower decay. Lower right panel: the time series of news, generated as a white noise, which can nevertheless exhibit runs of same-sign values.

estimate the value of C_1 just from the normal price dynamics is essentially impossible. The occurrence of a bubble/crash is an event that has drastically different statistical characteristics than the normal price fluctuations, exemplifying the occurrence of "outliers" (or "dragon-kings") that have been documented empirically for financial draw-downs (Johansen and Sornette, 1998, 2001; Sornette, 2003a).

Acknowledgments

We would like to thank Wei-Xing Zhou for invaluable discussions during the course of the project and Gilles Daniel and Ryan Woodard for a critical reading of the manuscript.

Appendix A. Detailed analysis of the emergence of a bubble

To clarify the mechanisms leading to the dynamics of the here presented model, we now illustrate some details on the micro-scale dynamics of the model. First we will go step-by-step through an occurrence of increased volatility, shown in Fig. A.1, with the system being in the efficient regime and explain in detail the relationships between the different variables.

Second, we will investigate the emergence of a bubble in the excitable regime and compare it to the dynamics resulting from the same stream of news in the efficient regime in Figs. A.1 and A.2.

A.1. Step-by-step description of the dynamics in the efficient regime

Fig. A.1 displays the dynamics of the key variables around the time $t = 800$, where the price suddenly, crashes, rebounds and then slowly relaxes to its pre-existing level. An increase of $u(t)$, the news' performance and $\langle k_{ij} \rangle (t)$, the average weight used by the agents to assess the relevance of the information stemming from their neighbors, is occurring at the same time.

The origin of this burst can be traced back to the random occurrence of a sequence of same signed news, shown in the lower right panel of Fig. A.1. Recall that we assume that the news are independently and identically distributed. Thus the dip structure in the news' realization is purely "bad luck", i.e. a stream of small bad news impact the market. The response of the agents to these run of bad news develops as follows. The observation of the news $n(t)$ gives the agents an information about the next return $r(t+1)$, but in order to profit from this insight, the agents have to act before $t+1$, i.e. they use $n(t)$ to buy or sell at time t. Therefore, a burst of activity, which has its origin in the news, can only occur if the sign of the news is, by chance, the same for several time steps as it is the case from $t = 799$ to 809.

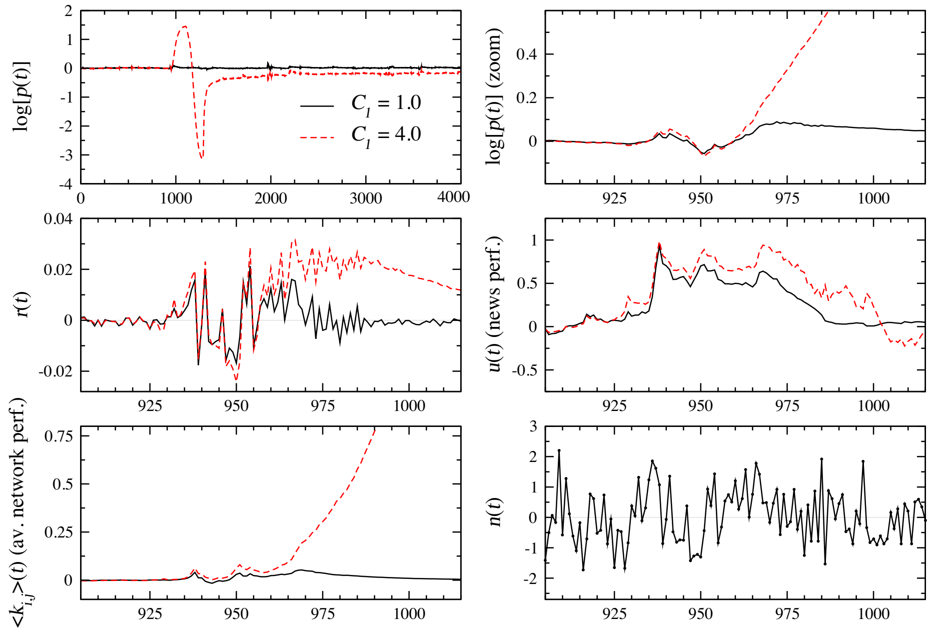

Fig. A.2. Time series of several key variables showing the response to the news for $C_1 = C_2 = C_3 = 1.0$ (efficient regime, in black) and $C_1 = 4$ with $C_2 = C_3 = 1.0$ (excitable regime, red dashed). We can observe that the random occurrence of persistence in the news stating around $t = 930$, initiates a bubble if agents give too much weight to their neighbors' actions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Let us report minutely the micro dynamics of the model to better understand this burst of activity. At $t = 799$, the news turns out to be negative, which suggests to the agents that the price may drop from $t = 800$ to 801. To prevent their portfolio from losing in value from time $t = 800$ to 801, some agents reduce their exposure to the market and sell a fraction gof their assets at $t = 799$. If enough agents listen to the news, as it is in this case, this selling will result in a negative return from $t = 799$ to 800. Then at $t = 800$, the news is, by chance, again negative resulting again in a negative return from $t = 800$ to 801. This negative return confirms the negative news from $t = 799$, leading agents to increase $u(t)$, the weight they attribute to the news.

The exponential growth of the weight u continues as long as the sequence of negative news goes on, further amplifying the impact of the news on the agents' decision and therefore on the price. Note that the average weight $\langle k_{ij} \rangle (t)$ of the propensity to imitate also exhibits a fast acceleration. This is due to the fact that the agents find that imitation is also a good predictor of the returns, since a majority of agents are following the news and are trading into the same direction. By this process, there is an amplification of the response of the whole herd to the exogenous news. When the run of bad news stops, it takes about $\simeq 1/\ln(\alpha) \approx 20$ times steps for u to relax back to its previous value. In this example, the maximum of $u(t)$ occurs at $t = 807$. At $t = 808$, u decreases lightly due to the small amplitude of the news at $t = 807$. Once u has reached a certain level, the news completely dominates agents opinion and thus also determines the returns.

At $t = 810$, the sequence of negative news is terminated by positive news, resulting in a large positive return due to the large value of $u(t)$. Furthermore, as the news at $t = 809$ predicted a negative return at $t = 810$, the predictive power of the news seems to have decreased, having a a decrease of $u(t)$ as consequence. Now that the news resumes its usual random switching signs and it is no longer a good predictor of the return, u decreases exponentially.

This case study illustrates that the occurrence of bursts of price variations is nothing but the amplification of runs of same-sign news, which leads to an exponential growth of the news weighting factor u, which itself increases dramatically the sensitivity of the agents to all future news. This heightened sensitivity lasts over a characteristic scale determined by the coefficient α governing the memory of the adaptation process (Eq. (7)). The process of agents' adaptation to the news and information from their neighbors, together with the random lucky or unlucky occurrence of runs of news of the same quality, is at the origin of the occurrence of this period of increase volatility.

A.2. Efficient regime vs excitable regime

In Fig. A.2, we plot the detailed dynamics during the nucleation of a bubble. The black continuous lines display the evolution of several variables with the system being in the efficient regime, i.e. $C_1 = C_2 = C_3 = 1.0$. The red dashed lines represent the same variables, with all parameters unchanged (including random seed), except that agents susceptibility to their neighbors'

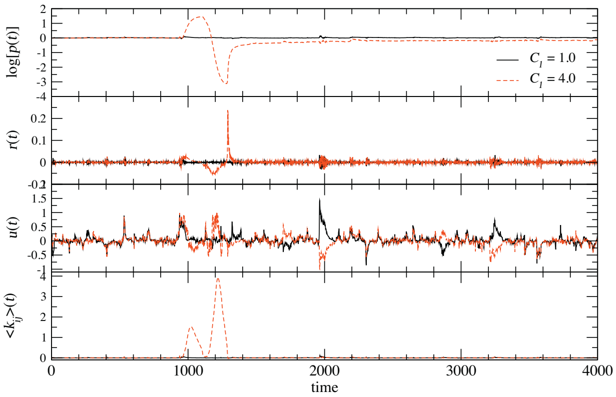

Fig. A.3. Magnification of the realization of the crash shown in Fig. A.2. Top to bottom: plots of the price, return, activity, news weight factor and average imitation factor, as a function of time.

actions is increased ($C_1 = 4.0$) such that the system is in the excitable regime. Fig. A.2 shows the detailed nucleation of the bubble, whereas Fig. A.3 shows the dynamics on a larger time scale.

In the second panel on the left in Fig. A.2, we plot the evolution of the return and witness a burst of volatility starting around $t = 930$. The origin of this volatility can be attributed to a random occurrence of some 'persistence' in the news $n(t)$, as explained in the previous section. This persistence increases the news' prediction power $u(t)$ and, because all agents are subject to the same news, agents' actions tend to synchronize, inducing an increase of the prediction power of their neighbors, $\langle k_{ij} \rangle (t)$.

Up to $t = 957$, the dynamics of the system in the efficient regime only deviates marginally from those of the excitable regime. After $t = 957$, their differences become apparent. In the excitable regime, where the neighbors' influence is a stronger factor in the opinion formation, the long string of positive news from $t = 957$ to 969 is able to increase the average interaction weight $\langle k_{ij} \rangle (t)$ up to a level high enough, such that the opinion of the agents, and therefore also their actions, are dominated by their neighbors'. As a consequence, the price continues to increase even after the positive news sequence has ended. In the efficient regime, on the other hand, the volatility of the returns, $u(t)$ and $\langle k_{ij} \rangle (t)$ return to their normal values after the 'luck streak' of positive news.

In the excitable regime, as a consequence of the strong propensity to interact, once a price rally or a crash is started, the dominating impact of the herd both in the adaption process and in the price impact makes the price trend self-reinforcing and basically independent of the sign of the news. This explains the very large amplitude of the price deviation compared with the price in the efficient regime in Figs. A.2 and A.3.

References

Abreu, D., Brunnermeier, M.K., 2003. Bubbles and crashes. Econometrica 71 (1), 173–204.

Barro, R., Fama, E., Fischel, D., Meltzer, A., Roll, R., Telser, L., 1989. In: Kamphuis Jr., R.W., Kormendi, R.C., Watson, J.W.H. (Eds.), Black Monday and the Future of Financial Markets. Mid American Institute for Public Policy Research, Inc. and Dow Jones-Irwin, Inc..

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

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. The Journal of Political Economy 100 (5), 992–1026.

Bischi, G.-I., Gallegati, M., Gardini, L., Leombruni, R., Palestrini, A., 2006. Herd behavior and nonfundamental asset price fluctuations in financial markets. Macroeconomic Dynamics 10 (04), 502–528.

Brock, W.A., Durlauf, S.N., 2001. Discrete choice with social interactions. The Review of Economic Studies 68 (2), 235–260.

Caginalp, G., Porter, D., Smith, V.L., 2001. Financial bubbles: excess cash, momentum, and incomplete information. The Journal of Psychology and Financial Markets 2 (2), 80–99.

Callen, E., Shapero, D., 1974. A theory of social imitation. Physics Today (July), 23–28.

Chaitin, G.J., 1987. Algorithmic Information Theory, Cambridge Tracts in Theoretical Computer Science. Cambridge University Press.

Chari, V.V., Kehoe, P.J., 2003. Hot money. Journal of Political Economy 111 (December (6)), 1262–1292.

Chen, J.S., Hong, H.G., Stein, J.C., 2002. Breadth of ownership and stock returns. Journal of Financial Economics 66, 171–205.

Crane, R., Schweitzer, F., Sornette, D., 2010. New power law signature of media exposure in human response waiting time distributions. Physical Review E 81, 056101.

Cutler, D.M., Poterba, J.M., Summers, L.H., July 1989. What Moves Stock Prices? Tech. Rep. 2538. National Bureau of Economic Research.

Dudley, W.C., 2010. Asset Bubbles and the Implications for Central Bank Policy. Speech by Mr William C Dudley, President and Chief Executive Officer of the Federal Reserve Bank of New York, at the Economic Club of New York, New York City, 7 April 2010 , p. 1.

Follmer, H., 1974. Random economies with many interacting agents. Journal of Mathematical Economics 1 (March (1)), 51–62.

Galbraith, J., 1997. The Great Crash, 1929. Houghton Mifflin Co., Boston.

Hong, H., Kubik, J.D., Stein, J.C., 2005. Thy neighbor's portfolio: word-of-mouth effects in the holdings and trades of money managers. Journal of Finance 60 (December (6)), 2801–2824.

Hussam, R.N., Porter, D., Smith, V., 2008. Thar she blows: can bubbles be rekindled with experienced subjects? American Economic Review 98 (3), 924–937.

Johansen, A., Sornette, D., 1998. Stock market crashes are outliers. European Physical Journal B 1 (February (2)), 141–143.

Johansen, A., Sornette, D., 2001. Large stock market price drawdowns are outliers. Journal of Risk 4 (2), 69–110.

Kaminsky, G., December 1998. Currency and Banking Crises: The Early Warnings of Distress. Social Science Research Network Working Paper Series.

Kaufman, G.G., 2001. Asset Price Bubbles: Implications Monetary and Regulatory Policies (Research in Financial Services: Private and Public Policy). JAI Press.
Kindleberger, C.P., Aliber, R., October 2005. Manias, Panics and Crashes: A History of Financial Crises (Wiley Investment Classics), 5th edition. Wiley.
Körding, K., 2006. Decision theory: what "should" the nervous system do? Science 318, 606–610.
Lei, V., Noussair, C.N., Plott, C.R., 2001. Nonspeculative bubbles in experimental asset markets: lack of common knowledge of rationality vs. actual irrationality. Econometrica 69 (4), 831–859.
Levine, S.S., Zajac, E.J., 2007. The Institutional Nature of Price Bubbles. SSRN eLibrary.
Malkiel, B.G., 1990. A Random Walk Down Wall Street. WW Norton & Company, New York.
McCoy, B.M., Wu, T.T., 1973. The Two-dimensional Ising Model. Harvard University, Cambridge, MA.
Miller, E.M., 1977. Risk, uncertainty, and divergence of opinion. The Journal of Finance 32 (4), 1151–1168.
Montroll, E.W., Badger, W.W., 1974. Introduction to Quantitative Aspects of Social Phenomena. Gordon & Breach Science Pub.
Ofek, E., Richardson, M., 2003. Dotcom mania: the rise and fall of internet stock prices. Journal of Finance 48, 1113–1137.
Orléan, A., 1984. Mimétisme et anticipations rationnelles: une perspective keynesienne. Recherches Economiques de Louvain 52, 45–66.
Orléan, A., 1986. L'auto-référence dans la théorie keynesienne de la spéculation. Cahiers d'Economie Politique, 14–15.
Orléan, A., 1989a. Comportements mimétiques et diversité des opinions sur les marchés financiers. In: Bourguinat, H., Artus, P. (Eds.), Théorie économique et crises des marchés financiers. Economica, Paris, pp. 45–65 (Chapter III).
Orléan, A., 1989b. Mimetic contagion and speculative bubbles. Theory and Decision 27, 63–93.
Orléan, A., 1991. Disorder in the stock market. La Recherche 22, 668–672 (in French).
Orléan, A., 1995. Bayesian interactions and collective dynamics of opinion – herd behavior and mimetic contagion. Journal of Economic Behavior & Organization 28, 257–274.
Plerou, V., Gopikrishnan, P., Gabaix, X., Stanley, H.E., 2002. Quantifying stock–price response to demand fluctuations. Physical Review E 66 (August (2)), 027104+.
Roehner, B.M., Sornette, D., 2000. thermometers of speculative frenzy. European Physical Journal B 16, 729.
Sandroni, A., 1998. Learning, rare events, and recurrent market crashes in frictionless economies without intrinsic uncertainty. Journal of Economic Theory 82 (September (1)), 1–18.
Satinover, J.B., Sornette, D., 2007a. Illusion of control in a Brownian game. Physica A December (386), 339–344.
Satinover, J.B., Sornette, D., 2007b. "Illusion of control" in time-horizon minority and Parrondo games. European Physical Journal B December (60), 369–384.
Satinover, J.B., Sornette, D., 2009. Illusory versus genuine control in agent-based games. European Physical Journal B 67, 357–367.
Sheffrin, H., 2005. A Behavioral Approach to Asset Pricing. Academic Press.
Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? The American Economic Review 71 (3), 421–436.
Shiller, R.J., 2000. Irrational Exuberance. Princeton University Press, New York.
Smith, V.L., Suchanek, G.L., Williams, A.W., 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. Econometrica 56 (5), 1119–1151.
Sornette, D., 2003a. Why Stock Markets Crash. Princeton University Press, Princeton, NJ.
Sornette, D., 2003b. Why Stock Markets Crash. Princeton University Press, Princeton, NJ (Chapter 3).
Sornette, D., 2009. Dragon-kings, black swans and the prediction of crises. International Journal of Terraspace Science and Engineering 2 (1), 1–18.
Sornette, D., Zhou, W.-X., 2004. Evidence of fueling of the 2000 new economy bubble by foreign capital inflow: Implications for the future of the US economy and its stock market. Physica A 332, 412–440.
Taleb, N.N., October 2008. Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets, 2nd edition. Random House.
Topol, R., 1991. Bubbles and volatility of stock prices: effect of mimetic contagion. The Economic Journal 101 (407), 786–800.
Trichet, J.-C., 2010. What Can Central Banks do in a Financial Crisis? Speech by Mr Jean-Claude Trichet, President of the European Central Bank, on the occasion of the Susan Bies Lecture, Kellogg Distinguished Lecture Series , Evanston, Illinois, 27 April 2010, pp. 1–12.
Umpleby, S., 2007. Reflexivity in social systems: the theories of george soros. Systems Research and Behavioral Science 24 (24), 515–522.
Vazquez, A.J., Oliveira, G., Dezso, Z., Goh, K.I., Kondor, I., Barabasi, A.L., 2006. Modeling bursts and heavy tails in human dynamics. Physical Review E 73, 036127.
Veldkamp, L., 2005. Slow boom, sudden crash. Journal of Economic Theory 124 (October (2)), 230–257.
Wyart, M, Bouchaud, J.-P., 2007. Self-referential behaviour, overreaction and conventions in financial markets. Journal of Economic Behavior & Organization 63 (1), 1–24.
Zhou, W.-X, Sornette, D., 2006. Is there a real-estate bubble in the us? Physica A 361, 297–308.
Zhou, W.-X., Sornette, D., 2007. Self-fulfilling ising model of financial markets. European Physical Journal B 55, 175–181.
Zhou, W.-X., Sornette, D., 2008. Analysis of the real estate market in las vegas: bubble, seasonal patterns, and prediction of the csw indexes. Physica A 387, 243–260.

4.3 Finalizing Comments

Due to space limitations, Harras and Sornette (2011) could only cover a limited among of the literature related to financial bubbles and crashes models. Next, further comparisons with selected models presented in Chapter 3 will be performed and additional evidence for the validity of our assumptions and conclusions are provided. The model of Harras and Sornette (2011) will be referred to as the HS-model.

4.3.1 Comparison with selected models

One of the most important distinctions, compared to the models of positive feedback trading (Sec. 3.4) and herding (Sec. 3.5.3), is that the agents in the HS-model are not restrained to one specific strategy to which they have to stick and which controls their behavior. They are given three sources of information, where the impact of the news and social interaction are controlled by the agents themselves. This is in stark contrast to the models of de Long et al. (1990); Daniel et al. (1998); Hong and Stein (1999) or Kirman (1993); Lux (1995); Kaizoji (2000), where the behavior of every type of agent is fixed by the modeler, independent of the strategy's performance.

Agents ability to adapt to different market regime leads to an other important distinction, which is the dynamic, multi-regime, non-cyclic characteristic of the HS-model. With the exception of the model by de Long et al. (1990), which is only a three-period model, the previously listed models can generate price dynamics of many time steps. However, these dynamics either only focus on the absorption of one news event into the price (Daniel et al., 1998; Hong and Stein, 1999), have a periodic nature (Lux, 1995; Kaizoji, 2000), or a quasi-periodic nature, with agents switching symmetrically between an optimistic and pessimistic view on the future (Kirman, 1993). None of these models is able to generate, with one set of parameters, the crossover from an efficient market to the emergence of a bubble and crash, back to the efficient market regime, as the HS-model is able to (cf. Fig. 2 for $C_1 = 4$). Nor are their models able to reproduce the major stylized facts as show in Fig. 1.

Of the previously mentioned models, the HS-model is closest to the model by Hong and Stein (1999) in terms the bubble (or momentum) growing mechanism. For both, it is the news induced price pressure in one direction, which creates a temporary momentum or convergence in opinions, that leads to profitable trend-following or herding. This self-reinforcing behavior lets the price significantly deviating from its fundamental value. Hong et al. assume slow diffusion of news across investors but do specifically state the mechanism behind this diffusion. One way such a percolation of news could be achieved is via cascading investors, enabled by social interaction. Such an interpretation reveals an additional connection between the two models. The major difference between the models, is that for Hong et al., the slowly diffusing news is build by hand into the model, whereas the herding mechanism in the HS-model is only a possibility to the agents, in most situation they will not indulge in such a trading behavior. The bubbling prices are not build in, but emerge out of the myopic adaption of the investors. Consequently, we provide a different ultimate reason behind the bubble (or the medium-term momentum), which only occurs if the news push of a several time-steps into the same direction, not for every release of a news piece of information as for the Hong et al. model[1].

For the discussed bubbles and crashes models based on social imitation and the Ising model (Kirman, 1993; Lux, 1995; Kaizoji, 2000), their bubbles and crashes are symmetric, and often periodic, deviations from the fundamental price, one in the positive the other in the negative direction. Both phenomena have the same time-scale and the same underlying origin. This is not the case for these phenomena in the HS-model, where bubbles are slowly building up, fueled by over-optimistic investors, and the crash happens suddenly and as result of the unsustainable and bloated price.

4.3.2 Additional empirical evidence

As shown in Section 3.5, one of the basic assumptions of the model, social influence of investors and analysts, is well confirmed by empirical data, so the is resulting momentum trading (c.f. Section 3.4). More specifically, the update-mechanism that agents give more weight to peers with higher predictive power in the past is confirmed by Welch (2000). He reports that the influence that one analyst has on others increases with the accuracy of his predictions of future security returns and that this influence is stronger the more recent his revisions were, which is exactly what Eq. (6) describes.

The interpretation that the regime, where agents give a lot of importance to their well-performing neighbors ($C_1 > 3$), corresponds to situations of high uncertainty is not only in agreement with rational models on herding, but also with empirical evidence. Several studies show that the degree of herding decreases with seniority (Clement and Tse, 2005; Hong et al., 2000; Greenwood and Nagel, 2009), which by considering age as a proxy for uncertainty, confirms the interpretation. This is clear as older investors, having already dealt with a larger variety of situations, will feel more secure, than an rookie, experiencing it for the first time. The same is true for uncertain periods in time, as faced with the emergence of a "new economy" in the late 90s, a lot of fund managers and analysts are reported to having increased their tendency for herding (Guedj and Bouchaud, 2004; Greenwood and Nagel, 2009; Griffin et al., 2011).

The range, over which investor compute their correlations (controlled by α), in order to estimated the predictive power of an information source, is crucial for the emergence of bubbles. It is the fact that they conceive a random fluctuation as a genuine signal, which starts the overreaction to a temporary trend in the news. This behavior is also found in experiments, where the participants are found to generalized a pattern based on a very small sample (Andreassen and Kraus, 1990; Offerman and Sonnemans, 2004).

The fact that the paper restricts itself to positive bubbles can be further justified by the empirical evidence of an optimistic bias in analysts earning predictions (cf. Section 3.5.2) and the limitations on short selling (cf. Section. 3.2), making positive deviations from the fundamentals much likelier then persistent undervaluations. An addition piece of evidence comes from Welch (2000), who shows that herding towards the consensus (which has a positive bias) is stronger during bull markets.

[1]This criticism is not entirely fair, as the model of Hong et al. is not aimed at explaining bubbles and crashes but momentum and reversal, a phenomenon with is more common. By assuming that bubbles are a case of amplified momentum, the models become comparable, although not in every point.

Evidence that a string of successive positive news will create a positive momentum in returns, which exceeds the period of positive news, is been reported by Lakonishok et al. (1994) and La Porta et al. (1997). Although these studies document momentum and not bubbles, the evidence points in the same direction then the conclusion of the HS-paper.

Also the mechanism behind the crash can be confirmed by empirical evidence. As the experiments by Smith et al. (1988) find that the period before a crash is characterized by a deceleration where participants realized the overvaluation. At this point a small price drop can destabilize the bubble and lead to a fast correction of the price. Also the end of the housing bubble in the US seems to have the same origin as predicted by the HS-model, that the bubble stops growing once investors liquidity drys up, as it was the reduction of foreign capital influx and the increase of interest rates of the Federal reserve that marked the end of the increasing real Estate prices.

Chapter 5

Noise-induced volatility

5.1 Introductory comments

In the previous chapter, a model of adapting agents under the influence of their social network, global news and private information was studied. As that model was designed to represent investors, interacting on a financial market, the agents were characterized by features special to that environment, such as a certain risk aversion and an amount of cash and assets. More importantly, due to the strong competition in the financial world, they were constantly adapting their strategy to the environment. In this chapter, a far simpler model will be studies, only keeping the same three basic influences, together with the random utility approach. By simplifying the model, it is possible to apply it to a very wide range of systems, generalizing the interpretation of the individual agent from an investor to any entity susceptible to its surrounding and a common driving force. Examples are magnetic spins, neurons in the brain, immune system activity or voting individuals.

The originating idea leading to this paper was the explanation of the phenomenon of excess volatility (first studies by Shiller (1981)) and is treated in Section V.b of the paper. There it is shown that the proposed model is not only able to generate strong price fluctuations, largely exceeding the amplitude of the news (or dividend flows in the scenario of Shiller (1981)), but also that this fluctuation amplification has an endogenous origin leading to a small correlation between news and price variations, a property also seen in financial markets (Parr, 1985; Cutler et al., 1989; Joulin et al., 2008).

The model studies in Section 5.2 has effectively two parameters, the amplitude of the private information, given by the standard deviation D of the noise term, and the coupling strength k, controlling the impact of the surrounding on the single unit. Given that the work was aimed at, and published in, a physics journal, the model is studied with D being the control parameter. The reason for this is that the noise term's standard deviation can be related to the temperature of the system (cf. Eq. (2.40) and it discussion), which is the natural control parameter in such systems. However, very similar results could be obtained by studying the system in function of the coupling strength, k. In that case, the name of the phenomenon would change from *noise-induced volatility* to *interaction-induced volatility*, leaving the entire conclusion of Section 5.2

unchanged.

Without further ado, here is the paper, as it was published in Physical Review E: Statistical, Nonlinear, and Soft Matter Physics.

5.2 The paper

PHYSICAL REVIEW E **85**, 011150 (2012)

Noise-induced volatility of collective dynamics

Georges Harras,[1] Claudio J. Tessone,[2] and Didier Sornette[1]
[1]*Chair of Entrepreneurial Risks, D-MTEC, ETH Zurich, CH-8092 Zurich, Switzerland*
[2]*Chair of Systems Design, D-MTEC, ETH Zurich, CH-8092 Zurich, Switzerland*
(Received 9 August 2011; revised manuscript received 10 October 2011; published 31 January 2012)

Noise-induced volatility refers to a phenomenon of increased level of fluctuations in the collective dynamics of bistable units in the presence of a rapidly varying external signal, and intermediate noise levels. The archetypical signature of this phenomenon is that—beyond the increase in the level of fluctuations—the response of the system becomes uncorrelated with the external driving force, making it different from stochastic resonance. Numerical simulations and an analytical theory of a stochastic dynamical version of the Ising model on regular and random networks demonstrate the ubiquity and robustness of this phenomenon, which is argued to be a possible cause of excess volatility in financial markets, of enhanced effective temperatures in a variety of out-of-equilibrium systems, and of strong selective responses of immune systems of complex biological organisms. Extensive numerical simulations are compared with a mean-field theory for different network topologies.

DOI: 10.1103/PhysRevE.85.011150 PACS number(s): 05.40.−a, 89.65.−s

I. INTRODUCTION

Noise has effects *a priori* unexpected on the organization of complex systems made of interacting elements, as shown by *stochastic resonance* (SR) [1], *coherence resonance* [2], *noise-induced phase transitions* [3], noise-induced transport [4], and its game theoretical version, the *Parrondo's Paradox* [5]. SR occurs in a system when a small applied (subthreshold) periodic signal is amplified by the addition of noise and the maximum of amplification is found for intermediate noise strengths. More generally, SR refers to the situation where noise and nonlinearity combine to increase the strength in the system response. Among others, SR was shown to appear in optical [6] and magnetic systems [7,8], and was thought to be relevant in various fields, ranging from Earth climate [9] and the dynamics of ice ages [10], to neurobiology [11,12] and visual perception [13]. Generally SR is studied in bistable systems, where the amplification of a subthreshold periodic signal is achieved through the synchronization of noise-induced interwell hopping of the dynamic variable and the driving signal. The signal is maximally amplified when the level of noise is such that the Kramers time, which is the intrinsic lifetime associated with the noise-induced transition between the two stable states, equals half of the period of the external forcing. The bistability of the dynamic variable can be given either explicitly [14]—for low dimensional systems—or emerge from the interaction of the many constituents as for the magnetization in the Ising model in the ferromagnetic case [15,16]. However, systems of many interacting constituents may depart from the paradigmatic setting of SR, as there is no interwell hopping for the macroscopic observable, and thus the bistability is only preserved at the microscopic level.

The Ising model, driven by a periodic signal, has been extensively studied in the realm of the kinetic Ising model and dynamical phase transitions [15–20]. When a spatially extended Ising system—for temperatures below the Curie temperature—is forced by a weak periodic influence, the magnetization performs dynamics around its nonzero equilibrium fixed point, resulting in a nonzero time-averaged magnetization. For a given temperature, by increasing either the field strength or the period of the signal, the system becomes able to hop between the two symmetric equilibrium fixed points inducing a zero average magnetization. The nature of this transition, from a nonzero to a zero average magnetization, can be manifold and depends on the control parameter. For weak, subthreshold forcing strength and finite-size systems—where interwell hopping is enabled by the fluctuations—SR is at the origin of the transition. In large enough systems, such that finite-size fluctuations can be neglected, and for intermediate field strengths, they experience a dynamical phase transition through a nucleation process [16]. For stronger forcing, the transition is forced by the exogenous field, with no contribution from endogenous factors.

In this paper, we investigate the behavior of systems composed of many interacting constituents under the influence of a time-varying external forcing. The Ising model framework is used as a generic example of such systems. In contrast to the classical SR studies, where the period of the periodic forcing is of the order of the Kramers time, we are interested in much faster signals. Additionally, aperiodic signals are included to the external forcing, a setup much closer to real world examples, which will yield some surprising differences to the cases involving periodic forcing.

We find, for periodic and aperiodic signals alike, that for intermediate values of the noise intensity, the system dynamics shows a maximum in amplitude [8,21]. Interestingly, the phenomenon of increased amplitude, which consists of an amplification of the signal for a periodic forcing, morphs into an increase of the system-wide fluctuations, uncorrelated with the signal for an aperiodic forcing. We call this phenomenon "*noise-induced volatility*" (NIV).

There are many examples of systems composed of a large number of interacting units that are subjected to a rapidly varying—periodic or aperiodic—common forcing. A first example refers to the empirical observations of strong amplifications of thermal noise into effective renormalized temperatures by quenched heterogeneities in materials [22], in organized flows in liquids [23] and in granular media near jamming [24]. We argue that NIV also provides a conceptual framework to model the immune systems of complex biological organisms, viewed as multistable complexes, which switch their mode of operation under the influence of noisy perturbations by pathogens and other stress factors [25–27].

Another important application of the proposed mechanism of volatility amplification can be found in financial markets. The phenomenon of "excess volatility" [28] constitutes one of the major unsolved puzzles in financial economics and refers to the ubiquitous observation that financial prices fluctuate with much larger amplitudes than they should if they obeyed the fundamental valuation formula, linking the share price of a company to its expected future dividends and discount factors [29]. The model described below can be applied to represent a market of interacting investors, where the external forcing represents the news (i.e., the publicly available information about the traded assets) that investors use to update their estimates of the asset's fair value. In addition to the phenomenon of the increased volatility compared to the news amplitude, our framework allows us to address two other well-known phenomena of financial markets: namely, the fact that the news is a poor predictor of future price changes [30] and the phenomenon of clustered volatility, quantified by the slowly decaying temporal dependence of volatility [31].

We document the phenomenon of noise-induced volatility by numerical and theoretical calculations on a stochastic dynamical version of the Ising model on fully connected, regular as well as random networks, in the presence of rapidly varying periodic and aperiodic signal. NIV also constitutes a new indicator for an approaching phase transition [32].

This paper is organized as follows. In the following section, we introduce the model studied and the measures chosen to quantify the phenomenon studied. In Sec. III, we revisit the case where the system is driven by a periodic forcing, focusing on the case of fast signals, by means of Monte Carlo simulations and by means of an analytical approach. In Sec. IV, we present the main contribution of the paper: the study of the system driven by an aperiodic forcing. In Sec. V, we go beyond the fully connected case, focusing on different network topologies and on a paradigmatic example of this phenomenon: the excess volatility in financial markets. Finally, Sec. VI presents a discussion and conclusions of the obtained results.

II. MODEL DESCRIPTION AND DIAGNOSTIC VARIABLES

Consider a system composed of N interacting units that can be in one of two states: $s = \pm 1$. The units are updated sequentially, randomly chosen at each unit micro-time $\delta = 1/N$ (i.e., N updates are equivalent to one time unit at the macroscopic level, that is, one Monte Carlo step [33–35]). The update of the state s_i of a given unit i from t to $t + \delta$ is given by

$$s_i(t+\delta) = \mathrm{sgn}\left(f(t) + \xi_i(t) + K(t) \sum_{j=1}^{N} \omega_{ij}\, s_j(t) \right). \qquad (1)$$

The value $s_i(t+\delta)$ is determined by three competing contributions: (i) a common external dynamic forcing term $f(t)$ (force, pathogens abundance, news); (ii) an annealed unit-specific term $\xi_i(t)$ that we will call *noise* (thermal fluctuations or threshold, intrinsic susceptibility of a unit immune system compartment, investor idiosyncratic opinion, or private information); (iii) an interaction term between units controlled by the amplitude $K(t)$ (elastic coupling, feedback loops between immune system elements, social impact).

The system's behavior will be investigated under the influence of two different types of external signals. To relate to the existing literature, we will use a smooth periodic signal, $f_p(t) = A \sin(\omega t)$, with period $2\pi/\omega$ and strength A; this implies that, when averaged over time, the standard deviation of the signal is $\sigma_{f_p} = A/\sqrt{2}$. Along this paper, we denote by *signal amplitude* the standard deviation of the signal σ_f. As a periodic signal is a rather stylized and artificial setup, we will, in a later section, also analyze the response of such a system to a stochastic process. The simplest choice of a stochastic process with tunable characteristic time scale is the Ornstein-Uhlenbeck (OU) process, which has exponentially decaying memory and is defined by $df_{ap} = -\theta f_{ap} dt + A\, dW_t$, with 0 mean, strength A, inverse time scale $\theta > 0$ and W_t is a Wiener process with normalized variance and zero mean. The asymptotic solution of the OU process is

$$f_{ap}(t) = A \int_{-\infty}^{t} e^{-\theta(t-\tau)}\, dW_\tau, \qquad (2)$$

which gives a signal amplitude $\sigma_{f_{ap}} = A/\sqrt{2\theta}$.

The noise term $\xi_i(t)$ of each unit in Eq. (1) follows an independent stochastic process, whose values are, at every micro-time-step, drawn from the cumulative distribution function $G(0,D)$, with zero mean ($\langle \xi_i(t) \rangle = 0$) and variance D^2. Thus, $\langle \xi_i(t)\, \xi_j(t+n\delta) \rangle = D^2 \delta_n \delta_{ij}$. If $f(t) = 0$ and $G(0,d)$ corresponds to a logistic distribution, the dynamical rule of Eq. (1) is equivalent to the kinetic Ising model with Glauber dynamics (cf. appendix) where D^2 is related to the temperature.

In the interaction term in Eq. (1), the matrix of weights ω_{ij} defines the network connectivity between units, both in topology and relative strength. We assume that the interactions between units are governed by connections that evolve much slower than the dynamics of the whole system. This amounts to considering a static network with fixed normalized weights $\sum_j \omega_{ij} = 1$. The effective coupling strength is given by $K(t)$, which may depend on time to reflect global softening-hardening in rupture processes, evolving physiological states of immune systems, and changes of social cohesiveness and/or social influence in financial markets.

The macroscopic dynamics of the system is captured by the instantaneous "magnetization":

$$m(t) = \frac{1}{N} \sum_i s_i(t), \qquad (3)$$

which fluctuates around its time-average Q, which is computed as

$$Q = \langle m(t) \rangle_t = \frac{1}{T} \int_0^T m(t) dt, \qquad (4)$$

where T is the duration of the simulation. We study the normalized standard deviation:

$$\tilde{\sigma} = \frac{\sigma_m}{\sigma_f} = \frac{\sqrt{\langle [m(t) - Q]^2 \rangle_t}}{\sqrt{\langle f(t)^2 \rangle_t}}, \qquad (5)$$

of $m(t)$, describing the "volatility" of the system dynamics scaled by the signal amplitude, σ_f. As the response of the system to an external influence is not instantaneous, the time-lagged correlation between the input signal and the magnetization, defined by

$$\rho(\tau) = \frac{\langle [m(t+\tau) - Q] f(t) \rangle_t}{\sigma_m \sigma_f}, \tag{6}$$

provides an additional insight on the level of synchronization between the external influence and the overall system dynamics at a lag of τ. The lag where the correlation is maximal will be called optimal lag, $\tau^* = \max_\tau \rho(\tau)$.

In the case of the periodic signal, a common measure in stochastic resonance research is the spectral amplification factor (SAF) [36],

$$R = \frac{S_\omega[m(t)]}{S_\omega[f_p(t)]} = \frac{S_\omega[m(t)]}{\sigma_f^2}, \tag{7}$$

which is the ratio of the power spectrum density of the magnetization $S_\omega[m(t)]$ over the power spectrum density of the driving signal, $S_\omega[f_p(t)] = \sigma_f^2 = A^2/4$, both at the driving frequency ω.

III. PERIODIC SIGNAL

A. Simulations results

First, we consider an homogeneous, complete, network ($\omega_{ij} = 1/(N-1)$) and a constant coupling strength $K(t) = k = 1$. The results reported below are not significantly different for random graphs with large average connectivity or when the connections allow for an unbiased statistical sampling within the population. As previously said, we set G to be a Gaussian distribution with standard deviation D, and zero mean. Even though the system loses its equivalence to the kinetic Ising model with Glauber dynamics, all the qualitative properties of the system remain unchanged. Without external forcing ($A = 0$), the system experiences, as for the equilibrium Ising model, a continuous phase transition at $D_c \simeq 0.80k$, separating the ordered phase with two stable fixed points at $\pm Q(D)$, from the disordered phase, with a single stable fixed point at $Q(D) = 0$. For the equilibrium case ($A = 0$), the dependence of $Q(D)$ as a function of D is shown by the continuous line in Fig. 1(b).

In Fig. 1(a), we plot the spectral amplification factor as a function of noise strength for signals with different periods. The symbols are obtained by simulations of the model with 10^6 units. We observe that, even for relatively small periods, an increase of amplification exceeding one order of magnitude is achieved for a broad range of intermediate value of noise, the hallmark of stochastic resonance. Figure 1(b) shows the average magnetization, $Q(D)$, which is the usual order parameter in the kinetic Ising model studies, for the same signals as in Fig. 1(a). In Sec. III B, we will develop a theory which shows that the global dynamics can be assimilated to a motion within a potential that exhibits a transition from a monostable to a bistable regime. For fast signals, $Q(D)$ has a value close to the equilibrium fixed point (continuous line), suggesting that the dynamics of $m(t)$ can be well described as fluctuations around this stable (or metastable) point, i.e., $m(t)$ performs intrawell dynamics. For slower signals, larger fluctuations around the equilibrium fixed point are observed,

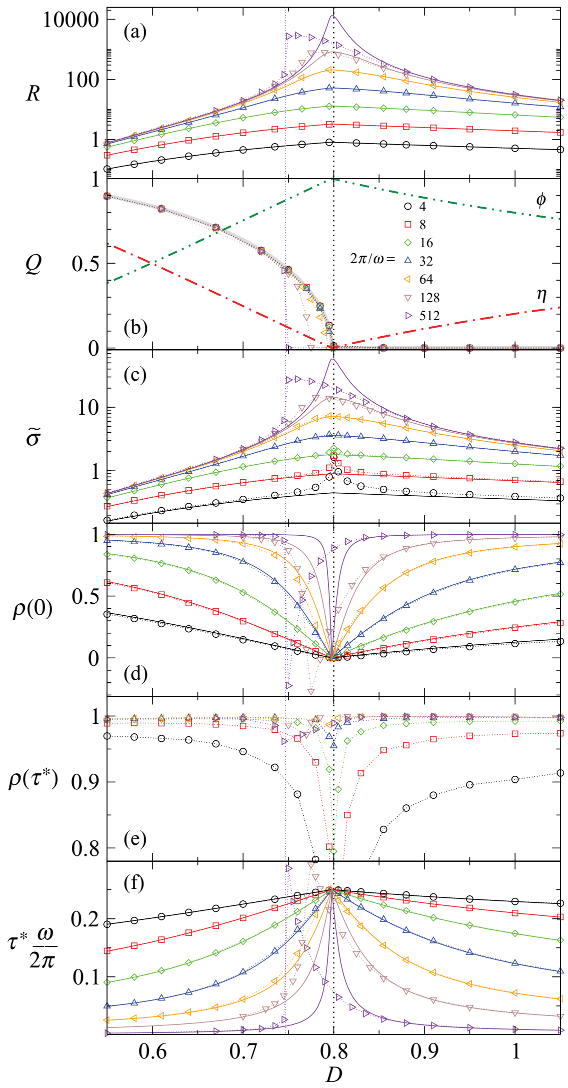

FIG. 1. (Color online) (For periodic signals) Different measures as a function of the standard deviation of the noise in Eq. (1), D, for weak periodic signals with amplitude $\sigma_{f_p} = 0.01$. The period of the signal is in the legend of (b). Symbols are obtained for simulations with a system size $N = 10^6$ and the lines are the result of the linear approximation presented in the main text. The vertical thin dotted line indicates D_T for $\omega = 2\pi/512$, where the transition from intra- to interwell dynamics occurs. The vertical thick dotted line indicates the critical value of noise D_c, where the phase transition takes place. For $D \in]D_T, D_c[$, $m(t)$ performs interwell dynamics. (a) Spectral amplification factor R, as defined by Eq. (7). (b) Average magnetization Q. For $D < D_c$, $Q = 0$ indicates that the system performs interwell dynamics. (c) The normalized standard deviation, $\tilde{\sigma}_p$. (d) Instantaneous correlation between m and f, $\rho(0)$, defined by Eq. (6). (e) Correlation between m and f at the optimal lag. (f) Optimal lag, τ^*, normalized by the period of the driving force.

as $Q(D)$ vanishes already for $D < D_c$, indicating symmetric oscillations around $m(t) = 0$, i.e., $m(t)$ performs interwell dynamics. By $D_T(A,\omega,N)$, we denote the threshold value of the noise strength at which the potential barrier between the two equilibrium fixed points become small enough, such that the system performs interwell hopping and thus $Q(D)$ goes to zero. The noise strength threshold D_T approaches D_c as either ω or N are increased or A is decreased. In the opposite limits, it will tend to zero.

Independent on the driving frequency, the maximum in the amplification is always found at D_T. For fast signals where $D_T \sim D_c$, this maximum is observed at the equilibrium phase transition. This happens in the presence of two competing phenomena near the equilibrium phase transition: on the one hand, a divergence in the susceptibility, making the system very sensitive to small changes in the external influences; on the other, critical slowing down, which inhibits the reaction of the system. For slow signals, where $D_T < D_c$, together with a more abrupt vanishing of Q, a pronounced jump in the spectral amplification factor R, defined in Eq. (7), is observed at $D = D_T$, where the response of the system is greatly increased by the transition from intra- to interwell dynamics. For $D_T < D < D_c$, the amplification decreases with D, as the position of the minimum ($\pm m_0(D)$) approaches 0 for D approaching D_c from the left.

Figure 1(d) shows the dependence of the correlation at zero lag on the noise strength. A minimum of instantaneous correlation is observed at the same values of D, where the maximum in R occurs. This result confirms the existence of a double peak of the non-normalized instantaneous covariance, as was found by Leung *et al.* [7,8].

The effect of the signal frequency on the system behavior is shown in Fig. 2, where we plot the amplification R as a function of the period of the signal for different values of D. For

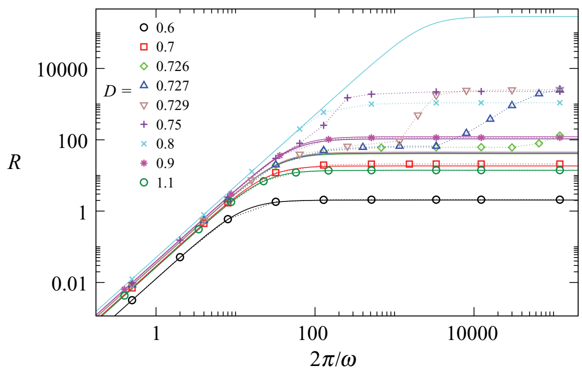

FIG. 2. (Color online) For periodic signals with amplitude $\sigma_{f_p} = 0.02/\sqrt{2}$ and different values for the standard deviation of the noise D specified in the legend. The spectral amplification factor R as a function of the period, $2\pi/\omega$. Symbols are obtained for simulations of a system with $N = 9 \times 10^4$ and the lines represent the linear approximation given by Eq. (15). Three regimes of R can be identified as a function of the period: (1) increasing amplification with increasing period, (2) plateau with intrawell dynamics, (3) for $D \lesssim D_c$, stark increase of amplification due to interwell dynamics. Similar results are found for larger system sizes, where the third regime appears for larger periods as the finite-size fluctuations are reduced.

$D > D_c$, where the macroscopic system dynamics is described by a monostable potential, the dependence is composed of two regimes. The first regime is where the amplitude of the oscillations of $m(t)$ increases with the period as $m(t)$ is pushed for longer durations into one direction, allowing for greater deviations from the origin. For larger signal periods, R reaches a plateau, which constitutes the second regime, where the diffusive motion of $m(t)$ is confined by the potential. The same behavior is observed for $D \ll D_c$, where the potential barrier between the two minima cannot be overcome by the system, restricting the dynamics to intrawell motions. Finally, for $D \lesssim D_c$, the dependence shows a transition into a third regime. If the potential barrier is not too high compared with the noise intensity and the finite-size fluctuations, the system is able to perform interwell dynamics for large enough periods of the external driving. These interwell dynamics are observed as a second rapid increase in R. The period at which this transition happens is the double of the Kramers time.

B. Analytical approach

In order to understand these results, we now develop a mean-field theory, which becomes exact in the thermodynamic limit and for weak signal amplitudes. As our system is composed of many interconnected units, we can rewrite Eq. (1) by replacing the interaction term by the global instantaneous magnetization and by explicitly writing down Eq. (1) in the form of

$$s_i(t+\delta) = \begin{cases} +1 & \text{if} \quad \xi_i(t) \geqslant -k\,m(t) - f(t) \\ -1 & \text{if} \quad \xi_i(t) < -k\,m(t) - f(t) \end{cases}.$$

Averaging over multiple noise realizations, the expected value for the state of the ith unit, at time $t+\delta$ is thus given by

$$\begin{aligned} \langle s_i(t+\delta)\rangle_\xi &= 1 - G(-k\,m(t) - f(t)) - G(-k\,m(t)) - f(t)) \\ &= 1 - 2G(-k\,m(t) - f(t)), \end{aligned} \tag{8}$$

where $G(\theta)$ is the cumulative distribution function of the noise term $\xi_i(t)$, i.e., $G(\theta) = \int_{-\infty}^{\theta} d\theta' g(\theta')$. Summing over all the units and given that only spin i is updated over the micro-time-step δ, we get that the updated instantaneous magnetization is exactly

$$m(t+\delta) = m(t) + \frac{1}{N}[s_i(t+\delta) - s_i(t)]. \tag{9}$$

By averaging over the complete population and identifying $1/N = \delta$ as dt in the thermodynamic limit, Eq. (9) transforms into a continuous process, which reads

$$\left\langle \frac{dm(t)}{dt} \right\rangle = \langle s_i(t+dt)\rangle - m(t). \tag{10}$$

With $\dot{m}(t) = \langle dm(t)/dt\rangle$ and substituting Eq. (8) into Eq. (10), we get

$$\dot{m}(t) = -m(t) + 1 - 2G(-k\,m(t) - f(t)), \tag{11}$$

which constitutes a general closed form evolution equation for the magnetization of the system.

From the point of view of the mean-field limit, the noise $\xi_i(t)$ can be either quenched or annealed, as the complete noise is condensed into the last term in Eq. (11). For $A = 0$ (no external driving), the stationary solution of Eq. (11) gives the dependence of the equilibrium fixed point $m_0(D)$ as the solution of the implicit equation,

$$m_0(D) = 1 - 2G(-k\, m_0(D)). \tag{12}$$

This solution exhibits a supercritical pitchfork bifurcation as a function of D, as expected for an Ising-like system, which is displayed by the continuous line in Fig. 1(b). The critical parameter is found equal to $D_c = k\sqrt{2/\pi}$, when $\xi_i(t)$ is drawn from a Gaussian distribution.

The emphasis of this paper is on the system's reaction to fast, subthreshold ($A \ll 1$) signals, so that interwell dynamics can be neglected. Thus, a perturbation expansion $m(t) = m_0 + m_1(t)$ up to first order yields

$$\frac{d}{dt}m_1(t) = -\eta(D)\, m_1(t) + \phi(D)\, f(t) + O\big(m_1^2\big), \tag{13}$$

where $\phi(D) \equiv 2g(-km_0)$, $\eta(D) \equiv 1 - 2kg(-km_0)$, and $g = dG/d\xi$. The dependence of ϕ and η as a function of the noise strength is displayed by the dash-dotted lines in Fig. 1(b). Based on Eq. (13), $\phi(D)$ can be interpreted as the attenuation of the signal by the noise in the individual constituents of the system as $\phi(D) \leqslant 1/k$ for any D. The value of $\phi(D)$ weights the impact of the external forcing on the global dynamics. The parameter $\eta(D)$ can be understood as the strength of the restoring force that tends to bring $m(t)$ back to its equilibrium value m_0, after being driven away by the influence of $f(t)$. The larger η, the closer the dynamics of $m(t)$ will be to m_0 and the shorter will be the memory of $m(t)$. The value of $\eta(D)$ controls the contribution of the endogenous part of the dynamics. In the particular case where $f(t)$ is constant, $m_1(t)$ approaches the fixed point $f\,\phi(D)/\eta(D)$. Since $\phi(D)$ remains finite when D passes through D_c, it is the vanishing of $\eta(D)$ at $D = D_c$ and its smallness in the vicinity of D_c that is at the origin of the amplified volatility. Based on Eq. (13), we can now compute the approximate value of the different measures for the external signals and compare them to the simulations of the actual system.

For the periodic forcing, the dynamics of the magnetization yields

$$m_p(t) = \frac{A\,\phi}{\eta^2+\omega^2}[-\omega\cos(\omega t) + \eta\sin(\omega t)] + m_0. \tag{14}$$

Together with Eq. (7), this gives a spectral amplification factor equal to

$$R_p = \frac{4}{A^2}\left(\frac{A\,\phi}{\eta^2+\omega^2}\right)^2\frac{\omega^2+\eta^2}{4} = \frac{\phi^2}{\eta^2+\omega^2}. \tag{15}$$

Figure 1(a) shows that, for fast signals (where $D_T \simeq D_c$), the value of R obtained from this approximation matches well with the simulation results. Deviations from the approximation appear for slower signals when D_T does not coincide with D_c and nonlinear effects cannot be neglected anymore.

As can be seen in Fig. 1(a), the spectral amplification factor can reach values above 100, showing that this system, even without considering interwell dynamics, is able to show remarkable reactions to a weak forcing. Two distinct amplification mechanisms of subthreshold periodic signals can be identified by comparing the simulation with approximation results. The first mechanism, being present for finite and infinite systems, is the increase of the output amplitude by the decrease of the value of η: By reducing the restoring force of $m(t)$, such that it can be further displaced from m_0, the oscillation amplitudes are increased. The second mechanism is the amplification through interwell jumps, which is only present in finite systems, as a subthreshold driving force cannot overcome the potential barrier without the existence of a source of fluctuations, like finite-size effects. Note that in the thermodynamical limit, where the approximation is valid for any frequency, at D_c where $\eta(D_c) = 0$, it follows from Eq. (15) that the fluctuations of $m(t)$ for nonzero frequencies will always be finite.

In addition to the spectral amplification factor R, we also measure the normalized variance of $m(t)$, $\tilde{\sigma}$, which measures the volatility of the dynamics, independent of the exact shape of the power spectrum. This measure is convenient as it can be used for comparison with the aperiodic signal, for which R is not defined. From Eq. (14), the normalized variance of $r(t)$ is

$$\tilde{\sigma_p}^2 = \frac{2}{A^2}\langle m_p(t)^2\rangle_t = \frac{2}{A^2}\frac{\omega}{2\pi}\int_0^{\frac{2\pi}{\omega}} m_p(t)^2 dt = \frac{\phi^2}{\eta^2+\omega^2}. \tag{16}$$

The equivalence between Eqs. (15) and (16) is due to the use of a linear response approximation in the macroscopic dynamics of the systems, neglecting the response at higher order harmonics of the driving signal. As a consequence, the approximation of the spectral amplification factor R_p is better fitted by the simulations than $\tilde{\sigma}_p$.

From Fig. 1(c), we see that the mean-field approximation matches well the values of $\tilde{\sigma}$ obtained by simulations for intermediate signal periods. For large periods, the interwell dynamics destroys the match, and for fast signals (small periods), the finite-size fluctuations overshadow the fluctuations induced by the signals.

The correlation between $m_p(t)$ and $f_p(t)$ is given by

$$\rho_p(\tau) = \frac{\eta\cos(\omega\tau) + \omega\sin(\omega\tau)}{\sqrt{\eta^2+\omega^2}}, \tag{17}$$

and, for the optimal lag, we obtain

$$\tau_p^* = \frac{\arctan(\frac{\omega}{\eta})}{\omega}, \tag{18}$$

which follows directly from Eq. (14). The correlation at zero lag is shown in Fig. 1(d). As for the results of R, the simulation results are well captured by the mean-field approximation for high frequencies and deviate due to interwell dynamics for lower frequencies. As was observed in [8], a dip in correlation is observed for intermediate values of the noise amplitude. This dip occurs at D_T, concomitant with the maximum in the amplification measured by R. This apparent contradiction can be explained by the results shown in Fig. 1(f), which plots the optimal lag between $m(t)$ and $f(t)$ normalized by the period of the signal. At D_T, $m(t)$ and $f(t)$ are maximally lagged, reducing the correlation at lag zero. On the other hand, the correlation at the optimal lag has a value close to one, explaining the high amplification of the signal. This behavior can also be found in the approximation, although there the maximum of amplification and minimum of instantaneous correlation is found at D_c as our approximation neglects interwell dynamics. From Eqs. (17) and (18) it follows that,

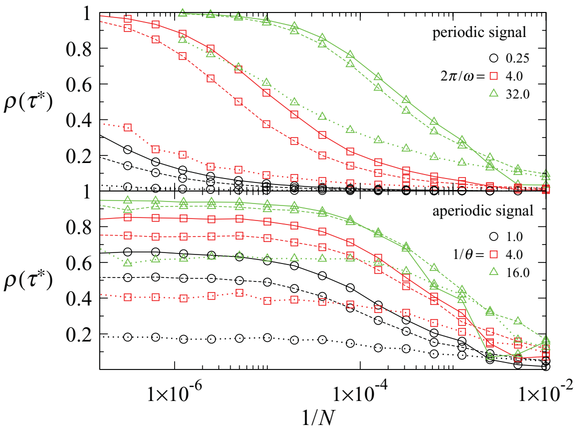

FIG. 3. (Color online) Size dependence of the maximal correlation between the magnetization and the signal with $\sigma_f = 0.01$ and different time scales for several noise intensities: $D = 0.7$ (solid lines), $D = 0.8$ (dotted lines), $D = 0.9$ (dashed lines). Characteristic time scales of the signal are specified in the legends. (top) Periodic signals, where $\rho(\tau^*)$ converges to 1 for any value of D in the thermodynamic limit. The rate of convergence increases with the signal period. (bottom) In case of aperiodic signals, the value of $\rho(\tau^*)$ does not converge to 1 in the thermodynamic limit and depends on D and θ. At $D = D_c$, the maximal correlation $\rho(\tau^*)$ is significantly reduced compared to $D \neq D_c$.

in the thermodynamic limit, there exists an optimal lag τ_p^*, for which perfect correlation is achieved for any frequency and any noise strength, i.e., $\rho_p(\tau_p^*) = 1$.

However, perfect correlation is not achieved for any frequency in finite systems as shown by the dependence of $\rho(\tau^*)$ in Fig. 1(e), where the deterioration of the correlation with increasing driving frequencies is observed. The origin of this effect lies in the finite-size fluctuations, which vanish in the thermodynamic limit, as Fig. 3 (top) shows. $\rho(\tau^*)$ converges to 1 for infinite systems, at a rate of convergence depending on D and the driving period.

IV. APERIODIC SIGNAL

We now turn our attention to the case where the common forcing is aperiodic. In this section, we will consider an external force, which is described by the OU process introduced in Eq. (2). Figure 4 illustrates the typical dynamic behaviors of $m(t)$ for different values of noise strengths D for a single realization of the driving force $f_p(t)$. For $D \geqslant D_c$, the magnetization fluctuates around $m_o(D) = 0$, with increasing amplitudes as D approaches D_c. The fluctuation amplitude deceases again when $m(t)$ performs intrawell dynamics for $D < D_c$, where $Q(D) \neq 0$.

We will compute the same observables as for the periodic signal and investigate the differences. The formal solution of the linearized version of the dynamics, Eq. (13), is given by

$$m_1(t) = A\,\phi \int_{-\infty}^{t} e^{-\eta(t-\tau)} \int_{-\infty}^{\tau} e^{-\theta(\tau-\tau')}\, dW_{\tau'}\, d\tau, \qquad (19)$$

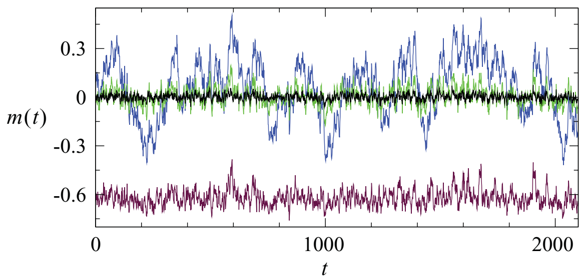

FIG. 4. (Color online) For aperiodic signals: time evolution of the magnetization $m(t)$ for different noise intensities D obtained with the same realization of the driving force $f(t)$ and $\xi_i(t)$. $N = 10^4$, $\sigma_{ap} = 0.04$, $\theta = 1.0$, $k = 1.0$, $D = 2.0, 1.0, 0.8$ (smaller to larger amplitude of $m(t)$'s fluctuations), and $D = 0.7$ (bottom curve fluctuating around $m(t) = -0.6$).

which describes the dynamics around m_0. The normalized variance of $m_1(t)$ described by Eq. (19) is now given by

$$\tilde{\sigma}_{ap}^2 = \frac{\phi^2}{\eta\,(\theta+\eta)}. \qquad (20)$$

In Fig. 5(a), we compare the normalized variance obtained by means of numerical simulations with this theoretical result. We find a very good agreement between the two for fast signals, with the same deficits due to inter-well dynamics as for the periodic case with slower signals. However, by comparing $\tilde{\sigma}_{ap}$ for the aperiodic signal with the periodic case $\tilde{\sigma}_p$, we observe a major difference. Whereas for the periodic case, the volatility of the dynamics shows a finite maximum value at D_c, the volatility diverges if the system is driven by an aperiodic signal as $\eta(D_c) = 0$. This divergence of the normalized volatility $\tilde{\sigma}_{ap}$ is not to be understood as an explosion of the dynamics, as $m(t)$ cannot exceed $[-1, +1]$. It reflects the immensely amplified reaction to a weak external forcing, consistent with the diverging susceptibility in equilibrium phase transitions. The fact that this divergence is absent for a periodic forcing stems from the discreteness of the power spectrum of the input signal. It is worth mentioning that the good match between the analytical approach—which becomes exact in the thermodynamic limit—and the numerical simulations shows that the phenomenon is not due to finite-size fluctuations, but is an emergent property of the system.

The correlation between the forcing $f_{ap}(t)$ and the magnetization, $m(t+\tau)$, is given by

$$\rho_{ap}(\tau) = \frac{\sqrt{\eta^2+\theta\eta}}{\eta^2-\theta^2}\left[(\eta+\theta)\,e^{-\theta\tau} - 2\theta\,e^{-\eta\tau}\right] \qquad (21)$$

and is shown in Fig. 5(b) for zero lag, together with the simulation results, which are found in good agreement. The optimal lag for which $\rho_{ap}(\tau)$ is maximum occurs at

$$\tau_{ap}^* = \frac{\ln\left(\frac{\eta+\theta}{2\eta}\right)}{\theta-\eta}, \qquad (22)$$

yielding a maximal correlation of

$$\rho_{ap}(\tau_{ap}^*) = 2^{\frac{\theta}{\theta-\eta}}\left(\frac{\eta}{\theta+\eta}\right)^{\frac{\eta+\theta}{2(\theta-\eta)}}. \qquad (23)$$

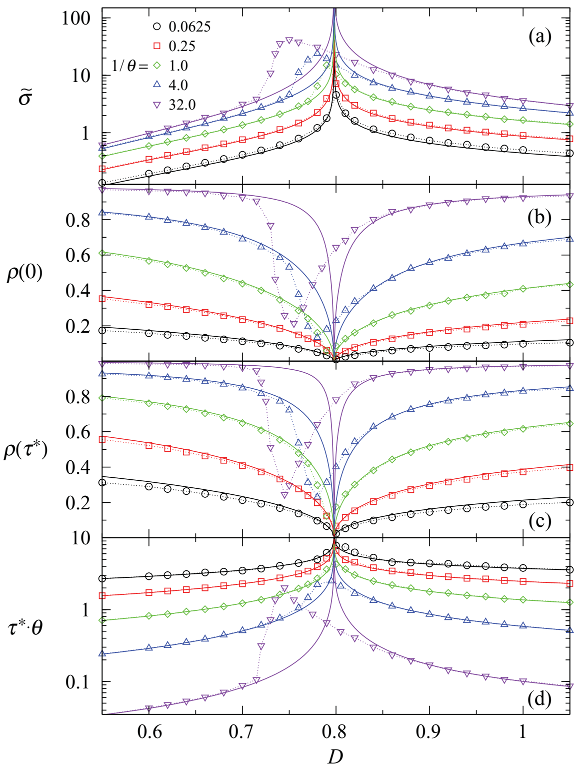

FIG. 5. (Color online) For aperiodic signals: different measures as a function of the standard deviation of the noise in Eq. (1), D, for weak aperiodic signals with amplitude $\sigma_{f_p} = 0.01$. The inverse of the signal's time scale θ is given in the legend. Symbols are obtained for simulations with a system size $N = 10^6$ and the lines are the result of the linear approximation presented in the main text. (a) Normalized standard deviation $\tilde{\sigma}$, measuring the volatility amplification. (b) Instantaneous correlation between m and f, $\rho(0)$, defined by Eq. (6). (c) Correlation between m and f at the optimal lag. In contrast to the periodic signal, a system driven by an aperiodic signal is not able to follow the signal, even at the optimal lag, which is well described by the mean-field approximation. (d) Optimal lag τ^* normalized by the time scale of the signal.

Here, we find the second major difference between the periodic and aperiodic driving. For the periodic signal, it is always possible to find a lag at which the correlation between the forcing and the system's response is perfect. For the aperiodic signal (see Figs. 5(c) and 3 (bottom)), on the other hand, even in the thermodynamic limit, the dynamics of the system can be almost unrelated to the forcing. Perfect correlation is only reachable for very slow signals, i.e., $\lim_{\theta\to\infty} \rho_{ap}(\tau^*_{ap}) = 1$. Given that the forcing—an OU process—has a continuous power spectrum, and that the response of the system is frequency dependent, the spectrum of the macroscopic dynamics is distorted when compared to the one of the forcing, which has the effect of decreasing the correlation. Indeed, the system is only able to follow the part of the signal spectrum with frequencies lower than $\eta(D)$, which describes the rate at which the system can effectively react to external stimuli. As with decreasing θ, the contribution of lower frequencies in the signal's spectrum is higher, the correlation for fixed D increases with decreasing θ.

For $D \approx D_c$, the volatility amplifies many times that of the driving signal $f(t)$. Concomitantly, ρ vanishes for every value of the lag τ, indicating that the volatility of the system is generated by an internal collective behavior. It is important to note that, even though the system dynamics are endogenously generated, they are initiated by an exogenous driving of the system. This is further confirmed by the good agreement between the approximation and simulations for fast signals. It is the shadow of the diverging susceptibility together with the vanishing rate of the reaction of the equilibrium model at D_c, which is responsible for the observed NIV phenomenon.

V. EXTENSIONS OF THE PHENOMENON STUDIED

A. Different networks

To show that the NIV phenomenon, characterized by the increase of volatility and decrease of correlation to the aperiodic forcing, is robust with respect to the structure of the network, Fig. 6 shows the normalized volatility $\tilde{\sigma}$ and the maximum in correlation $\rho(\tau^*)$ as a function of D (the standard deviation of the noise term in Eq. (1)) for different networks. We consider a two-dimensional regular grid with Moore neighborhood and random small-world connections with varying concentration p_w. Changing p_w from 0 to 1 interpolates between the regular two-dimensional (2D) lattice and the completely random network. For each p_w, the peak in volatility is still concomitant with the vanishing of ρ at the threshold value $D_T(p_w)$. The noise intensity threshold $D_T(p_w)$ is increasing in p_w, as larger global interconnection enhances the cooperative organization, and larger noise is needed to destroy the ferromagnetic state.

For one-dimensional (1D) lattices, the NIV phenomenon is still present, with a minimum in the cross correlation and a maximum in the volatility at D_T. As is well known, the

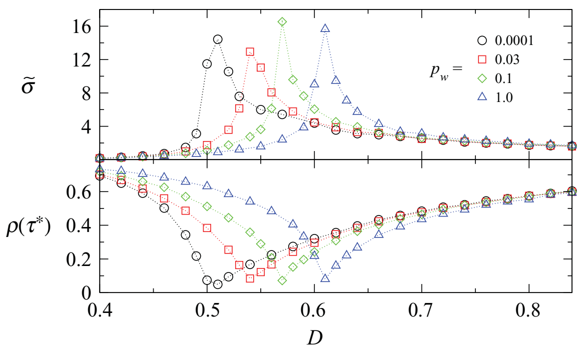

FIG. 6. (Color online) For an aperiodic signal: (top) the normalized volatility $\tilde{\sigma}$ and (bottom) the correlation at optimal lag $\rho(\tau^*)$ as a function of the standard deviation of the noise D for different small-world random connection concentrations p_w of a two-dimensional regular grid with Moore neighborhood. For $p_w = 0$, the network is a 2D regular grid with Moore neighborhood. For $p_w = 1.0$, the network is a random graph with an average degree, $d = 4$. The other system parameters are $N = 10^6$, $\sigma_f = 0.01$, $\theta = 1$, $k = 1$.

one-dimensional Ising model undergoes a first-order phase transition at zero temperature, and D_T converges to 0 in the thermodynamic limit. Notwithstanding the absence of a continuous phase transition, the susceptibility and relaxation time still diverge, exhibiting an essential singularity at zero temperature [37,38], which explains the survival of the NIV phenomenon in 1D.

B. Excess volatility in financial markets

By the definition of our model given by Eq. (1), it is clear that it is also interpretable as a model of opinion dynamics, where $s_i(t)$ is the opinion of agent i at time t, in line with the established literature on discrete choice [39]. The external forcing $f(t)$ can be seen as the flux of news, which is common to all agents, the noise $\epsilon_i(t)$ contains the agents' private information and the coupling term represents the social interaction between agents. The dynamics of the global opinion is then given by $m(t)$.

When applied to the social system of financial markets, the agents are investors and $s_i(t)$ corresponds to their opinion on whether the asset is over- or underpriced and hence to their willingness to buy (+1) or to sell (−1). The global demand is the given by $m(t)$, which impacts on the price as

$$\log[p(t+1)] = \log[p(t)] + \frac{m(t+1)}{\lambda}. \quad (24)$$

Here λ represents the liquidity depth of the market, which is assumed constant and $m(t)/\lambda$ is the financial return $r(t)$ from period t to period $t+1$. This equation expresses a linear market impact of the demands, which is a common hypothesis in stylized models of financial markets [40,41]. The results below do not change qualitatively for more general nonlinear impact functions [42].

To apply our model to the financial markets, we use the coupling strength k instead of D as the control parameter. Rather than assuming a fixed coupling strength for investors, we propose that the impact of colleagues' opinions on a given investor may be slowly varying with time. This effect reflects the fact that, in times of greater uncertainty, investors tend to be more influenceable by their surrounding [43]. There are many varying sources of uncertainty that impact financial markets, including the economic and geopolitical climate and past stock market performance. In the spirit of Ref. [44], all these factors are embodied into the notion that $K(t)$ undergoes a slow random walk with i.i.d. increments $K(t+\delta t) - K(t) \sim N(0,\sigma_k)$, which is confined in the interval $[k-\Delta k; k+\Delta k]$. This later constraint ensures that social imitation remains bounded. We could have used an Ornstein-Uhlenbeck process or any other such confining dynamics, without changing the crucial results presented below. More complex models of sophisticated investors involve the strategic adaptations of the traders' propensity to imitate to the reliability of their colleagues in recent outcomes [45–47].

By the mechanism of sweeping of the coupling strength $K(t)$ close to the critical coupling strength k_c (for fixed noise strength D) [48,49], we expect and find a transient burst of volatility in response to the aperiodic driving force $f(t)$ with constant amplitude and time scale. Figure 7 shows a typical realization, where the return $r(t)$ exhibits transient bursts,

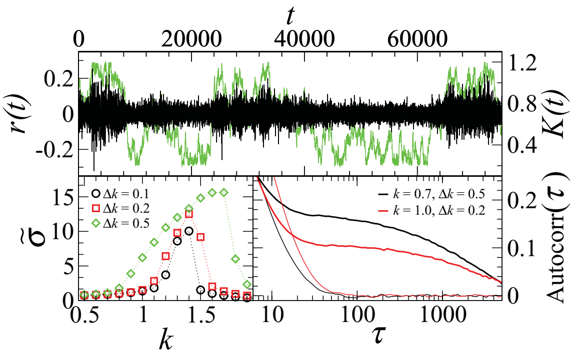

FIG. 7. (Color online) (Upper panel) Sample dynamics $r(t)$ (black bursty line) when the coupling strength $K(t)$ of the interactions between agents undergoes a confined random walk (green) in $[k-\Delta k; k+\Delta k]$ with $\Delta k = 0.5$ and step size $\sigma_k = \Delta k/\sqrt{5000}$. (Lower right panel) Quickly vanishing (respectively, long memory of) the autocorrelation of $r(t)$ (thin lines) [respectively, $|r(t)|$ (thick line)] for two values of k and of Δk. (Lower left panel) NIV resonance in the presence of a time varying $K(t)$, with $\Delta k = 0.1$ (circles), 0.2 (squares), 0.5 (diamonds). The other parameters are $N = 10^4$, $\sigma_{ap} = 0.04$, $D = 1$.

associated with excursion of $K(t)$ in the neighborhood of k_c. The lower left panel of Fig. 7 shows the robustness of the NIV phenomenon as a function of the average coupling k: Even with a fluctuating $K(t)$, a large volatility peak appears for intermediate values of k. The lower right panel shows very short-range correlations of $r(t)$ but very long-range correlations of the financial volatility $|r(t)|$ (another equivalent proxy for volatility), very similar to empirical observations of financial returns [31]. Such long persistence of the volatility can be traced back to the slow diffusive nature of $K(t)$ in line with the investors' slowly changing trust in the economy. From the previous section, it also follows that during times of crisis and strong social interaction (k close to k_c), the dynamics is generated mostly exogenously, well in line with the documented inability of news events to explain large price movements [30].

VI. CONCLUSIONS

In this paper, we have investigated the behavior of a system composed of coupled bistable units under the influence of a—rapidly varying—common exogenous forcing and independent noise sources. Independently of the shape of the driving force, intermediate noise strengths trigger a strong level of fluctuations of the macroscopic dynamics around the critical value separating the ordered from the disordered phases. For a periodic forcing, this peak corresponds to a pronounced amplification of the signal, with a strong correlation between the macroscopic dynamics and the driving force at the optimal lag, the paradigmatic signatures of stochastic resonance.

When the driving force is aperiodic a similar peak appears, but here the amplitude of the fluctuations exceeds by far those observed for periodic signals. Coincidental with the increase of fluctuations, the correlation between the driving force and the system dynamics is completely destroyed. This shows that even though these fluctuations are induced by the common forcing, the macroscopic dynamics has an endogenous origin. This phenomenon of *noise-induced volatility* contrasts with that of stochastic resonance, with the major difference being that it is not the signal, but the fluctuations that are amplified.

Moreover, this phenomenon of *noise-induced volatility* also constitutes a new indicator for the approaching of a phase transition [32], and it applies to a broader range of real-world systems due to the more common setup given of a coupled system driving by an aperiodic forcing and its robustness with respect to changes in the underlying network of interactions.

As an example of a system where this phenomenon can be observed, we have proposed the social system of stock markets, in which we have been able to not only explain the excess of volatility observed in stock prices, but also the apparent absence of correlation between news and price changes and the persistence of volatility during times of crises.

APPENDIX: EQUIVALENCE TO THE KINETIC ISING MODEL WITH GLAUBER DYNAMICS

A popular update mechanism in the kinetic Ising model literature was introduced by Glauber [50]. In it, the probability for a spin to flip is given by

$$p_{\text{flip}} = \frac{1}{e^{\beta \Delta E_i} + 1}, \tag{A1}$$

where ΔE_i is the energy gained by the system through the spin-flip and $\beta = 1/kT$. With $s_i = \pm 1$ and $E_i = -s_i(\sum_j K_{ij} s_j + f)$, which is the energy of the state s_i, Eq. (A1) can be rewritten as

$$p_{\text{flip}} = p_{s_i \to -s_i} = \frac{1}{e^{s_i 2\beta[\sum_j K_{ij} s_j + f(t)]} + 1} = \frac{1}{e^{s_i 2\beta\Lambda} + 1}, \tag{A2}$$

where $\Lambda = \sum_j K_{ij} s_j + f$. With the transition rate given by Eq. (A2), we can compute the probability of being in state s_i at time $t + \delta$ by

$$p(s_i;\, t+\delta) = p(s_i;\, t) p_{s_i \to s_i} + p(-s_i;\, t) p_{-s_i \to s_i} \tag{A3}$$

$$\begin{aligned} &= p(s_i;t)\left(1 - \frac{1}{e^{s_i 2\beta\Lambda}+1}\right) + p(-s_i;t)\frac{1}{e^{-s_i 2\beta\Lambda}+1} \\ &= [p(s_i;\, t) + p(-s_i;\, t)]\frac{1}{e^{-s_i 2\beta\Lambda}+1} \\ &= \frac{1}{e^{-s_i 2\beta\Lambda}+1} = p(s_i), \end{aligned} \tag{A4}$$

which is independent of time and gives us the probability of finding spin i in state s_i. Equation (A4) can be rewritten as

$$s_i(t+\delta) = \begin{cases} +1 & \text{with Prob} = (e^{-2\beta\Lambda} + 1)^{-1} \\ -1 & \text{with Prob} = (e^{2\beta\Lambda} + 1)^{-1} \end{cases}$$

$$= \begin{cases} +1 & \text{with Prob} = 1 - F(-\Lambda) \\ -1 & \text{with Prob} = F(-\Lambda), \end{cases} \tag{A5}$$

where $F(x)$ is the cumulative density function (CDF) of a logistic distribution with zero mean and variance $\pi^2/12\beta^2$.

The model studied in this paper, defined by Eq. (1), can be rewritten as

$$s_i(t+\delta) = \begin{cases} +1 & \text{if} \quad \xi_i(t) \geqslant -\Lambda \\ -1 & \text{if} \quad \xi_i(t) < -\Lambda \end{cases}$$

$$= \begin{cases} +1 & \text{with Prob} = 1 - G(-\Lambda) \\ -1 & \text{with Prob} = G(-\Lambda), \end{cases} \tag{A6}$$

with $G(x)$ being the CDF of the probability density function of $\xi_i(t)$, with zero mean and variance D^2. By direct comparison of Eqs. (A5) and (A6), one can see that the model defined by Eq. (1) is equivalent to the kinetic Ising model with Glauber dynamics if the distribution of the noise is chosen to be a logistic distribution.

[1] L. Gammaitoni, P. Hänggi, P. Jung, and F. Marchesoni, Rev. Mod. Phys. **70**, 223 (1998).
[2] A. S. Pikovsky and J. Kurths, Phys. Rev. Lett. **78**, 775 (1997).
[3] C. Van den Broeck, J. M. R. Parrondo, and R. Toral, Phys. Rev. Lett. **73**, 3395 (1994).
[4] P. Hänggi and F. Marchesoni, Rev. Mod. Phys. **81**, 387 (2009).
[5] D. Abbott, Fluctuation and Noise Letters **9**, 129 (2010).
[6] B. McNamara, K. Wiesenfeld, and R. Roy, Phys. Rev. Lett. **60**, 2626 (1988).
[7] K. T. Leung and Z. Néda, Phys. Lett. A **246**, 505 (1998).
[8] K. T. Leung and Z. Néda, Phys. Rev. E **59**, 2730 (1999).
[9] C. Nicolis, J. Stat. Phys. **70**, 3 (1993).
[10] R. Benzi, G. Parisi, A. Sutera, and A. Vulpiani, SIAM J Appl. Math. **43**, 565 (1983).
[11] J. K. Douglass, L. Wilkens, E. Pantazelou, and F. Moss, Nature (London) **365**, 337 (1993).
[12] J. E. Levin and J. P. Miller, Nature (London) **380**, 165 (1996).
[13] E. Simonotto, M. Riani, C. Seife, M. Roberts, J. Twitty, and F. Moss, Phys. Rev. Lett. **78**, 1186 (1997).
[14] L. Gammaitoni, F. Marchesoni, and S. Santucci, Phys. Rev. Lett. **74**, 1052 (1995).
[15] M. Acharyya, International Journal of Modern Physics C **16**, 1631 (2005).
[16] G. Korniss, P. A. Rikvold, and M. A. Novotny, Phys. Rev. E **66**, 056127 (2002).
[17] T. Tomé and M. J. de Oliveira, Phys. Rev. A **41**, 4251 (1990).
[18] G. M. Buendía and P. A. Rikvold, Phys. Rev. E **78**, 051108 (2008).
[19] S. W. Sides, P. A. Rikvold, and M. A. Novotny, Phys. Rev. Lett. **81**, 834 (1998).
[20] J. J. Brey and A. Prados, Phys. Lett. A **216**, 240 (1996).
[21] P. Jung, U. Behn, E. Pantazelou, and F. Moss, Phys. Rev. A **46**, R1709 (1992).
[22] S. Ciliberto, Physica D: Nonlinear Phenomena **158**, 83 (2001).
[23] M. C. Cross and P. C. Hohenberg, Rev. Mod. Phys. **65**, 851 (1993).
[24] I. K. Ono, C. S. O'Hern, D. J. Durian, S. A. Langer, A. J. Liu, and S. R. Nagel, Phys. Rev. Lett. **89**, 095703 (2002).
[25] R. B. Pandey and D. Stauffer, J. Stat. Phys. **61**, 235 (1990).
[26] S. Bewick, R. Yang, and M. Zhang, PLoS ONE **4**, e8112 (2009).
[27] D. Sornette, V. I. Yukalov, E. P. Yukalova, J. Y. Henry, D. Schwab, and J. P. Cobb, J. Biol. Syst. **17**, 225 (2009).
[28] R. J. Shiller, The American Economic Review **71**, 421 (1981).

[29] R. Brealey, S. Myers, and A. Marcus, *Fundamentals of Corporate Finance*, 6th ed. (McGraw-Hill Irwin, New York, 2009).
[30] D. M. Cutler, J. M. Poterba, and L. H. Summers, *What Moves Stock Prices?*Tech. Rep. 2538 (National Bureau of Economic Research, 1989).
[31] Z. Ding, C. W. J. Granger, and R. F. Engle, Journal of Empirical Finance **1**, 83 (1993).
[32] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara, Nature (London) **461**, 53 (2009).
[33] E. Stoll, K. Binder, and T. Schneider, Phys. Rev. B **8**, 3266 (1973).
[34] K. Binder, Rep. Prog. Phys. **60**, 487 (1997).
[35] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, New York, 1999).
[36] P. Jung and P. Hänggi, Phys. Rev. A **44**, 8032 (1991).
[37] S. J. Cornell, *Nonequilibrium Statistical Mechanics in One Dimension*, Cambridge Lecture Notes in Physics, Chap. 6 (Cambridge University Press, Cambridge, 1997), pp. 111–139.
[38] W. Zwerger, Phys. Lett. A **84**, 269 (1981).
[39] W. A. Brock and S. N. Durlauf, The Review of Economic Studies **68**, 235 (2001).
[40] A. Beja and M. B. Goldman, The Journal of Finance **35**, 235 (1980).
[41] M. Wyart and J.-P. Bouchaud, J. Econ. Behav. Organ. **63**, 1 (2007).
[42] F. Lillo, J. D. Farmer, and R. N. Mantegna, Nature (London) **421**, 129 (2003).
[43] S. Bikhchandani, D. Hirshleifer, and I. Welch, The Journal of Political Economy **100**, 992 (1992).
[44] D. Stauffer and D. Sornette, Physica A: Statistical Mechanics and its Applications **271**, 496 (1999).
[45] D. Sornette and W.-X. Zhou, Physica A: Statistical Mechanics and its Applications **370**, 704 (2006).
[46] W. X. Zhou and D. Sornette, The European Physical Journal B–Condensed Matter and Complex Systems **55**, 175 (2007).
[47] G. Harras and D. Sornette, J. Econ. Behav. Organ. **80**, 137 (2011).
[48] D. Sornette, Phys. Rev. Lett. **72**, 2306 (1994).
[49] D. Sornette, J. Phys. I **4**, 209 (1994).
[50] R. J. Glauber, J. Math. Phys. **4**, 294 (1963).

5.3 Finalizing Comments

The subject of an Ising model subjected to a changing magnetic field is a popular area of research in statistical physics. However, with the exception of very few studies[1] which consider a stochastic forcing, the majority of such studies investigate the system under periodic forcing only. Also, most studies concentrate on slowly varying signal, which contrasts with our study, where the emphasis is on rapidly varying external signals. These ingredients of a fast stochastic driving, together with the very general set-up of the Ising model, or likewise a discrete choice with social interactions, lead to a wide applicability of the model, especially in the area of *approaching phase transitions*, which is used a lot in ecology and the study of climate change.

In the context of financial markets, several studies confirm our conclusion, that a large part of the volatility of prices can be traced back to noise traders. For instance French and Roll (1986) show that the volatility in stock returns is larger when the stock market is open compared to when it is closed, i.e., trading generates volatility. Foucault et al. (2011) deliver an even clearer picture, as they can show that specifically retail trading activity increases volatility, which involves, in average, the least sophisticated investors on the exchange and the most likely to indulge in momentum trading and herding.

[1] To the best of my knowledge there is one such study, namely the one by Hausmann and Ruján (1997).

Chapter 6

Volatility-induced overreaction

6.1 Introductory comments

Compared to the previous two chapters, which contained theoretical studies, the paper presented in this chapter is an empirical study, investigating the dependencies of the first-order autocorrelation of daily return. The study is applied to price dynamics of equities, traded on the New York Stock Exchange over a time-span of 21 years. The autocorrelation is estimated on moving windows, enabling me to study its dynamics and relations to other moving window statistics, such as volatility, illiquidity or market trend.

A priori one would assume, based on the EMH, that the return autocorrelation, especially of daily returns, would fluctuate around zero in a random fashion, as any significant autocorrelation should be arbitraged away. Surprisingly this is not the case and among others, a strong correlation between the volatility and return autocorrelation, both estimated on moving windows, is observed.

The paper presented in the next section is not yet submitted to any journal, and has also not been peer-reviewed. However, a thorough literature search on the subject of return autocorrelation was conducted and no similar phenomenon was found, such that I can say with strong confidence that the reported phenomenon of volatility induced return autocorrelation is a novel and yet unpublished finding. In order to refute the criticism of weak statistical treatment, an extended list of robustness tests are performed. Also data-snooping can be ruled out, as the phenomenon is highly persistent through-out the full time-span. A systematic error in the data can also be rejected, as the results, obtained with data from Bloomberg, were exactly replicated with data retrieved from Yahoo! Finance, which is an alternative financial data provider[1].

[1]For NYSE traded stocks, Yahoo! Finance get the price dynamics direct from the exchange.

6.2 The paper

Volatility induced Overreaction: Evidence from daily return autocorrelation

Georges Harras & Didier Sornette

April 19, 2012

Chair of Entrepreneurial Risks
Department of Management, Technology and Economics
ETH Zurich, CH-8001 Zurich, Switzerland

Abstract

For individual stocks traded on the New York Stock Exchange, we investigate the dynamics of the cross-sectional average of the first order autocorrelation of their daily returns and show that changes in the average autocorrelation of returns strongly correlate with prior changes in the cross-sectional volatility and market trends. It is found that return autocorrelation relates negatively to past volatilty changes and positively to past market trends. This observation, which is a market-wide phenomenon, is persistent for over 20 years of data and also present in individual stocks. In contrast to the existing literature on return autocorrelation, we can reject the illiquidity and bid-ask bounce as driving forces behind the return autocorrelation dynamics. We propose a behavioral origin of the phenomenon, where high volatility and bear markets lead to uncertainty and panic, reflected in overreacted behavior on a daily scale, whereas low volatility and bull markets lead to overconfidence, identified by price momentum. In order to address the non-stationarity of some of the analyzed time-series, we have developed a very powerful and yet intuitive method to compute meaningful correlations between time-series with various memories.

Tags: *autocorrelation, volatility, overreaction, herding, non-stationary process, leverage effect, LeBaron effect, illiquidity, volume*

1 Introduction

The autocorrelation of financial returns is a subject of great interest to academics and practitioners alike. The existence of an exploitable autocorrelation would challenge the "Efficient Market Hypothesis" (EMH), introduced by Samuelson (1965) and Fama (1970), which states in its weak form, that it should be impossible to make abnormal economic profit, on a risk-adjusted basis, by trading based on information contained in past price dynamics. This makes the return autocorrelation an ideal candidate to investigate the validity of the EMH, while avoiding the joint hypothesis problem.

In the present study, the dynamics of the first-order autocorrelation of daily equity returns is investigated. We find that future return autocorrelation is negatively related to past changes in volatility and positively related to market trends. We propose a behavioral origin of the phenomenon, where high volatility and bear markets are a sign of uncertainty and induce panic, leading to intraday overreaction via trend-following behavior. On the other hand, a decrease of volatility and a bull market, calms investors and induces overconfidence, resulting in smoother and persistent price dynamics. In order to put our work into perspective, we first review the literature on the much studied subject of return autocorrelation and its origins.

For individual stocks, studies on return autocorrelation generally have observed negative autocorrelation on short time scales. Negative autocorrelations were found for daily stock returns of individual companies at a lag of 2 to 13 days (French and Roll, 1986) and at the first lag for weekly and monthly returns (Lehmann, 1990; Nam et al., 2006; Jegadeesh, 1990). On intermediate time-scales, returns are found to be trending, as positive autocorrelations are found for three and twelve months returns (Jegadeesh, 1990; Jegadeesh and Titman, 1993). On the long-term (3 to 5 years), mean-reversion of the prices, i.e. negative return autocorrelation, is again observed (Bondt and Thaler, 1985; Jegadeesh and Titman, 1993). Similar results are found by Cutler et al. (1990), who confirm positive return autocorrelation for returns up to one year for markets for stocks, bonds, foreign exchange, and various real assets and negative return autocorrelation for returns stocks, bonds, and foreign exchange in the time-scale of ~ 2 years.

For index returns, the structure of the return autocorrelation is slightly different compared to individual stocks. Strongly significant positive autocorrelations are found for daily, weekly and monthly returns for value- and equal-weighted indices (Lo and MacKinlay, 1988; Campbell et al., 1997). This conversion from negative to positive return autocorrelation on small time-scales can be explained by the cross-correlation between individual stocks, as larger capitalization stocks lead and smaller capitalization stocks lag (Lo and MacKinlay, 1990).

On longer time-scales, index returns also show anti-persistence. Fama and French (1988) find negative autocorrelation for 3-5 year returns with equal- and value-weighted portfolios. For national stock market indices, Richards (1997) finds persistent returns on time scales smaller than one year and negative return autocorrelation for larger time scales, with the strongest reversal for 3- and 4-year returns. Similar results for equal- and value-weighted CRSP portfolios are reported by Poterba and Summers (1989), who find positive autocorrelation for returns shorter than one year and negative on the longer scale (~ 8 years). Such results are confirmed by Chopra et al. (1992), who report negative autocorrelation for 5 year returns.

The literature provides two explanations[1] for the occurrence of negative autocorrelation in financial returns. The first one is overreaction to news, either due to momentum traders, who amplify the actions of news-based traders (Hong and Stein, 1999) or behavioral reason such as overconfidence (Daniel et al., 1998), pushing the price too far into one direction and leaving the stock over- or under-valued after the excess demand has calmed off. It is the correction of the price in the opposite direction, which then leads to the mean-reverting characteristic, i.e. negative autocorrelation, of the returns. For the mean-reversion on the long time-scales, this seems to be a valid explanation as the presence of positive autocorrelated returns on intermediate time-scales is a sign of trend-following and herding behavior, leading to prices above or under their fundamental value.

The alternative explanation for the return reversals is given by the bid-ask bounce, making negatively correlated returns a characteristic of illiquid assets. This alternative was shown to be a valid explanation for the anti-persistence at short time-scales (daily to monthly) (Grossman and Miller, 1988; Kaul and Nimalendran, 1990; Ball et al., 1995; Jegadeesh and Titman, 1995; Conrad et al., 1997), where the negative autocorrelation disappears, once the bid-ask spread is controlled for. This indicates that even with significant negative return autocorrelation, it is impossible for regular investors (not for market makers) to make profit from this knowledge, an observation in line with the weak form of the EMH.

A more elaborate way of investigating the return autocorrelation is to look at the conditional autocorrelation, relating the structure of financial returns to other dynamical variables. In a more

[1] Earlier, non-synchronous trading was proposed as a third explanation. However, as the activity on stock markets has increased tremendously this explanation does not hold for recent and future studies.

recent study on autocorrelation in individual stock returns, Avramov et al. (2006) show a strong relationship between weekly return reversal and prior illiquidity. Contrary to short time-scale price reversals reported in the mid-90s, this negative autocorrelation is not a result of the bid-ask bounce, as the spread for NYSE-traded stocks has decreased significantly in the last 20 years (Chordia et al., 2001, 2008). The reversals are reactions to price pressures by non-informational demands, which is too high compared to the availability of investors willing to take the other side of the trade. During high price pressure periods, liquidity providers demand prices deviating from the security's fair price, to which the price will return after the price pressure is reduced. This intuition is confirmed by the negative relation between (i) daily volume and subsequent autocorrelation (Campbell et al., 1993) and (ii) trading activity and subsequent autocorrelation in weekly returns (Conrad et al., 1994). Similar price reversions based on daily and weekly return have been documented by Nam et al. (Nam et al., 2006), whose analysis is however unconditional.

The relation between autocorrelation and volatility for stock returns was first investigated by LeBaron (LeBaron, 1992), who found a negative relationship between volatility and first-order return autocorrelation in daily and weekly returns of indices. The relationship is investigated by relating the volatility at a given day (or week), estimated by GARCH-like models, to the relation between the returns at the same and next day (or week). Similar results for individual stocks and indices were found in other studies (Koutmos, 1997; Booth and Koutmos, 1998; Venetis and Peel, 2005; Bohl and Siklos, 2008), which employed a similar statistical approach. Contrarily to LeBaron's original paper, these studies also document negative autocorrelations, which they interpret as confirming evidence for the model proposed by Sentana and Wadhwani (1992), a simple model, which can relate return autocorrelation to the volatility-depended behavior of trend-following and fundamental investors.

Also, in the most recent study of the LeBaron-effect (Bianco et al., 2009), the relation between volatility and autocorrelation is restricted to immediate impacts from the former to the latter, comparing the t-to-$t+1$ return relationship with the volatility at time $t-1$ and t. The study, which is applied to the price dynamics of futures on the Standard and Poor 500 stock index from 1993 to 2007, states that the effect has significantly diminished for daily returns since LeBaron's paper. However, they find that serial correlation of 5-minute intra-day returns is negatively correlated with expected volatility and positively correlated with unexpected volatility, a finding not completely compatible with the feedback trading model of Sentana and Wadhwani (1992).

The main interest of the present study lies in the impact of various time-series on the dynamics of the first order autocorrelation of daily returns. In contrast to previous studies, we will not investigate the short-term impact but will focus on the low-frequency dependences of the intermittent return autocorrelation. The dynamics of return autocorrelation are obtained by computing the sample autocorrelation of daily returns within a moving-window and will be referred to as "local" return autocorrelation. The proposed explanatory variables are volatility, returns, illiquidity and transaction volume, in line with the existing literature. As all those variables, with the exception of the returns, exhibit a long memory, it seems only natural that, in case of a genuine dependency, the local autocorrelation dynamics is bound to exhibit slowly varying characteristics as well. By focusing on time-scales of several months instead of days or weeks, we study the relation between different regimes as opposed to the immediate impact from one onto another variable. Consequently, we do not constrain ourself to cross-correlations at a fixed small lag, but take a more data-driven approach and extend the possible dependencies from an explanatory variable onto the return autocorrelation to greater lags.

Another difference to the existing literature lies in the fact that we not only consider the various dependences independently for each stock, but that the main analysis will be performed on daily cross-sectional averages, a procedure also employed in a study of liquidity and trading activity (Chordia et al., 2001). By doing so, we investigate market-wide changes in the return structure and their origins, which will not only significantly boost the observed signal, but also help to differentiate stock-specific from market-generic effects.

By computing the cross-correlation between the moving window time-series of return autocorrelation and volatility, a strongly significant negative correlation between volatility and autocorrelation can be observed. However, in contrast to previous findings, the correlation is strongest if the moving return autocorrelation is lagged by 3-4 month with respect to the volatility, i.e. a change of volatility will negatively impact the daily return autocorrelation 3-4 months later. The robustness of the result is impressive, as it is consistent over time, over a large number of individual stocks and robust with respect to statistical methods and their parameters. At zero lag, their correlation is statistically insignificant, indicating that, by looking at the lower end of the spectrum of the dynamics, we find a qualitatively different phenomenon to those reported in the existing literature.

The relation between price movements and moving window return autocorrelation yields a positive correlation at the 3-4 month lag. Given the impact of the volatility, this observation is however not surprising as there is a strong negative correlation between past returns and future volatility, a stylized fact known as the “leverage effect”(Black, 1976), which was however later found to not be caused by leverage (Hasanhodzic and Lo, 2011; Hens and Steude, 2009). Much weaker and lagged impacts onto return autocorrelation is observed for the moving window dynamics of illiquidity and average transaction volume, leading us to the conclusion that, on the larger scale, the impact of these two dynamics is by far outweighed by the volatility's and market trend's impact. In contrast to the feedback model of Sentana and Wadhwani (1992), we propose a behavioral origin of the phenomenon, well confirmed by our analysis.

The remainder of the paper is organized as follows. Section 2 presents the data and used metrics to create the various moving-window time-series. A simple analysis, with some preliminary results, and a motivation for the detailed analysis is given in Section 3. As many of the analyzed variables have a non-negligible autocorrelation, which can create a spurious cross-correlation, we introduce in Section 4 a novel and intuitive method to reduce the memory of the variables, which vastly simplifies the investigation of dependencies between slowly varying variables. The main results of the analysis are presented in Section 5 and their interpretation is given in Section 6. A detailed analysis of the result's robustness is shown in Section 7. Section 8 lists some possibilities for future research based on our findings and Section 9 concludes.

2 Data and correlation measures

For this study, we use historical data of approximately 3000 stocks traded on the New York Stock Exchange (NYSE) obtained from Bloomberg (Historical End of Day Data). We collected data for open, high, low, close, closing ask, closing bid, volume and market capitalization[2] from January 1. 1984 to September 9. 2011.

The stocks were selected by taking the union of all the stocks traded on the NYSE that populated the Russell 3000 Index from 1995 to 2011. Stocks with incomplete data in the various fields were removed. The data was considered incomplete when the time-series of any field contained consecutive holes, which exceeded 20 days. Smaller holes in closing bid and ask where filled by the closing price, holes in the other fields were filled by their previous value.

As the negative autocorrelation in daily and weekly returns was found to have its origin in the bid-ask bounce, we control for the spread by computing the return autocorrelation based on the

[2]The Bloomberg Mnemonics for these fields are PX_OPEN, PX_HIGH, PX_LOW, PX_LAST, PX_BID, PX_ASK, PX_VOLUME, CUR_MKT_CAP.

closing mid-price (p_{CM})

$$p_{CM}(t) = \frac{\text{closing bid price}(t) + \text{closing ask price}(t)}{2}$$
$$r(t) = \log\left[\frac{p_{CM}(t)}{p_{CM}(t-1)}\right], \tag{1}$$

which eliminates the bid-ask bounce.

As we are interested in investigating low-frequency dependence between different variables, we will compute the cross-correlations within moving-window time-series, such that the fluctuations on daily and weekly scales are averaged out and only the large-scale dynamics remain. The moving-window mean and variance of the daily returns, $r(t)$, with a window-size of Δ, are given by

$$\mu_r(t,\Delta) = \frac{1}{\Delta}\sum_{t'=t-\Delta+1}^{t} r(t') \tag{2}$$
$$\sigma_r^2(t,\Delta) = \frac{1}{\Delta}\sum_{t'=t-\Delta+1}^{t} \big(r(t') - \mu_r(t,\Delta)\big)^2. \tag{3}$$

The measure's value at time t is based on data from time $t-\Delta+1$ till t included. To estimate the relation between two time-series $x(t)$ and $y(t)$, with $t \in \{1,\ldots,T\}$, at a lag of τ, two measures of cross-dependence are used in this study:

- Pearson product-moment correlation coefficient $\rho(x,y,\tau)$,
- Kendall's τ rank correlation coefficient(Kendall, 1938; Kruskal, 1958), which we will denote as $K\tau(x,y,\tau)$.

For the estimated cross-dependence to give meaningful results, $x(t)$ and $y(t)$ have to be stationary time-series, a condition which is not satisfied for several of the analyzed time-series. In section 4, we will introduce a novel method to obtain stationary increments from these time-series, such that meaningful cross-dependence can be computed. By $\tilde{\rho}(x,y,\tau)$ and $\widetilde{K\tau}(x,y,\tau)$, we denote the cross-dependence between x and y, where the two time-series have first been subtracted by their non-stationary mean, before estimating their cross-correlation.

As the objective of this study is to investigate the dynamical dependencies of the return autocorrelation, the autocorrelation will be estimated in moving-windows over daily returns, similar to the moving-window mean and variance of Eq. (2), respectively Eq. (3). The statistical measures based on moving window estimates will be referred to as "local" estimates. In order to simplify some of the expressions later in the text, we introduce the following syntax

$$x[t_1,t_2,\omega] = \{x(t_1), x(t_1+\omega), x(t_1+2\omega)\ldots, x(t_1+n\omega)\}, \tag{4}$$

the subset of x, ranging from t_1 to t_2 in steps of ω, where n is specified by $n\,\omega \le t_2 - t_1 < (n+1)\,\omega$ and $\omega \in \mathbb{N}_+$. To estimate the autocorrelation of the daily returns, we extend the previously introduced correlation estimates with the addition of a third measure, the variance ratio, in order to test the robustness of the results:

- the autocorrelation at lag τ, based on Pearson's correlation coefficient is given by

$$A\rho_r(t,\Delta,\tau) = \rho(r[t-\Delta+1,t,1], r[t-\Delta+1,t,1],\tau) \tag{5}$$
$$= \frac{\frac{1}{\Delta}\sum_{t'=t-\Delta+1}^{t-\tau}\big(r(t') - \mu_r(t,\Delta)\big)\big(r(t'+\tau) - \mu_r(t,\Delta)\big)}{\sigma_r^2(t,\Delta)} \tag{6}$$

- the autocorrelation computed by the use of Kendall's τ (Kendall, 1938; Kruskal, 1958), which has to be corrected for a bias, that arises when the two input-samples are the same, but lagged, time-series (Ferguson et al., 2000),

$$K\tau AC_r(t,\Delta,\tau) = K\tau\big(r[t-\Delta+1,t,1], r[t-\Delta+1,t,1],\tau\big) - \frac{2}{3(\Delta-1)} \quad (7)$$

- the variance ratio as introduced by Lo and MacKinlay (1988); Campbell et al. (1997)

$$\begin{aligned} \mathrm{VR}_r(t,\Delta,q) &= \frac{1}{m\,\sigma_r^2(t,\Delta)} \sum_{k=t-\Delta+1}^{t-q} \left(\sum_{t'=k+1}^{k+q} r(t') - q\mu_r(t,\Delta) \right)^2 \\ m &= q(t-q+1)(1-\frac{q}{t}). \end{aligned} \quad (8)$$

Due to its non-parametric nature and its independence with respect to the processes' marginals[3], the autocorrelation based on Kendall's τ ($K\tau AC$) is a very robust measure, which performs well for heterogeneous time-series, and will be our measure of choice.[4]

The daily volatility is estimated by the method proposed by Yang and Zhang (2000), which uses multiple periods of open, high, low, and close prices in historical time series. The local volatility, estimated by this method using data from time $T-\Delta+1$ to T, will be referred to as $\sigma_{ohlc}(T,\Delta)$. Again, very similar results are obtained with alternative volatility estimators, such as σ_r^2 or the non-parametric scale estimator S_n, introduced by Croux and Rousseeuw (1992). For our main analysis, we will use the logarithm of the volatility to reduce the tails of the distribution and work with a variable which is closer to normal.

Another possible explanatory variable for the local autocorrelation is the absence of liquidity (Avramov et al., 2006), estimated by the illiquidity measure introduced by Amihud (2002). This estimator, which is also highly correlated with high-frequency measures of illiquidity, is based in the ratio between the daily absolute price change and dollar value of trading volume. The average illiquidity as a function of time is given by

$$illiq(t,\Delta) = \frac{1}{\Delta} \sum_{t'=t-\Delta+1}^{t} \frac{|r(t')|}{\frac{p_{\mathrm{high}}(t')+p_{\mathrm{low}}(t')}{2} \cdot \mathrm{volume}(t')} \cdot 10^6, \quad (9)$$

where the denominator estimates the dollar trading volume of day t'. The highest and lowest transaction price of day t' are given by $p_{\mathrm{high}}(t')$ and $p_{\mathrm{low}}(t')$, whose average estimates the average transaction price during day t'.

The different dynamic measures introduced in this Section have two common properties. First, they are daily time-series in units of working days. Second, the measure's value at day t is a condensation of the information contained in the last Δ working days, from working day $t-\Delta+1$ to and including working day t. As such, they are backward-looking local statistical measures, allowing us to analyze their dynamical properties instead of only investigating them on a global level. From the moving-window characteristic follows that, for any t and t' such that $|t-t'| < \Delta$, the measures' values at time-step t and t' are obtained using partly overlapping information, creating mechanically a memory over Δ days. The simplest measure is $\mu_r(t,\Delta)$, which is the moving window average return with a window-size of Δ working days.

The main analysis of the relation between the different time-series will be performed on their daily cross-sectional averages, i.e. the average of their daily value over all stocks. To test the

[3]Kendall's τ is a copula measure of dependence.

[4]Very similar results are obtained with the other two autocorrelation estimators, as is shown in Figure 15.

robustness of the results, the analysis will be repeated for different subsets of companies, grouped by their market capitalization and for the individual stocks themselves. The different capitalization groups are constructed on the first working day of every year and kept constant during that year. The number of stocks per capitalization group changes with time, as old companies leave and newer ones enter the exchange.

Once the different companies are grouped together, either in one or multiple groups, we can consider the group as a portfolio and compute the equal-weighed daily returns[5] of the portfolio and get the moving autocorrelation of the daily portfolio returns. By comparing the average moving autocorrelation of returns over many companies with the moving autocorrelation of the portfolio's returns, we can identify if the observed correlations are due to co-movements of the price, or co-movements of the individual stock's local autocorrelation.

3 Motivation

As stated in Section 1, most studies on return autocorrelation consider returns over periods longer than one day. A reason for the lack of studies on the autocorrelation of daily returns might be that, for anyone involved with financial markets, it seems "obvious" that such an autocorrelation should be absent or spurious, as markets are thought, at least on the short term, to be very efficient. In this section, we will show that even the simplest of analysis yields some indications that the structure of daily returns is richer then usually presumed.

Fig. 1 displays the evolution of first-order autocorrelation of daily returns based on yearly estimates. In every panel, the full line shows the normalized histogram of $K\tau AC(\tau = 1)$ for all available stocks, with the year being specified in the top left corner of the panel. The vertical full lines indicate the 20th and 80th percentile of the histogram. In order to compare the measures of the autocorrelation with the null hypothesis of zero autocorrelation, the dotted curves represent the histogram of the first-order autocorrelation of the randomized returns, with the thin dotted vertical lines indicating their 20th and 80th percentile. The randomization is implemented by multiplying the returns with random signs, which destroys the directional information but keeps the heteroscedastic nature unchanged[6]. For every stock, the autocorrelation of 10 randomized return realizations are measured, as to obtain smoother zero-hypothesis histograms.

By comparing the full lines from the real data with the dotted lines obtained from the randomized data, it becomes apparent that, for many years, there are strong deviations between real and randomized data. The average autocorrelation, indicated by the red dashed-dotted vertical line, is positive for 13 of the 21 years but also shows some strongly negative values, clustered around 2008, where the most negative value is measured and the financial credit crisis peaked. Also around 2003, where the dot-com crisis raged, a cluster of negative values is found. As indicated by the clustering of negative and positive values, the dynamics of the yearly mean autocorrelation shows a significant amount of persistence, with an autocorrelation of 0.47 at the first lag.

A preliminary investigation on the relation between return autocorrelation and volatility is performed by linearly regressing the yearly autocorrelation mean onto the median of the yearly OHLC-volatility, σ_{ohlc}, of the same set of stocks. The regression results are shown in Table 1. The significant negative slope (p-value = 0.03) confirms the intuition of the qualitative observation that negative daily return autocorrelations coincide with financial crises. The result is robust in respect to other autocorrelation estimators, c.f. Table 2 in the appendix for the regression results for Pearson's

[5] Discretely discounted returns are used to compute the portfolio returns as continuously discounted returns cannot be cross-sectionally averaged.

[6] Compared with reshuffled returns, which also destroys the heteroscedasticity, randomizing the signs of the returns is a less invasive method.

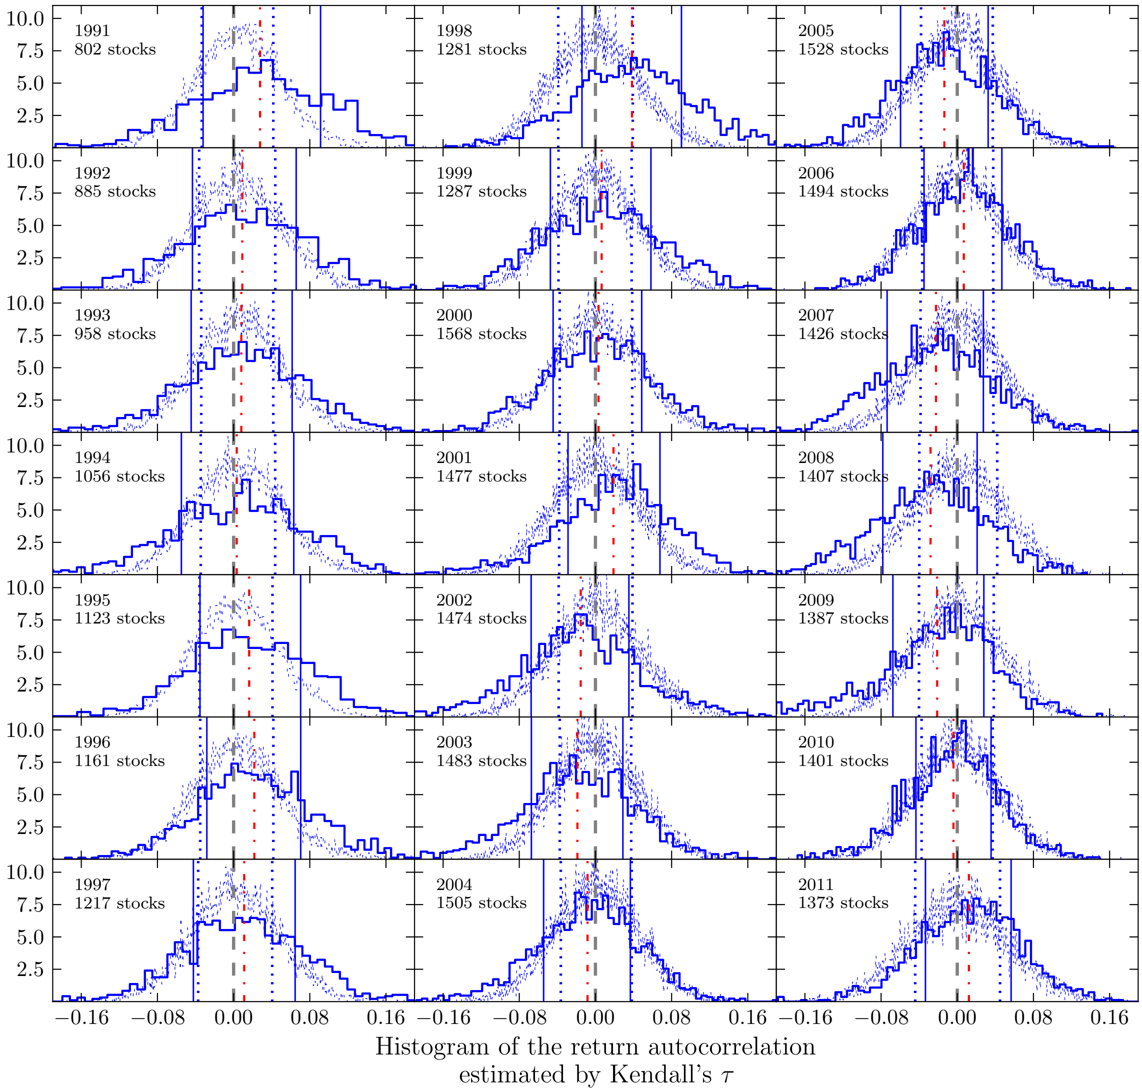

Figure 1: Normalized histogram of the first-order autocorrelation of daily returns, $K\tau AC(\tau = 1)$, for all available stocks (thick full lines). At the top left corner of every panel, the first number specifies the year (from first of January to last of December) over which the data is used for the autocorrelation estimation. The second number gives the number of stocks used to compute the histogram. The vertical full lines indicate the 20th and 80th percentile of the histogram. Dotted lines represent the histogram of the same metric over the same time, but where the returns were randomized by multiplying them by random signs. For every available stock, 10 randomized realizations of returns are generated and their autocorrelation is computed. These distributions represent the null-hypothesis of no return autocorrelation. The thin dotted vertical lines represents the 20th and 80th percentile of the autocorrelation histogram of the randomized returns. The thick dashed vertical line indicates $K\tau AC(\tau = 1) = 0$ and the thin red dashed-dotted line the average autocorrelation of that year. Fig. 17 displays the same information for Pearson's autocorrelation.

	Estimate	Std. Error	t value	Pr(>|t|)
Intercept	0.0273	0.0111	2.46	0.0237
Slope	-1.1821	0.5015	-2.36	0.0293

Table 1: Results from the linear regression of the mean of the yearly Kendall's τ autocorrelation versus the median of the yearly volatility, σ_{ohlc}. Residual standard error: 0.016 on 19 degrees of freedom, multiple R-squared: 0.23, p-value: 0.03. In the appendix, the results for Pearson's autocorrelation are reported in Table 2, whose regression on volatility has multiple R-squared of 0.30 and p-value of 0.01.

autocorrelation with a p-value of 0.01.

From Fig. 1 and the positive autocorrelation of the dynamics of the yearly autocorrelation, it follows that the return autocorrelation is a slowly moving process. Similar results hold for the volatility, which was found to be a long-memory process (Ding et al., 1993; Liu et al., 1999)[7]. In order to perform a more detailed analysis of the relation between volatility of daily returns and their autocorrelation, a novel statistical method will be developed in section 4, which is able to handle the slow moving dynamics of the time-series at hand.

4 Adaptive detrending of slowly moving time-series

4.1 Justification of the method

As noted in the previous section, the volatility and local return autocorrelation are measures with a non-negligible memory, for which one has to correct, before computing cross-correlations between these processes that are not spurious (Yule, 1926; Granger and Newbold, 1974). One way of correcting for memory of a stochastic process is to fit it by an ARMA-process and perform the cross-correlation analysis on the residuals. However, one hypothesis behind any ARMA process is constant (stationary) mean, which is not satisfied for the given time-series, at least not on the time-scale at which our analysis will be done. An alternative approach is to compute the cross-correlations on the first differences of the process, i.e. interpret the slowly moving time-series as unit root processes, as adopted by studies investigating time-series such as volatility, illiquidity, bid-ask spread and so one (Chordia et al., 2001; Hasanhodzic and Lo, 2011). However, as the unit-root can be rejected in most cases, the first difference processes will have a negative autocorrelation at the first lag, leading to deteriorating effects of the overall analysis.

The approach adopted in the present study is that the different time-series are not considered as ARMA-processes nor as unit-root processes, but as processes that fluctuate around a slowly evolving local mean. After estimating this local mean, deviations from that mean are interpreted as innovations, on which the cross-correlation is estimated. This method is optimal for processes, which are neither memory-less with constant mean, nor unit-root processes, but whose first differences constitute a stationary process. In order to show the robustness of the results, the analysis is repeated for the raw time-series, without any nonlocal demeaning and for their first-differences in Section 7.

Another major difference to the general practice used with ARMA-like models is that the calibration of our method will be based on past data only, transforming the method's parameters into dynamics variables. One advantage of this practice is that the local parameters are better adapted to changing environments. A more important second advantage is that the results are as if they

[7]Here we report an autocorrelation of 0.59 at the first lag for the dynamics of the yearly σ_{ohlc}.

would come from live measurements, eliminating any hind-sight bias. This is very different from the general practice with ARMA- or GARCH-like models for example, where the parameters of the statistical methods are estimated over the full data set, taking the strong assumption of stationarity, which we will not.

In the following, we will introduce our non-stationary demeaning procedure, which makes it possible to compare processes characterized by a long memory, like volatility or return autocorrelation, with short memory processes, like returns.

4.2 Local mean with exponentially decaying weights

One way to estimate a local mean of a time-series is Eq. (2), which uses a uniform weighting over Δ observations. An alternative way, which we will adopt to demean the various processes with memory, is to use weights that decrease exponentially, the further they are in the past

$$\begin{aligned}\tilde{\mu}_x(t) &= (1-\alpha)\sum_{t'=0}^{\infty} e^{t'\ln(\alpha)}x(t-t') \qquad (10)\\ &= \alpha\cdot\tilde{\mu}_x(t-1)+(1-\alpha)\cdot x(t), \qquad (11)\end{aligned}$$

with $\alpha \in [0,1]$, which controls the memory of the moving average process, and $-1/\ln(\alpha)$ being the characteristic memory length of the process. Eq. (11) gives a recursive way of estimating the local mean. For $\alpha = 1$, $\tilde{\mu}_x(t)$ is equal to the global mean $\mu_x(t,t)$ of $x[1,t,1]$, which is the natural extension for Eq. (10). Exponentially decaying weights in the estimation of a local mean is a well used practice (Hamilton, 1994) and has the advantage that (i) recent data is given more importance than past values and (ii) the impact of past data decreases continuously, resulting in smoother dynamics compared to mean estimations based on homogeneous weights.

4.3 Demeaning of time-series

Once the local mean is computed, the demeaned time-series of x is given by

$$x_\mu(t) = x(t) - \tilde{\mu}_x(t-1), \qquad (12)$$

which represent the time-series of innovations of $x(t)$. Here, it is important to note that $x(t)$ is demeaned based on information up to and including time $t-1$, which allows us to predict $x(t)$ at time $t-1$, by estimating the deviation from the mean, $x_\mu(t)$, and computing $\tilde{\mu}_x(t-1)$.

Our non-stationary demeaning method is finalized by the following criterion, which sets the value of α in Eq. (10). As our objective is to retrieve the innovations from a slowly evolving process, the optimal α, which will be denoted as $\hat{\alpha}$, is such that it reduces the memory of the demeaned process. As a proxy for the memory of a time-series, we use the sum of its autocorrelation over the first $\tau_{\hat{\alpha}}$ lags[8]. We will consider the reduction of the memory of $x_\mu(t)$ to be achieved when this sum approaches zero. Again, as we do not want to use any data from the future for the calibration, $\hat{\alpha}$ will be estimated based on $\Delta_{\hat{\alpha}}$ time-steps in the past, where $\Delta_{\hat{\alpha}}$ is large compared with the intrinsic time-scale of $x(t)$. With this estimation procedure, optimal α will become a dynamic variable, $\hat{\alpha}(t)$, given by

$$\begin{aligned}\hat{\alpha}(t) &= \max\left\{\alpha : \sum_{\tau=1}^{\tau_{\hat{\alpha}}} A\rho_{x_\mu}(t,\Delta_{\hat{\alpha}},\tau)=0\right\} \qquad (13)\\ &= \max\left\{\alpha : \sum_{\tau=1}^{\tau_{\hat{\alpha}}} A\rho(x_\mu[t-\Delta_{\hat{\alpha}}+1,t,1],\tau)=0\right\}.\end{aligned}$$

[8]We will set $\tau_{\hat{\alpha}} \gg \Delta$.

In order to reduce unnecessary optimization, $\hat{\alpha}(t)$ will only be updated every $\delta_{\hat{\alpha}}$ time-steps and will remain constant in between. From Eq. (13), it follows that $\hat{\alpha}(t) = 1$ when $x(t)$ has no memory, such that $x_\mu(t) \sim x(t) - \mu_x(t,t)$ as $\Delta_{\hat{\alpha}}$ is large compared to the time-scale of $x(t)$. Similarly, when $x(t)$ is a unit-root process, $\hat{\alpha}(t) = 0$ and $x_\mu(t)$ reduces to the first-differences process of $x(t)$. For a process being characterized by fluctuations around a slowly evolving mean, so being neither white nor brown, $\hat{\alpha}(t) \in]0,1[$.

For the analysis in this study, we set the range over which $\hat{\alpha}$ is estimated ($\Delta_{\hat{\alpha}}$) to 1750 working days, i.e. approximately 7 years, and the period over which the estimated $\hat{\alpha}$ will be used to compute the local mean ($\delta_{\hat{\alpha}}$) is set to 100 working days. The reported results are not sensitive to the exact values of $\Delta_{\hat{\alpha}}$ and $\delta_{\hat{\alpha}}$ as long as $\Delta_{\hat{\alpha}}$ exceeds the characteristic length of any of the involved time-series and $\delta_{\hat{\alpha}}$ is small enough, such that the used value of $\hat{\alpha}$ is up-to-date.

The number of lags of the autocorrelation over which α is optimized ($\tau_{\hat{\alpha}}$) will be set such that the maximal lag is approximately one year. The relation between the value of $\tau_{\hat{\alpha}}$ and $\hat{\alpha}$ is weakly positive. With larger $\tau_{\hat{\alpha}}$, the autocorrelation of the demeaned time-series can go to zero in a slower fashion, enabling slow motions of the time-series, which is equivalent to a longer memory, i.e. a value of $\hat{\alpha}$ closer to 1. Due to their weak relation and the robustness of the method, the exact value of $\tau_{\hat{\alpha}}$ does not have an impact on the overall results.

4.4 Demeaning of moving window time-series

As noted earlier, the measures introduced in Section 2 are moving window daily time-series and, as such, they have a minimum time-scale equal to the size of the moving window, Δ. On top of this mechanically introduced memory, the time-series such as volatility and moving autocorrelation also have an intrinsic memory as stated previously. If $x(t)$ is a daily moving window time-series, the spirious memory over Δ time steps introduced by overlapping time windows is removed by only considering the dependence measures every Δth time-step[9]. As such, the complete time-series of $x(t)$ for $t \in \{0, \ldots T\}$ can be represented as the union of Δ time-series, which increments in time-steps of Δ , i.e.:

$$x(t) = \bigcup_{i=1}^{\Delta} x[i, T, \Delta]. \tag{14}$$

The remaining memory in $x[i, T, \Delta]$ is the time-series' intrinsic memory, which will be removed by the previously introduced method[10].

As the optimal $\alpha(t)$ for every $x[i, T, \Delta]$ (denoted as $\hat{\alpha}_i(t)$) will be similar for all i, we choose to estimate only 4 optimal $\hat{\alpha}_i(t)$ for $i \in \{\frac{\Delta}{4}, \frac{2\Delta}{4}, \frac{3\Delta}{4}, \Delta\}$ in order to minimize unnecessary computations. For the computation of the local mean of $x(t)$, we use the average $\hat{\alpha}_i$, $\langle \hat{\alpha} \rangle(t) = 1/4 \sum_{i=1}^{4} \hat{\alpha}_{i\Delta/4}(t)$,

$$\tilde{\mu}_x(t) = \big(1 - \langle \hat{\alpha} \rangle(t)\big) \sum_{j=0}^{\infty} e^{j \ln(\langle \hat{\alpha} \rangle(t))} x(t - j\Delta). \tag{15}$$

Here it is to be noted that x is incremented in steps of Δ, such that only the intrinsic memory and not the mechanically introduced memory is removed. Later, for the cross-correlation analysis, the mechanical memory is of no concern, as only moving-window time-series with the same window-size will be compared.

[9] Every Δth element of a moving-window time-series, with window-size Δ, is estimated based on non-overlapping windows.

[10] As $x[i, T, \Delta]$ are incremented in steps of Δ working days and the units of $\Delta_{\hat{\alpha}}$ and $\delta_{\hat{\alpha}}$ are given in units of working days, the values of these window-sizes will be adjusted to the rounded values of $\Delta_{\hat{\alpha}}/\Delta$ and $\delta_{\hat{\alpha}}/\Delta$ before being used in Eq. (13).

The demeaned moving-window time-series will be computed as

$$x_\mu(t) = x(t) - \tilde{\mu}_x(t - \Delta), \tag{16}$$

where it should again be noted that there is no intersection of the data used for the instantaneous value of the measure, $x(t)$, and the subtracted mean, $\tilde{\mu}_x(t-\Delta)$.

As an example of the non-stationary demeaning, Figure 18 of the appendix shows the time-series of the volatility and local autocorrelation with their local mean and $\hat{\alpha}$.

5 Cross-correlation analysis between moving window time-series

As mentioned before, the time-series, for which the main cross-correlation analysis is performed, are cross-sectionally averaged time-series, indicated by the $\langle\ \rangle$ enclosing the variable names. The cross-section is either taken over all valid stocks available to us, or over a group of stocks, grouped by their market capitalization. These groups are reformed on the first working day of every year and kept fixed for the rest of the year. By using cross-sectional averages, market-wide dynamics are investigated as the idiosyncraticities of the individual stocks are averaged out and only the common characteristics remain.

In this section, the majority of the analysis will be performed on moving-window time-series with a window size of 55 working days. This size was found to be optimal for the present purpose as it is large enough to give a reliable estimate of the local return autocorrelation and small enough to witness its dynamics. The robustness of the results in respect to the window size is confirmed in Section 7.

Fig. 2 displays the cross-correlation between various measures as a function of the lag τ between the time-series. The colored full lines correspond to five different capitalization groups, from the lowest capitalizations (red) to the highest (yellow). The thick dashed lines corresponds to the cross-sectional average over all available stocks. The cross-correlations are estimated based on data ranging from January 2. 1991 to September 9. 2011. As $\Delta_{\hat{\alpha}} = 1750$, data from 1984 to 1991 is used to compute the value of $\langle\hat{\alpha}\rangle$ and the local means at the start of 1991. The advantage of using daily sliding windows time-series, instead of non-overlapping windows, results in a resolution of the cross-correlation, that exceeds the window-size and in a major increase of available data leading to better estimates and smoother lines. The time-series of the volatility and local autocorrelation are shown in Figure 18 of the appendix. The 95% and 99% confidence intervals, indicated by the dashed-dotted and dotted horizontal lines are however computed for cross-correlations of non-overlapping window time-series, leading to wider ranges then necessary[11].

Panel (a) of Fig. 2 shows the cross-correlation between the cross-sectional average of the local first-order autocorrelation of daily returns and the cross-sectional average of daily logarithm of the volatility, both non-stationarily demeaned and estimated in a moving window of 55 working days. As the non-stationarily demeaned time-series are compared with each other, this plot illustrates the impact of a deviation from the local mean in one variable onto the deviation from the local mean of the other variable, not the relation between their absolute values, as was investigated in Table 1. A strongly negative cross-correlation at $\tau \sim -80$ is observed with the minimum of -0.39 for cross-sectional averages over all available stocks. This indicates that a change in the global volatility at

[11] The distribution of the cross-correlation for the null-hypothesis of no serial correlation in the direction of the returns is shown in Figure 16, which confirms the exaggerated confidence bounds. The null-hypothesis is generated by computing the return autocorrelation with the original returns, multiplied by a random sign so as to destroy the memory of the direction and conserve their original heteroscedastic nature.

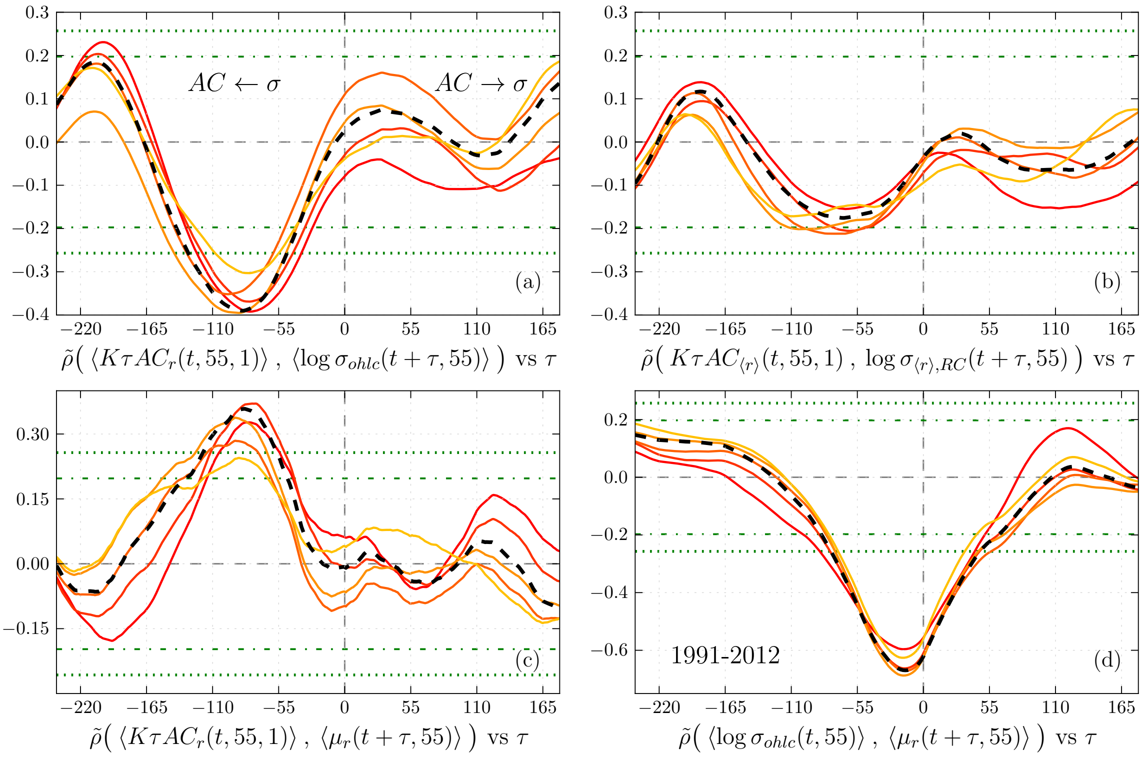

Figure 2: Cross-correlation between non-stationarily demeaned cross-sectional averages of various dynamic measures. The measures are: (a) return autocorrelation and volatility, (b) average-return autocorrelation and average-return volatility, (c) return autocorrelation and return, (d) volatility and return. The different time-series range from January 2 1991 to September 9 2011 and are estimated with a window-size of 55 working days. The dashed black line corresponds to the cross-correlation between measures cross-sectionally averaged over all available stocks, the colored lines correspond to the different capitalization groups, from the lowest (red) to the highest (yellow) capitalizations. Each capitalization group has on average 225 stocks. The dashed-dotted and the dotted horizontal lines represents the 95%, respectively 99% confidence intervals. The average optimal characteristic time-scale for the demeaning is 365 working days for $\langle K\tau AC_r\rangle$, and 135 working days for $\langle\sigma_{ohlc}\rangle$. For $\langle\mu_r\rangle$, $\langle\hat{\alpha}\rangle = 1$, meaning that it is demeaned by the average daily return over the last $\Delta_{\hat{\alpha}}$ days. $\Delta_{\hat{\alpha}} = 1750$ and $\delta_{\hat{\alpha}} = 100$.

some given day gives rise to a change in the overall autocorrelation of daily returns in the opposite direction approximately 80 working days later. This observation tells us that, when the average stock volatility is increasing, the daily price dynamics will show an increase of anti-persistence, approximately 4 month later. Inversely, a decrease of volatility, coincides with a later increase of momentum and smoother daily price dynamics.

For $\tau \sim -200$ in Fig. 2(a), a maximum of cross-correction is observed. The origin of the maximum, which does not exceed the 95%-confidence interval, is not entirely clear. One explanation might be that this maximum is spurious, resulting from the sometimes too strong detrending, leading to an negative autocorrelation of the return autocorrelation dynamics. Evidence for this hypothesis is given by the rather large lag of 10 months and the weakened maximum in Figure 14, where the time-series are not nonstationarily detrended. This maximum will not be further analyzed in this study.

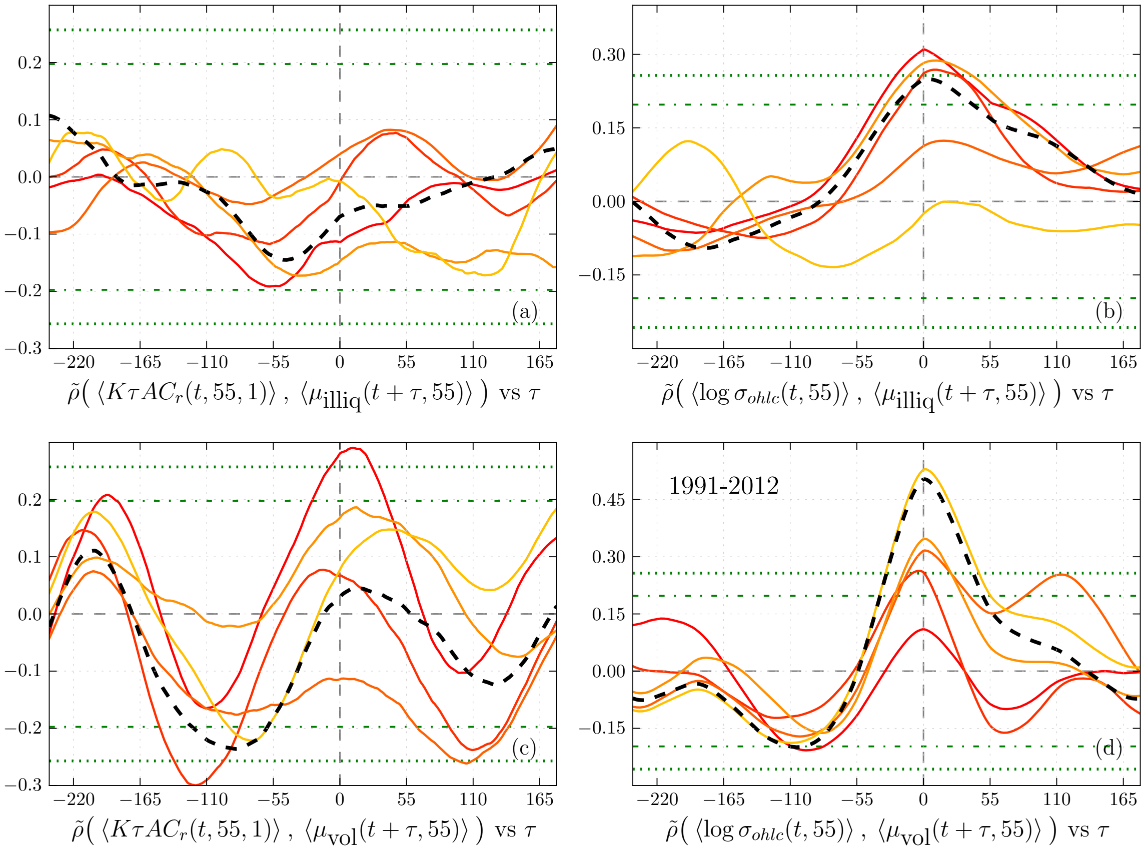

Figure 3: Cross-correlation between illiquidity (upper row), resp. volume (lower row), and return autocorrelation (left column), resp. volatility (right column). The data spans the same period and the results are obtained with the same methods as in Fig. 2.

As the analysis is performed on closing-mid-prices, the change of the return autocorrelation cannot be attributed to the bid-ask bounce, as was the case for price dynamics up to the early 90s. An alternative explanation of the dynamics of the local return autocorrelation could be that the negative autocorrelation is associated with a drop of liquidity, as was found by Avramov et al. (2006). They showed that the likelihood of weekly return reversion increased if the previous week was characterized by low liquidity. Such a relation between return autocorrelation and illiquidity would manifest itself as a significant negative cross-correlation at small negative lags ($\tau \lesssim 0$). To test this hypothesis, the cross-correlation between the moving window illiquidity, estimated via Eq. (9), and the return autocorrelation is shown in Figure 3(a). As predicted, a negative cross-correlation minimum is found for $\tau \sim -50$. The magnitude of the correlation is however statistically insignificant, leading to the conclusion that, for the daily return autocorrelation dynamics, illiquidity is not the right explanatory variable. The cross-correlation between volatility and illiquidity is shown in Figure 3(b), indicating the volatility and illiquidity are correlated at lag 0, even though this correlation is relatively weak, not exceeding the 99%-confidence interval.

As volume can be interpreted as an alternative proxy for liquidity, the previous analysis is repeated for moving window daily transaction volume in Figure 3(c) and (d). For the volume-autocorrelation relationship, a similar shape as for the volatility-autocorrelation is observed, but significantly weaker in magnitude. This can be explained by the well known instantaneous cross-correlation between volume and volatility (Karpoff, 1987; Gallant et al., 1992; Jones et al., 1994), which is displayed in Figure 3(d).

Consequently, we can reject the hypothesis that the dynamics of the daily return autocorrelation have their origin in the changes of the overall market liquidity, as for lags around 0, the cross-correlation between illiquidity, resp. volume, and return autocorrelation is indistinguishable from zero.

The results for the returns of equal-weighted portfolios, based on the capitalization groups, are shown in Fig. 2(b), where the cross-correlation between the moving-window autocorrelation and the volatility of the portfolio returns is shown. The general shape of the cross-correlation as a function of the lag is similar to those of Panel (a) but with a significant weaker negative correlation for $\tau \sim -80$. This indicates that the phenomenon observed in Panel (a) cannot be entirely traced back to the daily co-movement of stock prices but mainly arises through a common change in the auto-correlation of the individual stock return as a reaction to an overall change in volatility. This fact might also explain why this effect was not found by previous studies, which mostly concentrate on index dynamics.

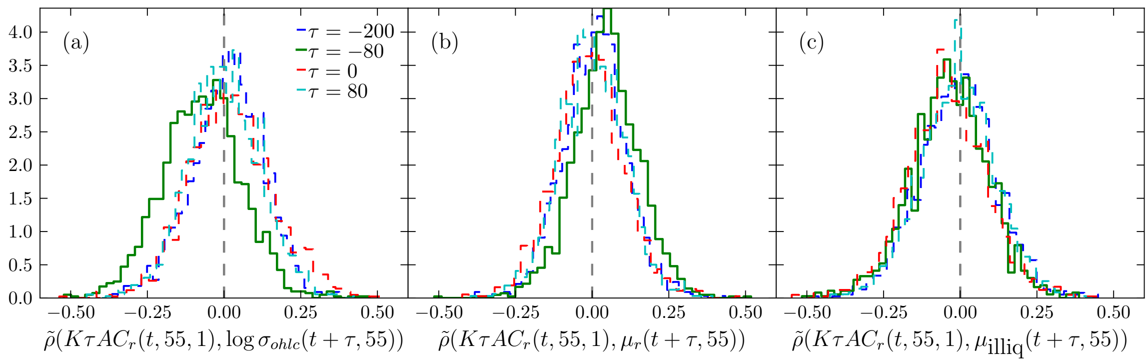

Figure 4: Histogram over all available stocks of the cross-correlation between various time-series of individual stocks at different lags, specified in the legend. The different time-series range from January 2. 1991 to September 9. 2011. All time-series are moving window processes with $\Delta = 55$ and are non-stationarily detrended with $\tau_{\hat{\alpha}} = 5$, $\Delta_{\hat{\alpha}} = 1750$ and $\delta_{\hat{\alpha}} = 100$.

In order to investigate the importance of the synchronicity of the changes of the return structure across the individual stocks, we compute the cross-correlation between the non-stationarily demeaned time-series independently for each available stock. Figure 4 displays the histogram over all stocks of the cross-correlation value at different lags. The investigated lags are $\tau = -80$, where the minimum cross-correlation in Figure 2(a) is found, $\tau = 0$ for instantaneous correction and $\tau = -200, 80$ to confirm the stationarity of the demeaned process. All single stocks are demeaned in the same fashion. From Figure 4(a), which shows the histograms of the cross-correlation between return autocorrelation and volatility, a clear distinction between the cross-correlation values at $\tau = -80$ and the other lags at -200, 0 and 80 working days becomes apparent. From this observation, two conclusions can be drawn. First, the previously reported effect also exists at the level of the individual stocks. Notwithstanding the strongly statistical significant negative average cross-correlation, the magnitude of the individual cross-correlation is far inferior to the values found for cross-sectional averages. This leads to the second conclusion that the dynamics of the local return autocorrelation are similar across a majority of stocks, such that by taking the cross-sectional average, the idiosyncraticities of the individual stocks are removed and only the common market-overarching value of daily return autocorrelation remains. Figure 4(b), which investigates the relation between returns and their autocorrelation at the same lags, will be discussed later in this section.

Figure 4(c) confirms the insignificance of illiquidity to explain return autocorrelation as already found from Figure 3. Also, on the individual stock level, for any lag, none of the histograms are significant different from one another.

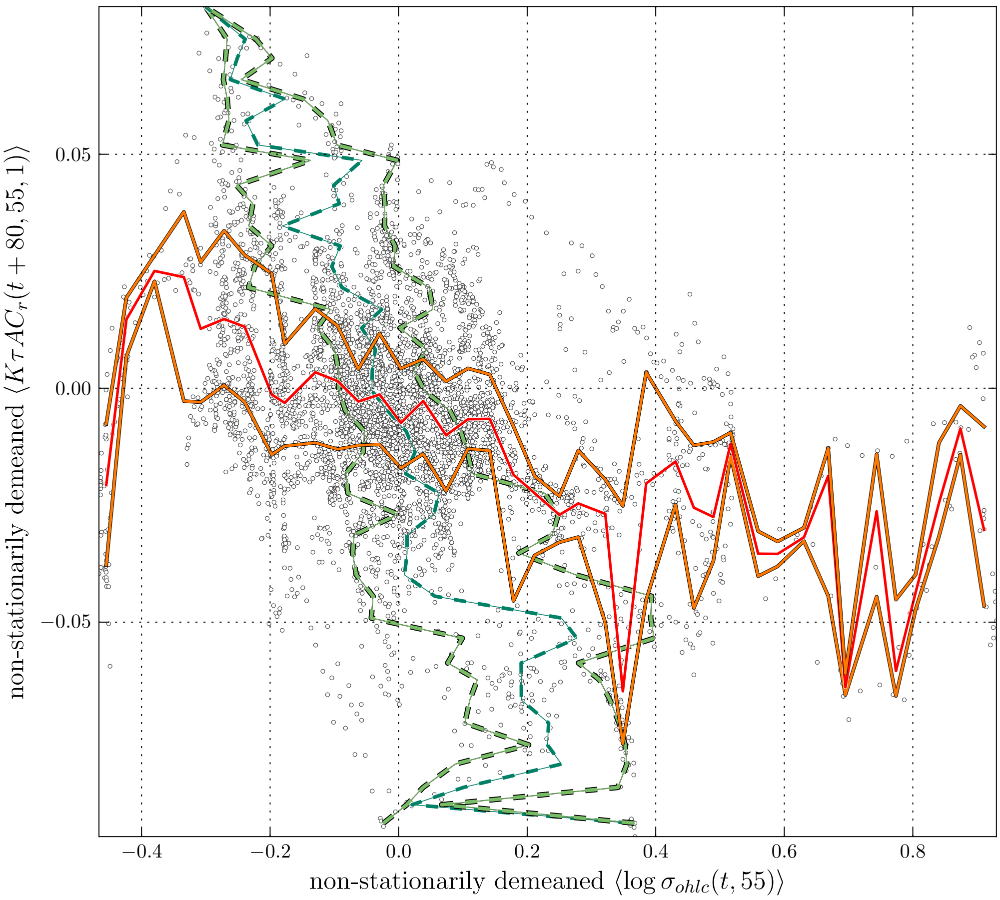

Figure 5: Scatter-plot of local volatility at time t versus the local return autocorrelation at time $t+80$ with $\Delta = 50$. Both are non-stationarily demeaned averages over all available companies. The lines correspond to the 25, 50, 75 percentile lines, where the lines ranging from left to right are based on equidistant volatility bins and the lines ranging from top to bottom are based on equidistant autocorrelation bins. The time-series range from January 2. 1991 to September 9. 2011. The time-series plot of the data plotted here is shown in Fig. 18(d) and (e).

As the main measure of cross-correlation used in this study is the Pearson product-moment, the implicit assumption of a linear relation between the two compared time-series is taken. The non-parametric relation between volatility and return autocorrelation is shown in Figure 5. This scatter-plot contains the same data as is used to compute Fig. 2(a) and displays the two time-series, lagged by 80 days, which is the lag for which the strongest cross-correlation is obtained. The three lines spanning from left to right are the 25%, 50%, 75% percentiles estimated on equidistant volatility-bins and which indicate, for a given volatility change, the change of return autocorrelation 4 months later. A clear negative relation between the two variables is visible, with an approximately linear relation at the origin, supporting the use of a linear correlation measure. Considering that the positive deviations are larger then the negative ones for the volatility, the observed phenomenon has a symmetric nature as an increase of volatility induces a decrease of return autocorrelation and vice-versa.

The three lines spanning from top to bottom are the 25%,50%,75% percentiles estimated on equidistant return autocorrelation bins. They indicate that, given a change of return autocorrelation, what was the change of volatility 4 months ago. As such they investigate events in the past, i.e. looking for an explanation once an autocorrelation change has occurred. This is complementary information to the left-to-right lines, which investigate the predictability of the autocorrelation change in the future for a given change of volatility in the present. From the negative slope of both group of lines, we can conclude that changes of global volatility predict and explain changes in the global return autocorrelation very well. However, from the postion of the percentile lines it seems that the negative autocorrelation is better predicted by positive volatility changes as the positive autocorrelation is better predicted by negative volatility changes, whereas the occurrence of persistence in returns is well explained by a decrease of volatility in the past as the occurrence of anti-persistence is explained by an past increase of volatility. In other words, these results point to the conclusion changes in market volatility are better in predicting future anti-persitence of returns and explaining past return persitence, then the other way around.

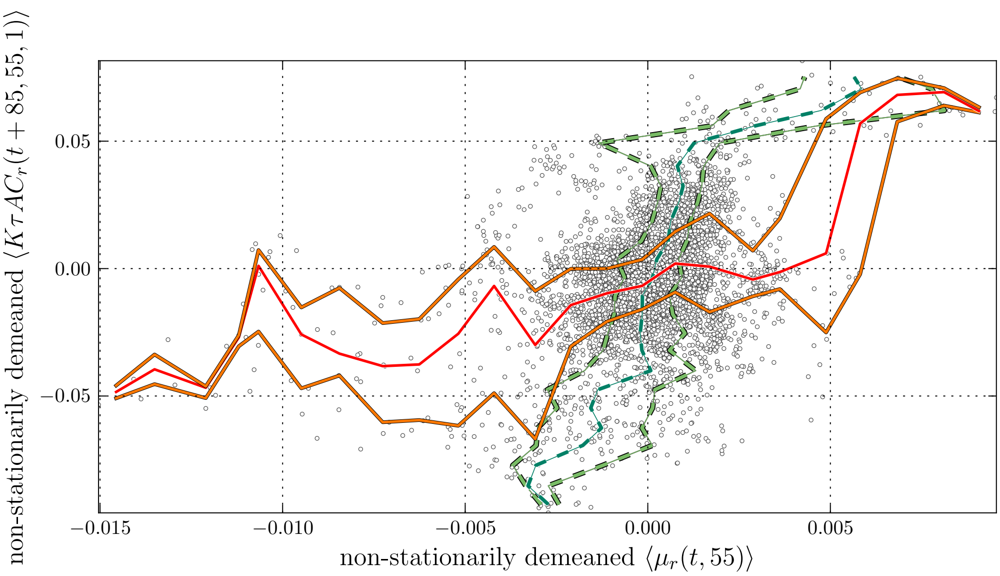

Figure 6: Similar as Figure 5 for the local return autocorrelation versus past cross-sectionally averaged Δ-day returns with $\tau = -85$.

The relation between stock-returns and autocorrelation changes is displayed in Figure 2(c). As the moving average return is computed from continuously discounted returns, $\mu_r(t, \Delta)$ is equivalent to $\log[p_{CM}(t)/p_{CM}(t-\Delta)]/\Delta$, the scaled return over Δ working days. Similar to Panel (a), there is a significant cross-correlations for $\tau \sim -80$, with the cross-correlation reaching a value of 0.36 for $\tau = -85$. Opposite to the volatility-relation, past returns are positively correlated with future return autocorrelation, indicating that a positive price drift induces a momentum in future price dynamics, while a large drop in the price results in anti-persistent return dynamics. From the non-parametric analysis in Figure 6 follows that the predictive power of positive returns is weaker then that of negative returns[12], the explanatory power of returns is however approximately symmetric.

This result is not surprising, given the impact of volatility on return autocorrelation. The negative cross-correlation between past price movements and future volatility is a well known stylized fact of equity markets, which was first observed by Black (1976). Figure 2(d) confirms this phenomenon by the strongly negative cross-correlation between volatility and returns with a value of -0.67 for $\tau = -15$. As the return-volatility-phenomenon is already apparent on the daily scale, i.e. the impact of the return on one day is observable in the magnitude of the return of the following day, the employed analysis method, designed to investigate effect on larger time-scales, cannot make a clear distinction between cause and effect. The reason for this is that for $|\tau| < \Delta$, the two windows, measuring volatility and returns, are overlapping, which is the origin of the unusually strong cross-correlation.

Also on the individual level, a positive relation between past returns and future return autocorrelation is observed in Figure 4(b). The histogram for $\tau = -80$ has a significant positive mean, whereas the means for the other lags are indistinguishable from zero. Compared with the distribution of the cross-correlation between volatility and return autocorrelation, the cross-correlation between returns and return autocorrelation is however weaker, indicating the relative importance of the volatility.

6 Interpretation of results

A popular model, which relates return autocorrelation and volatility, is the model proposed by Sentana and Wadhwani (1992). Although the general outcome of the model, which predicts a negative relation between volatility and return autocorrelation is partly confirmed by our results, we are opposed to their proposed generating mechanism of this relation. In this section, we will interpret our results and propose an alternative generating mechanism, although further analysis is needed to reject definitely either one of the models. In the following, the Sentana-Wadhwani model will first be introduced, and then discussed.

6.1 Sentana-Wadhwani model

The stock market model of Sentana and Wadhwani (1992) is based on the trading of two distinct types of investors. One group is constituted of positive feedback traders (trend follower), who take their trading decisions based on past stock price changes, ignoring any relation of the stock's price to its fundamental value. Their demand function for the shares is given by

$$Y(t) = \gamma\, r(t-1), \tag{17}$$

where $Y(t)$ is the fraction of shares that feedback traders hold, $r(t-1)$ is the stock return in the previous time-step and $\gamma > 0$ is the strength of the feedback.

[12] If the data set is restricted to the last 10 years, the predictive power of returns becomes approximately symmetric, with similar future impacts close to the origin on return autocorrelation for positive and negative returns, shown in Figure 19.

The second group consists of smart money investors, who use all available information to estimate future prices, resulting in a proportional demand for stocks in the form of

$$Q(t) = \frac{E_{t-1}[r(t)] - \alpha}{\eta(t)}, \tag{18}$$

where $E_{t-1}[r(t)]$ represents their expectation of the next return, based on information available up to and including time $t-1$, α is the return on a risk free asset and η is the risk of holding the risky asset. In their original paper, the authors do not put any restrictions on η, however we can safely assume that $\eta > 0$, otherwise the smart money traders would act contrarily to their expectations. It is assumed that this risk is a monotonous increasing function of the volatility, i.e.

$$\eta(t) = \eta\big(\sigma^2(t)\big) \tag{19}$$

with $\eta'(\cdot) > 0$ and $\sigma^2(t)$ being the conditional variance of the returns in period t (formed at time $t-1$). Market equilibrium and conservation of available shares require that

$$Q(t) + Y(t) = 1. \tag{20}$$

Sentana and Wadhwani (1992) assume rational expectations of the smart money traders, such that their predictions about future returns are on average correct,

$$r(t) = E_{t-1}[r(t)] + \epsilon(t), \tag{21}$$

with $\epsilon(t)$ being a zero-mean noise term. Combining this assumption with Eq. (17), (18) and (20) yields

$$r(t) = \alpha + \eta\big(\sigma^2(t)\big) - \gamma\,\eta\big(\sigma^2(t)\big)\,r(t-1) + \epsilon(t), \tag{22}$$

which determines the complete price dynamics.

From Eq. (22), it follows that the return dynamics exhibit autocorrelation as $r(t)$ is partly determined by $r(t-1)$, which enters the equation through the action of the feedback traders. Surprisingly, the trend-following behavior of the feedback traders leads to a negative autocorrelation of the returns, as past returns have a negative contribution ($\gamma > 0$) to future returns.

The origin of this counter-intuitive outcome can be traced back to the smart money traders which can easily be illustrated by the following example. Let us assume that, in the last time-step, the price dropped ($r(t-1) < 0$), which leads to $Y(t) < 0$, i.e. the feedback traders short the asset[13]. This short-selling is enabled by the smart money traders, which have to take the opposite direction, according to Eq. (20), leading to $Q(t) > 1$, which by rearranging Eq. (18) gives

$$E_{t-1}[r(t)] > \eta(t) + \alpha. \tag{23}$$

Together with the assumed rational expectation of Eq. (21), $\alpha > 0$ and $\eta(t) > 0$, the negative return in time-step $t-1$ results in a positive bias for the return in time-step t,

$$E_{t-1}[r(t)] > 0. \tag{24}$$

It follows from this example that it is the reaction of the smart money traders in the opposite direction to the demands of the feedback traders that is at the origin of the negative autocorrelation of the resulting return dynamics. The strengthening of anti-persistence by an increase of volatility stems from the smart money traders' requirement of larger returns in order to take a position in

[13] This is a very artificial assumption for general markets, as short-selling cannot be done for arbitrary volumes or is simply not possible.

risky times. This relation is also revealed in the previous example by Eq. (23), which shows that the expected future return increases with volatility, given a previous negative return.

Besides the fact that negative return autocorrelation is put into the model by hand, via the mechanical effect of Eq. (20) and the rational expectations of the smart money traders, there are several other issues with the model of Sentana and Wadhwani (1992). Given that the behavior of the feedback traders does not change in time ($\gamma > 0$), the model is unable to create a positive return autocorrelation as that factors of $r(t-1)$ in Eq. (22) are both positive. However, as it is shown in Fig. 1, for a majority of the yearly estimates the return autocorrelation is positive.

An other issue is that the model does not make any testable predictions on the mechanisms leading to the autocorrelation. It just assumes rational expectations, a feature which is hard to reject in the given circumstances. Next we will propose and alternative generating mechanism for the return autocorrelation with easily testable predictions.

6.2 Alternative interpretation

Our interpretation of the dependence between prior volatility changes and posterior changes in the return autocorrelation is more in line with the majority of the empirical literature reviewed in Section 1, relating negative return autocorrelation on a particular time-scale to overreaction on a smaller time-scale. As such, we propose a micro-founded interpretation, where the return autocorrelation is an emerging phenomenon of the interaction between many individual investors, in contrast to the interaction of two homogeneous groups in the Sentana-Wadhwani model.

The results of Section 5 have shown that changes in the local daily return autocorrelation are best explained by prior changes in local volatility. Neither illiquidity, nor transaction volume has a similar explanatory power. The phenomenon is found for individual stocks, but the impact of volatility is greatly amplified if cross-sectional averages are considered. As volatility is a sign of uncertainty, we conclude that a sense of certainty, or the lack of it, qualitatively changes the behavior of investors. The statistical insignificance of liquidity and volume, and the use of closing-mid prices leads us to reject mechanical origins, in contrast to studies performed in the early 90s.

As our analysis is applied to changes with respect to the local mean of the variables, a large positive value in $\langle \sigma_{ohlc} \rangle_\mu$ represents a sudden increase of market volatility, not only a high level of absolute volatility but a shock to the system. Due to the cross-sectional averaging, the shock has a market-wide nature, which is pointing to a global, instead of a company specific, crisis. If the higher levels of volatility are persistent for several months that follow, indicating consistent uncertainty about fundamental values and the possibility for an approaching change of regime, investors will start doubting their standard pricing mechanisms and will be looking for alternative or additional sources of information to form their trading decisions. By doing so, they will become more susceptible to the opinion of their collegues, who might have additional information. They can obtain this information either by direct interaction with their peers, leading to herding and possible informational cascades (Bikhchandani et al., 1992), or by looking at recent market moves, which represent the opinion of all active investors, thus leading to trend-following behavior. As during highly volatile periods, markets move fast[14], it is important for investors to gather this information as fast as possible, leading to overreaction to news or continuation of price-movements on an intra-day time-scale. The excess demand (supply) during one day will result in an overpriced (underpriced) asset at the close of the market, making the reversal more likely the next day. As this reversal on the following day engenders again trend-following by the same mechanism, the correction will be also amplified, resulting in an underpriced (overpriced) asset. It is this daily over- and under-shooting of the price, due to intra-day herding, which we propose as the origin of the negative daily return autocorrelation. An argument in

[14] Tautologies are tautological.

favor of this hypothesis is the robustness of the results, when the return autocorrelation is not based on close-to-close returns but on open-to-close returns, as reported by Figure 10. This confirms that the phenomenon is still apparent if the close-to-open price changes is disregarded and that indeed it is the trading activity during the day that pushed the price away from its “fundamental” value.

A decrease of volatility is followed by an increase of persistence in the price dynamics. Here the known reasoning in terms of “underreaction”, as discussed in the introduction, is valid. A decrease of volatility is a sign of reduced uncertainty asociated with the interpretation by traders that the apparent problem has been solved, leading to a smooth transition to the new “fundamental ” price. Simultaneously, the apparent reduction of uncertainty can lead to overconfidence of investors, who are not often revising their opinion and keep following the current trend, possibly leading to the growth of a financial bubble. In other words, the overreaction is now working on a different time-scale.

An other way of describing the same interprestion, is to state that volatility changes are positively related to the frequency of information-updates of investors. In highly volatile periods, the information-update-mechanism, which leads to momentum in the returns, works on small time-scales, leading to intra-day-momentum. Whereas in periods of low volatility, the momentum-generating process is active on larger time-scales as investors have more time to take their decision, leading to persistence in day-to-day returns.

7 Robustness of the results

In this section, we will investigate the robustness of our results and show that they persist over time and for independent companies, are insensitive to the employed statistical measures and can even be recovered without any non-stationarily demeaning (albeit with smaller statistically significance).

7.1 Robustness of firm selection

The discussion of Figure 2 and 3 in Section 5 was based on the cross-sectional average over all available stocks, displayed by the dashed black lines. Comparing this line with the colored lines, which represent the cross-correlation of the time-series averaged over capitalization groups, indicates that the dependencies also hold for subsets of all available companies. The similarities between the colored lines is a strong argument for the generality of the phenomena, as the different lines are obtained from non-intersecting sets of stocks, and thus are independent from each other. For all panels in Figure 2, the capitalization of the company seems to have only a negligible impact on the reported phenomenon, with the minimum in Figure 2(a) being only slightly less pronounced for the group with the highest and lowest capitalizations. The same can be said for the maximum in Figure 2(b). The differences between the capitalization groups are smallest for Figure 2(d), which displays the well-known “leverage effect”. On the other hand, the weak similarities of the full lines in Figure 3 confirm the minor importance of the liquidity and volume on return autocorrelations.

7.2 Robustness over time

An alternative evidence for the robustness of results is their persistence over time. In Figure 7, the same kind of analysis is performed as in Figure 2 but, instead of computing the cross-correlations for the whole 21-year period, the analysis is performed for 3 non-overlapping seven-year periods. The periods, specified in the legend, range from January 1 of the beginning year to January 1 of the ending year, except for 2012, which only ends September 9, 2011 and have been labeled 2012 for convinence[15]

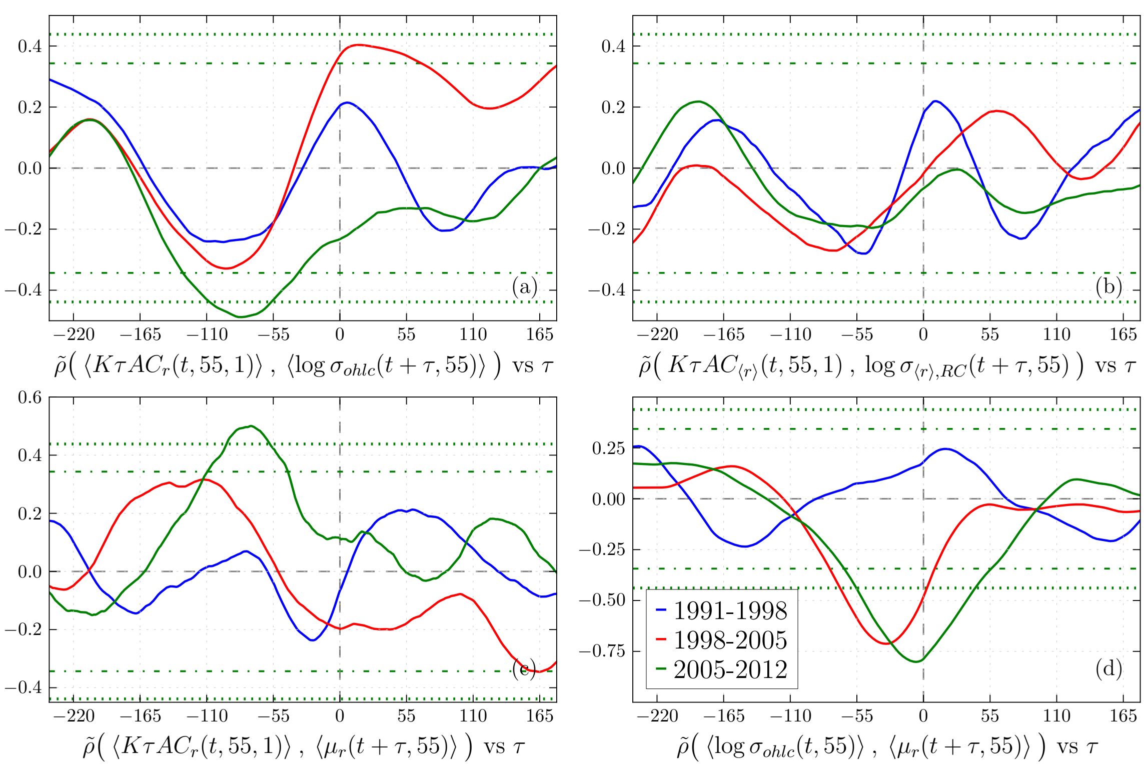

Figure 7: Cross-correlation analysis similar to Figure 2, but performed over three non-overlapping seven-year periods. The periods, specified in the legend, range from January 1 of the beginning year to January 1 of the ending year, or September 9, 2011 in the case of 2012. Only results for time-series of overall cross-sectional averages are shown, capitalization groups plots can be obtained from the author upon request. Figure 8 shows the same analysis but with 4-year periods instead of 7-year periods.

The cross-correlation between return autocorrelation and volatility, displayed in Fig. 7(a), shows a minimum of cross-correlation at $\tau \sim -80$ for all three periods, re-enforcing the validity of the results obtained for the entire 21-year period. For $\tau > 0$, the results for the three periods do not coincide, which could be explained by two reasons. First, the dependence for $\tau > 0$ changes over time, or second, these correlations are spurious.

More details are given in Figure 8, where the analysis is repeated for 4-year periods instead of 7-year periods. From Fig. 8(a), one can deduce three observations. First, the negative cross-correlation minimum at $\tau \sim -80$ is observed for every independent period. Even though the exact position of the minimum varies, its existence cannot be doubted. Second, the cross-correlation for $\tau \sim 0$ seems to be changing in time, as it is consistently positive for the 1992-2004 and negative for 2004-2012. And third, the two periods exhibiting the strongest minimum are those where markets were in turmoil, with the dot-com bubble and crash in 2000-2004, and the housing and credit crisis in the 2008-2012 period.

[15] Most likely an additional 3 months of data would not significantly change the results.

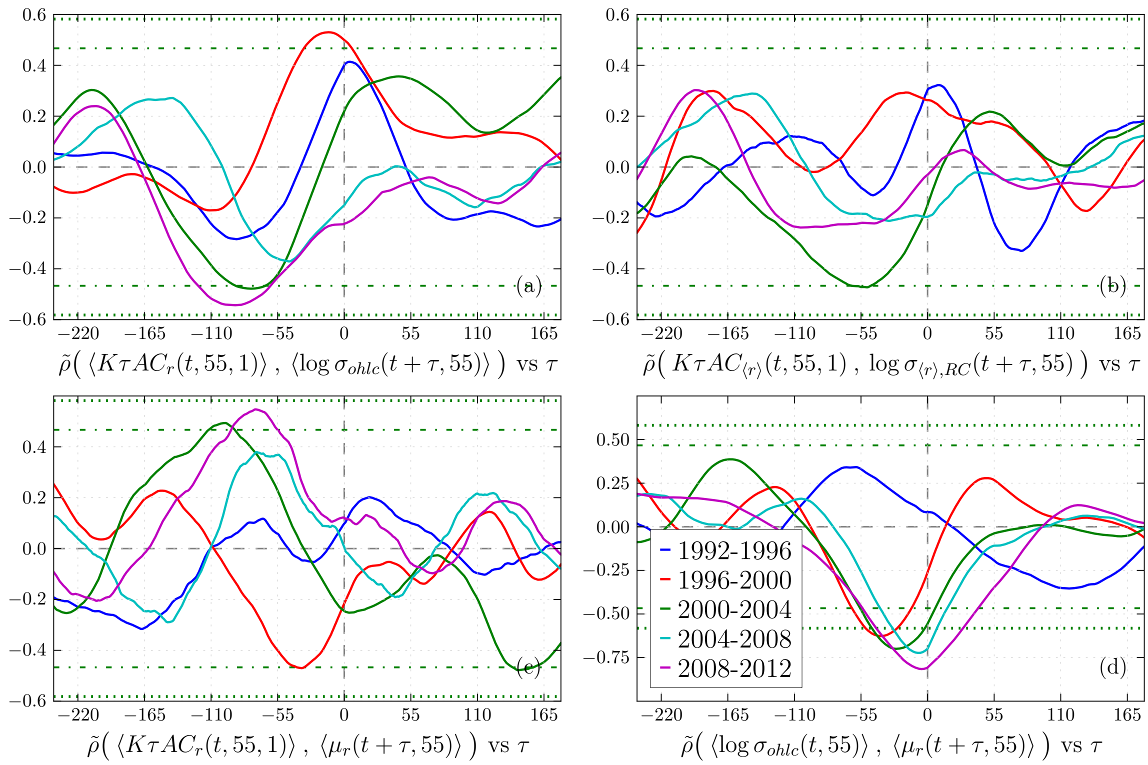

Figure 8: Same as Figure 7, but with 4-year periods instead of 7-year periods.

The positive autocorrelation between market trends and future return autocorrelation, displayed in Fig. 7(c) and Fig. 8(c) does not show the same strength of persistence as the volatility-autocorrelation relationship, as the maximum seems to be absent in the earlier periods. However for later periods, from 2000-2012, the positive impact of large-scale returns onto investors behavior seems to be persistent. Also the large-scale "leverage effect" shown in Fig. 7(d) and Fig. 8(d) is absent in the early 90s, but very well developed latter on. Here, it is important to note that the absence of this effect on the larger scale (with Δ large compared to 1) does not imply its absence at the daily scale (the effect is actually present at the daily scale).

From Figure 7(a) and 8(a), it follows that the cross-correlation values have a low statistical significance, with a majority of lines not crossing the 95-percentile bounds. The reason for this is however not that the signal is spurious or too noisy, as the minimum cross-correlation is around -0.4, similar to Fig. 2(a), but that the period over which the cross-correlation is computed is very small, such that the confidence interval becomes very large. The importance of the Figure is however not undermined by this, as the very strong statistical significance of the phenomenon is shown in Figure 2 and the objective of these figures is only to show its persistence over time.

The temporal evolution of the cross-correlations between volume and illiquidity is investigate in Figure 9. For the illiquidity, Figure 9(a) shows that no apparent similarities between the cross-correlations over the different periods can be observed, confirming the weak signal in Figure 3(a). However for the 2005-2012 period, the significant negative minimum in the cross-correlation for $\tau \sim -40$ might be an indication that the local illiquidity gained in strength, which is confirmed by its correlation with the volatility. The maximum of the volatility-illiquidity cross-correlation at

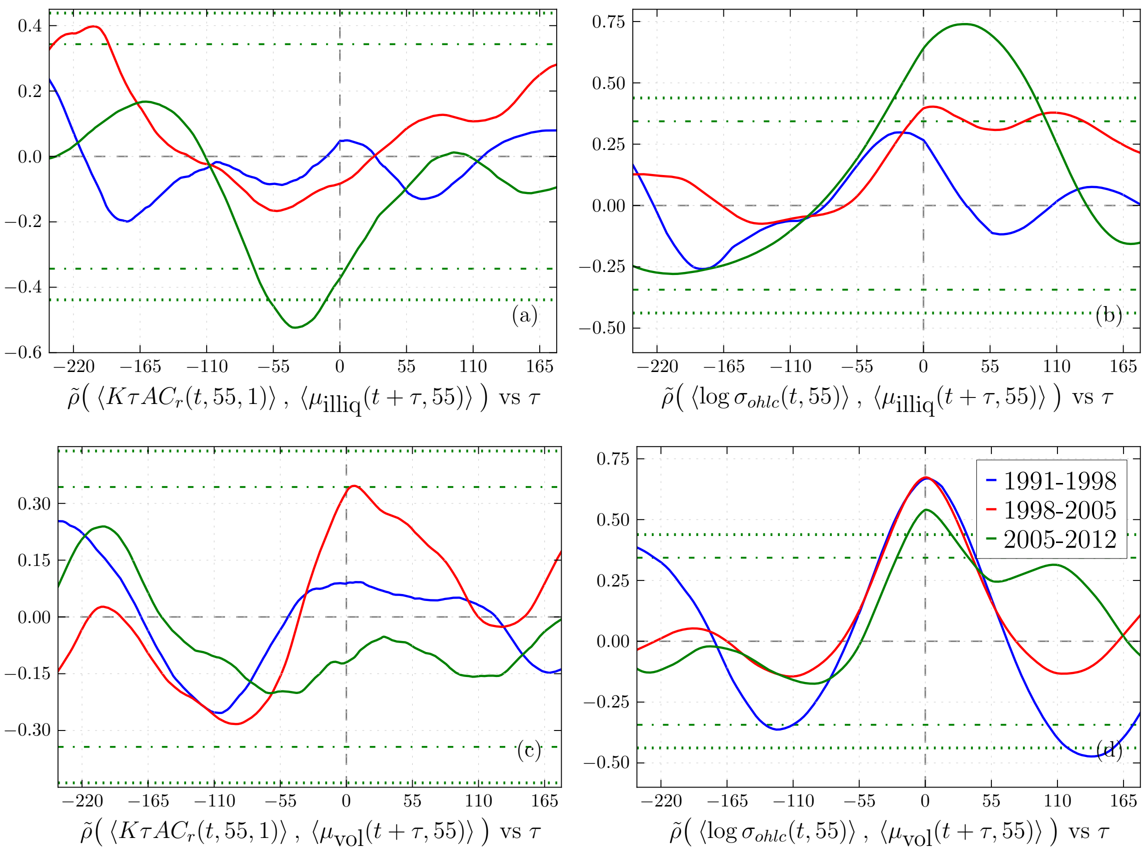

Figure 9: Volume Cross-correlation with $\Delta_{\hat{\alpha}} = 1750$ and $\delta_{\hat{\alpha}} = 100$.

$\tau > 0$ leads to the preliminary conclusion of the volatility impacting on the illiquidity, and not the other way around. Figure 9(d) shows that the volume-volatility instant correlation is stable and statistically significant over time, explaining persistent, but weak compared to the volatility, negative cross-correlation between past volume and future return autocorrelation.

7.3 Robustness with respect to the computation of returns

As all the return autocorrelations are computed based on closing-mid-price-to-closing-mid price returns and we have proposed an interpretation based on intra-day trading effects, Figure 10 investigates the dependence of the autocorrelation dynamics of open-to-close returns,

$$r_{\text{oc}}(t) = \frac{p_{\text{close}}(t) - p_{\text{open}}(t)}{p_{\text{open}}(t)}. \qquad (25)$$

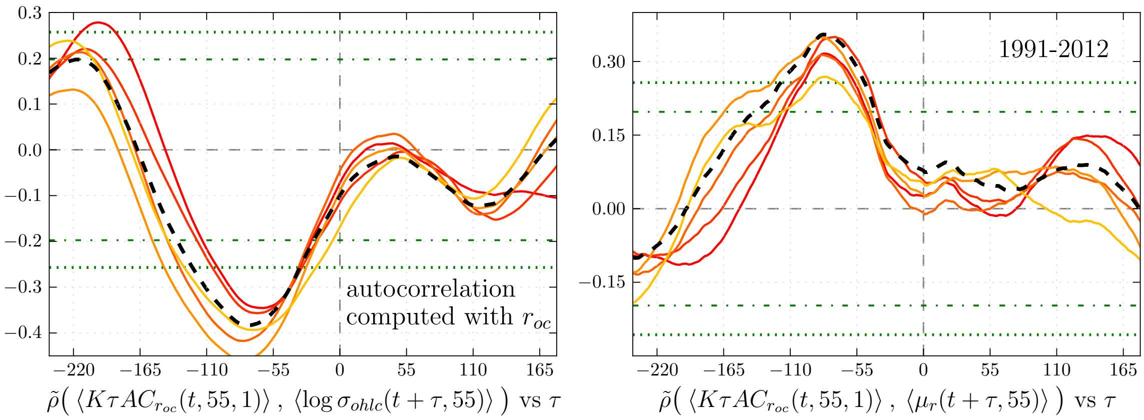

Figure 10: Same as Fig. 2(a) and (c) except that the return autocorrelation is computed based on open-to-close returns.

The recovery of the same dependencies with the same magnitude of cross-correlation is a first confirmation of our interpretation. As the phenomenon in Figure 10 disregards price changes that happen from the close of one day to the open of the exchange the next day, we can state with high confidence that the origin of the reported phenomenon is found during trading hours.

7.4 Robustness over moving window size

Up to this point, all the cross-correlation analyses were performed with $\Delta = 55$ working days and $\tau_{\hat{\alpha}} = 5$ working days. We found $\Delta = 55$ working days to be an optimal window-size, large enough to minimize the errors in the return autocorrelation estimation and small enough to give a usable temporal resolution for the cross-correlation analysis. In Figure 11 and 12, the cross-correlation analysis are repeated with a smaller ($\Delta = 35$) and a larger ($\Delta = 75$) window-size in order to confirm that the cross-correlation shapes do not depend on the specific parameters used to perform the analysis. In Figure 11 and 12, $\tau_{\hat{\alpha}}$ is adapted to the specific value of Δ, such that the maximal lag over which $\hat{\alpha}$ is optimized is approximately one year.

7.5 Robustness over statistical methods

For the previous cross-correlation analysis, all time-series where demeaned by the method introduced in Section 4. As this is a novel method, the truthfulness of the results might be doubted, interpreting them as a consequence of the unorthodox demeaning method. In order to address these concerns, Figures 13 and 14 show the cross-correlation analysis of the raw time-series, using simple methods to address the long-memory characteristics of some of the time-series.

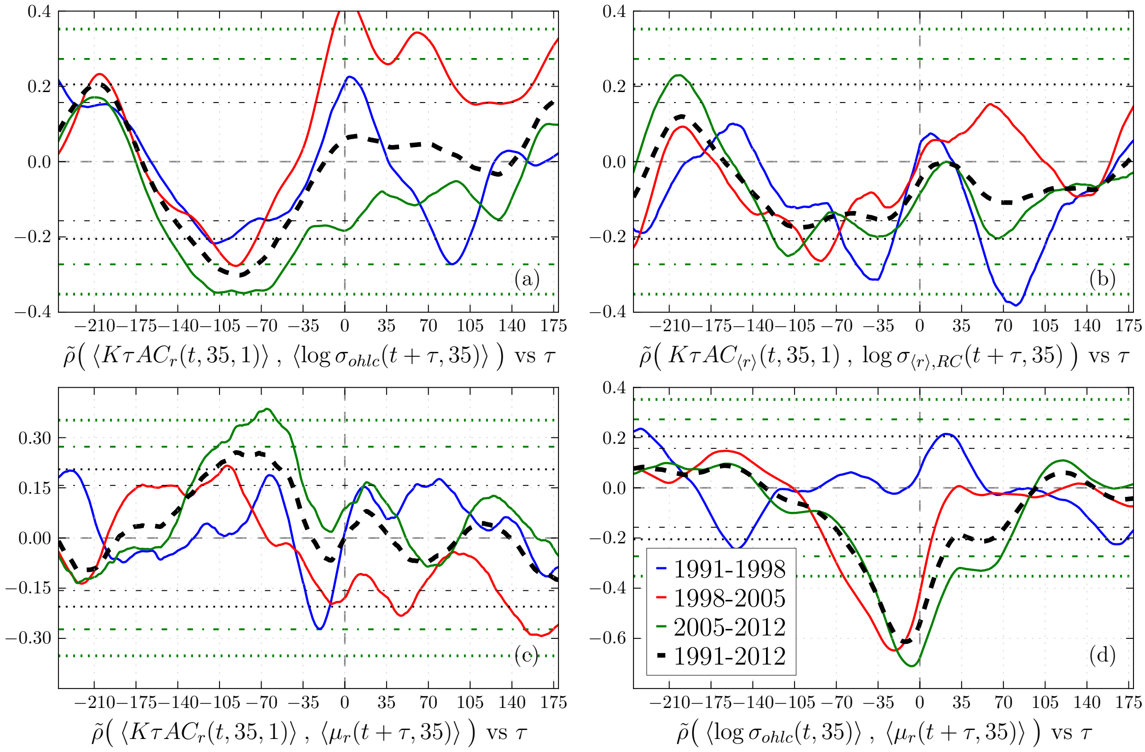

Figure 11: Same as Figure 7, but with $\Delta = 35$ working days and $\tau_{\hat{\alpha}} = 10$ instead of $\Delta = 55$ and $\tau_{\hat{\alpha}} = 5$. Cross-correlations are estimated over 7-year periods are displayed by full colored lines, with the green horizontal lines representing the 95% and 99% confidence intervals. The black dashed line is estimated over the full 21-year period, with black horizontal confidence intervals.

A standart alternative to correct for memory in the time-series is to perform the cross-correlation on their first differences. As the time-series present a memory of Δ days by construction of the windows, the first difference time-series is computed as

$$\mathrm{diff}\big(x(t,\Delta)\big) := x(t,\Delta) - x(t-\Delta,\Delta). \tag{26}$$

The full lines in Figure 13 shows the cross-correlation analysis estimated with all time-series untreated, except for the volatility, to which the first-differences transformation (Eq. (26)) is applied. The dashed lines are obtained with the first-differences applied to the volatility and return autocorrelation, leaving the average return untouched. From Figure 13(a) follows that, for any time-period and when taking the first differences of the local autocorrelation or not, the negative minimum for $\tau \sim -80$ is preserved in any setting. The exact location of the correlation minimum changes by taking the first differences of the autocorrelation as taking the first differences of a non-unit-root process distorts the time-series. However the great similitude between all the lines for $\tau < 0$ confirms that the reported phenomenon is not an artifact of the non-stationarily demeaning.

Another alternative to compute meaningful cross-correlation of long-memory processes is to estimate the correlation over a short time-scale, short enough such that the process can be considered stationary. The full lines in Figure 14 report the cross-correlations between the raw cross-sectionally averaged time-series, estimated on non-overlapping two-year periods. The black dashed lines represent the average cross-correlation over the obtained 20 cross-correlations estimated over the two-year

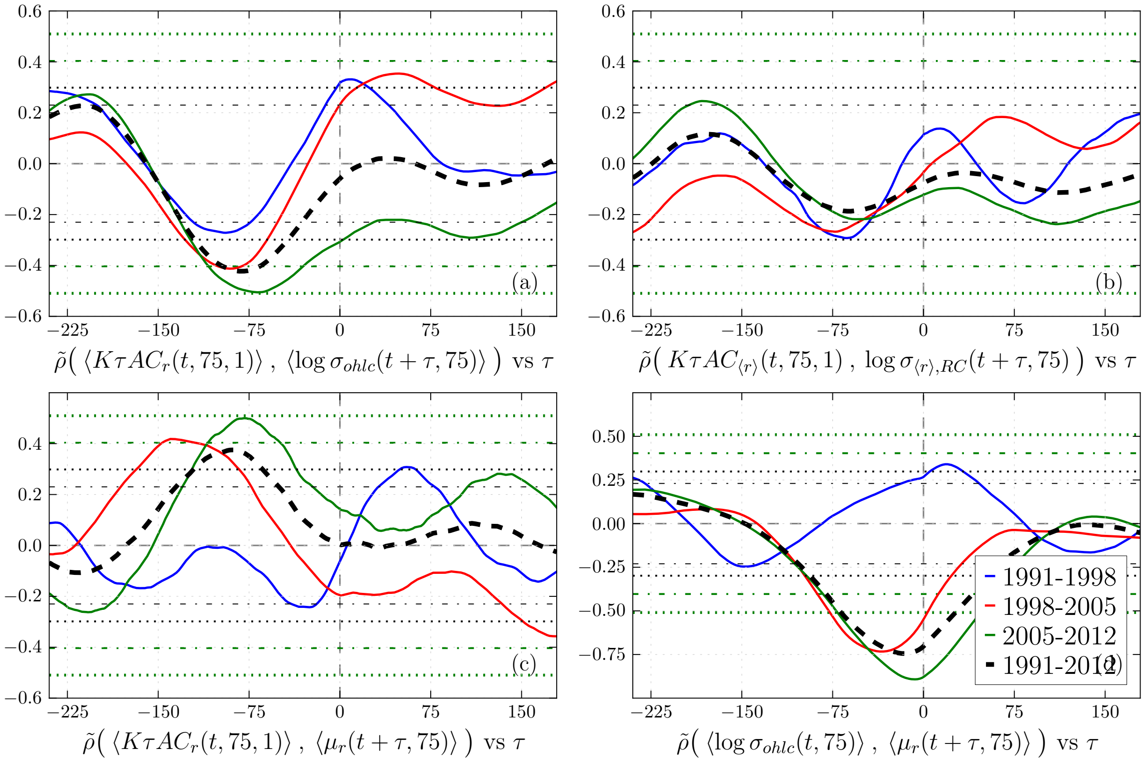

Figure 12: Same as Figure 11, but with $\Delta = 75$ working days and $\tau_{\hat{\alpha}} = 3$.

periods. As the Pearson correlation estimator is a linear measure, the 95% and 99% confidence intervals, indicated by the black horizontal lines, are based on $20 \cdot 250/\Delta$ data-points[16].

Again, as for the first-difference alternative, a retarded impact of volatility onto return autocorrelation manifests itself, even in the most parsimonious of all cross-correlation analysis. As the dashed black line in Figure 14(a) shows, not only is the volatility impact clearly visible, it is also highly statistically significant. As for Figure 14(a), Figure 14(b)-(d) do not deviate from the previously discussed results.

The distribution of the cross-correlation values, for the null-hypothesis of no serial correlation in the direction of the returns, is analyzed in Section A the appendix, giving further evidence for the validity of the used non-stationary detrending method.

7.6 Robustness over autocorrelation estimator

The measure of choice for the autocorrelation estimation in this study is the estimator based on Kendall's τ. Due to the daily returns' heteroscedastic nature and fat tailed distribution, we found this estimator best suited for our purposes. However, very similar results are found with alternative estimators, such as the Pearson autocorrelation or the variance ratio introduced in Section 2.

The different autocorrelation estimators are compared in Figure 15, where the cross-correlations between the local return autocorrelation and volatility, respectively market trends, are displayed. Only negligible differences are noticeable, confirming again the validity of our results.

[16] Every year has approximately 250 trading days.

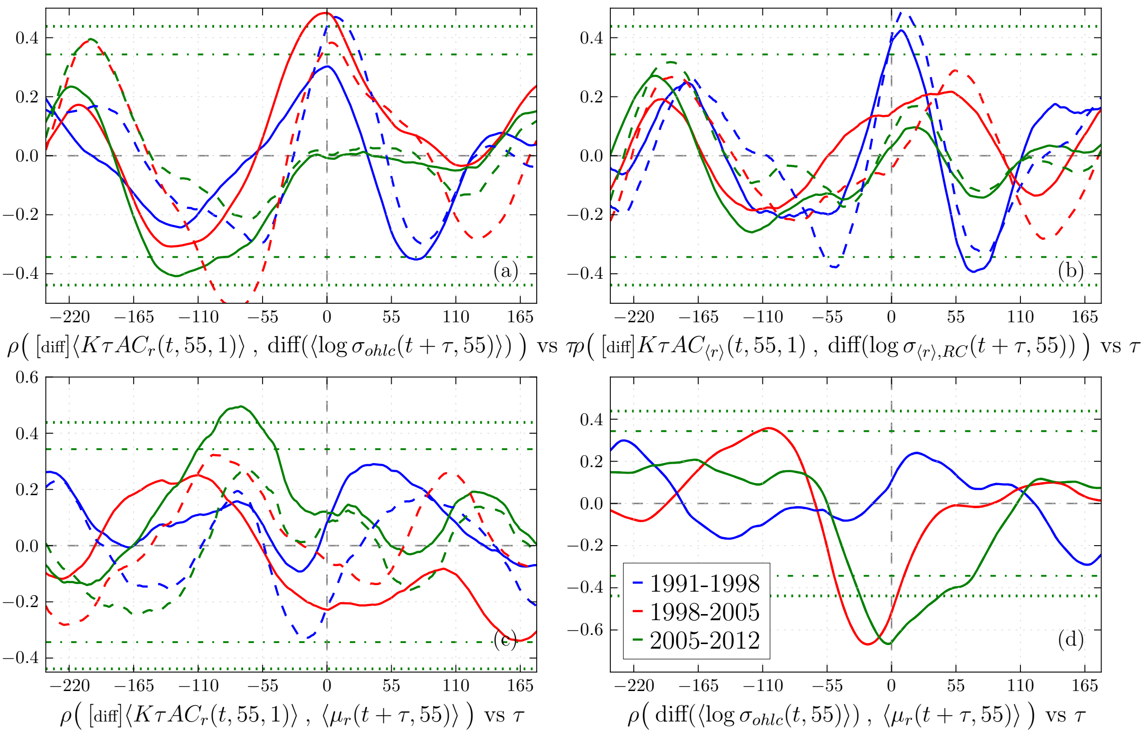

Figure 13: Cross-correlation analysis similar to the one in Figure 7, but performed on raw time-series. For the full lines, the first differences of the volatility is correlated with the untreated autocorrelation and average return. The dashed lines are obtained by applying the first-differences to the volatility and local return autocorrelation. The three colors represent the results for three non-overlapping 7-year periods, specified in the legend.

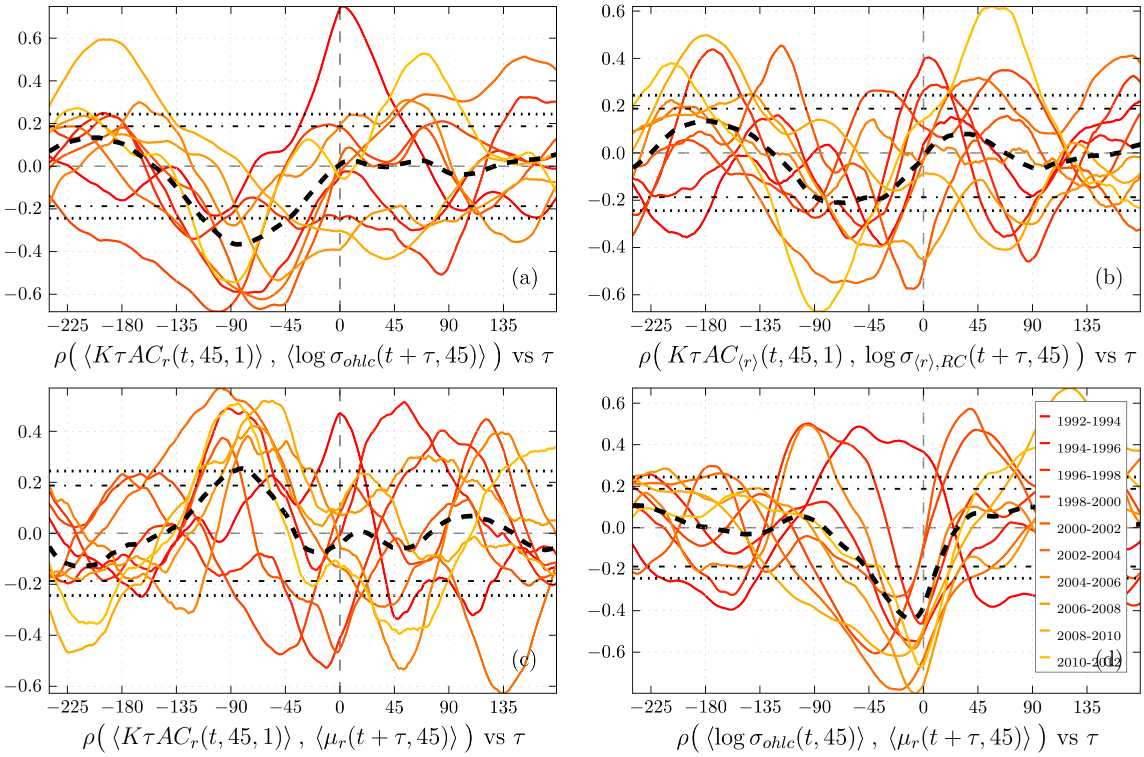

Figure 14: Cross-correlations between the raw cross-sectionally averaged time-series (no demeaning performed on them), estimated on non-overlapping two-year periods. The black dashed lines represent the average cross-correlation over the 20 two-year periods and the horizontal black lines indicate the 95% and 99% confidence intervals.

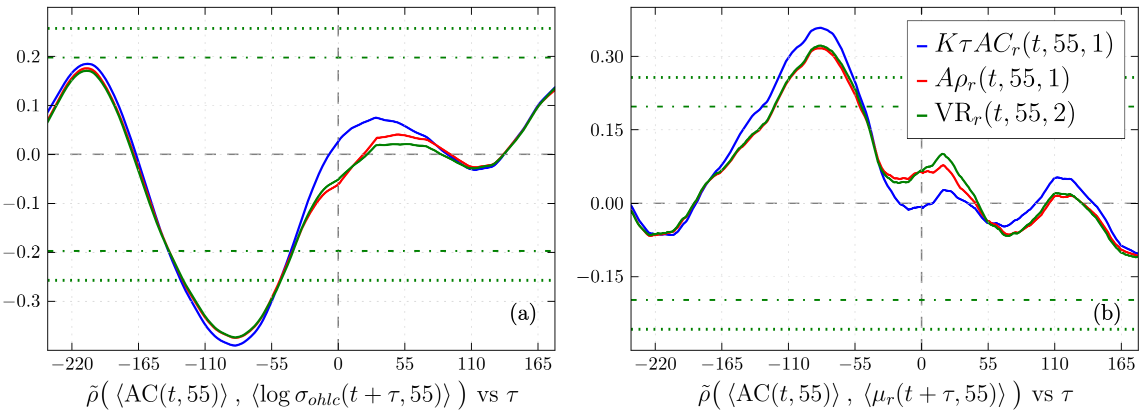

Figure 15: Same as the black dashed line of Fig. 2(a), resp. Fig. 2(c), with the addition of the return autocorrelation time-series estimated with two alternative autocorrelation estimators, specified in the legend. The autocorrelation estimator are based on: Kendall's τ ($K\tau AC_r$) in blue, Pearson's correlation ($A\rho_r$) in red and variance ratio (VR_r) with $q = 2$ in green. The cross-correlation is estimated over the full 21-year period with $\Delta = 55$ and $\tau_{\hat{\alpha}} = 5$. As the lines almost perfectly overay each other, the robustness in respect to the autocorrelation estimators is confirmed.

8 Future Works

This study being the first to investigate the lagged effects of volatility changes onto price dynamics and consequently investors behavior, there are multiple possibilities for extending the reported results. As our proposed interpretation is based on intra-day herding as a reaction to increased volatility, an analysis based on high-frequency data seems an obvious extension. Another interesting question would be to differentiate between price reversals after positive or negative daily returns and their dependence on past volatility, similarly to the study by Nam et al. (2006).

An alternative extension would be the implementation of simple trading strategies based on the here presented results, as the presence of return autocorrelation in financial returns by itself is not enough to reject the EMH. Only if this correlation leads to a profitable trading strategies with abnormal risk-adjusted returns can a rejection of the EMH be claimed.

As this study only investigates the daily returns structure, a generalization of the time-scale of the returns would also be highly interesting. Especially as the impact of the volatility is symmetric, the trending of the price after a volatility decrease will also lead to an over- or under-valuation of the company. This pricing error will at some point be corrected, creating a negative autocorrelation of returns over larger time-scales. As a result, the dependence of the local autocorrelation of returns at different time-scales might be very different from the presented results for daily returns. That there extist a wealth of momentum and reversal phenomena at lower frequencies is confirmed by the literature review in Section 1 and other recent studies, such as the one by Gutierrez and Kelley (2008), where weekly returns are investigated and an interesting observation of reversal followed by momentum is reported.

9 Conclusion

The impact of volatility changes and market trends on daily return autocorrelation is investigated in this paper. A very strong negative relation between market-wide volatility changes and market-wide daily return autocorrelation 3-4 months later is found. The robustness of the phenomenon is phenomenal, as it is found in every analyzed sub-period and across all kinds of stocks. The phenomenon is also not sensitive to estimation parameters or statistical methods, and it is even found for autocorrelations based on open-to-close returns.

For the dependency of the market-wide price trend on the market-wide daily return autocorrelation, a slightly weaker but still very strong positive relation is observed, also with a lag of 3-4 months. The weaker correlation is explained by the smaller robustness of that observation, which is persistent over the last 12 years.

As neither illiquidity nor volume are strongly related to the return autocorrelation dynamics, we can rule out a mechanical origin of the phenomenon and propose a behavioral explanation of this effect. We state that long periods of increased volatility and falling prices diminish investors' confidence and lead to insecurities regarding their trading decisions. As a result, investors try to gather more information, either from direct interaction with their peers, leading to herding, or by looking at recent market moves, which represent the opinion of all active investors, thus leading to trend-following behavior. In regimes of high volatility, this information-gathering process is performed through-out their daily trading activity, such that investors pursue intra-day trend-following strategies, which leave the stock price over- or under-valued at the end of the day. This over- or under-valuation increases the probability for a price correction on the following day, which due to the same positive feedback mechanisms will under- or over-shoot the stocks' "fundamental" value again. The fact, that the same results are found for the autocorrelation of open-to-close return, is a strong indication, that the origin of the volatility-return-autocorrelation dependency can be found in the intra-day behavior of traders.

It is found that the effect is approximately symmetric, such that for a decrease of volatility (an increase of prices), an increase of momentum in daily price dynamics is observed. This means that information takes more time to get absorbed by the price and that a trend following mechanism of a longer time-scale is at work. This points to an exaggerated increase of investors confidence, which clouds their view and leaves them unaware of a company's overvaluation.

This study is inconclusive in respect to the possibility of rejecting the EMH, as no trading strategies based on our results has been implemented. But even in the (likely) case of no abnormal profits, the importance of our results is not diminished, as our focus is on the identification of a volatility induced change of investor behavior, which is present, and persistent over at least 21 years of data, independent of an opportunity for profit or not.

Acknowledgement: We are greatful to Vladimir Filimonov for constructive discussions and Peter Cauwels for critical feedback on the manuscript.

A Distribution of the cross-correlations under the Null-Hypothesis

In section 5, it was stated that the confidence intervals given in this study are wider then necessary, as the moving-window time-series were not only sampled every Δth time-step but continuously, increasing the number of available data points.

Here, we will investigate the distribution of the cross-correction values for the null-hypothesis of no serial correlation in the direction of daily price movements. This will be done by repeating the analysis of Figure 2, with the only difference being that the returns are multiplied by random signs before the autocorrelation is computed, everything else being equal. By multiplying the returns by random signs, only the directional information is destroyed, keeping the heteroscedasticity intact[17]. A further consequence of this null-hypothesis cross-correction distribution generation is that it adds to the evidence that the reported correlations are not simple artifacts of the non-stationary demeaning. The obtained percentiles of the distributions as a function of the lags are plotted in Figure 16.

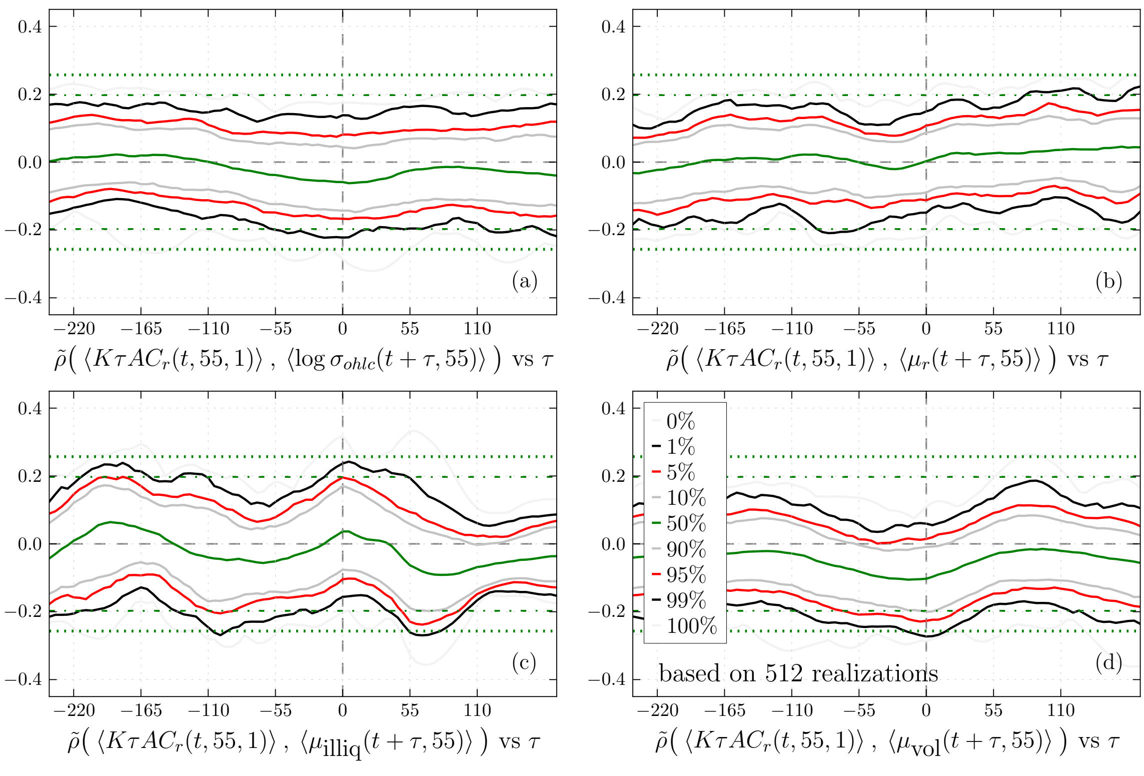

Figure 16: Percentiles of the cross-correlation distributions as a function of the lags. The cross-correlations are estimated in the same fashion as Fig. 2 with the only difference being that the returns are multiplied by random signs, before the local return autocorrelation is computed. The percentiles are based on 512 realizations of randomized signed returns.

From Figure 16, the exaggeration of the confidence bounds is clearly visible, as for most of the

[17] This procedure is less invasive than reshuffling the returns.

lags, the 99%-percentile of the cross-correlations with random signed return autocorrelations lie inside of the 95%-percentile bounds as used throughout this study.

The absence of any clear extrema in the cross-correction is also a further confirmation of the validity of our results, as the sign of the daily returns seems to be the crucial information to their recovery.

B Bloomberg Mnemonics details

PX_BID: If the market is closed, this will return the last bid from the last day the market was open. If the market is open, and there is not a bid in the market, this will return 'N.A.'

PX_ASK: If the market is closed, this will return the last ask from the last day the market was open. If the market is open, and there is not an ask in the market, this will return 'N.A.'

PX_LAST: Returns the last price provided by the exchange. For securities that trade Monday through Friday, this field will be populated only if such information has been provided by the exchange in the past 30 trading days. For all other securities, this field will be populated only if such information was provided by the exchange in the last 30 calendar days. This applies to common stocks, receipts, warrants, and real estate investment trusts (REITs).

CUR_MKT_CAP: Current monetary value of all outstanding shares stated in the pricing currency. Capitalization is a measure of corporate size. Current market capitalization is calculated as: Current Shares Outstanding * Last Price, Where 'Current Shares Outstanding' is DS124, EQY_SH_OUT, and 'Last Price' is PR005, PX_LAST. For companies which trade on multiple regional exchanges, the Composite Ticker is used in the calculation of the market cap.

PX_VOLUME: Trading benchmark calculated by adding up the value traded for every transaction (price times shares traded) and then dividing by the total shares traded for the time period. Volume Weighted Average Price (VWAP) can be used in conjunction with the VWAP Start Time (PR313, VWAP_START_TIME) and VWAP End Time (PR314, VWAP_END_TIME) overrides.

C Additional Material

	Estimate	Std. Error	t value	Pr(>|t|)
Intercept	0.0433	0.0180	2.41	0.0265
Slope	-2.2995	0.8118	-2.83	0.0106

Table 2: Results from the linear regression of mean of the yearly Pearson's autocorrelation unto the OHLC-volatility median. Residual standard error: 0.026 on 19 degrees of freedom Multiple R-squared: 0.30, p-value: 0.01

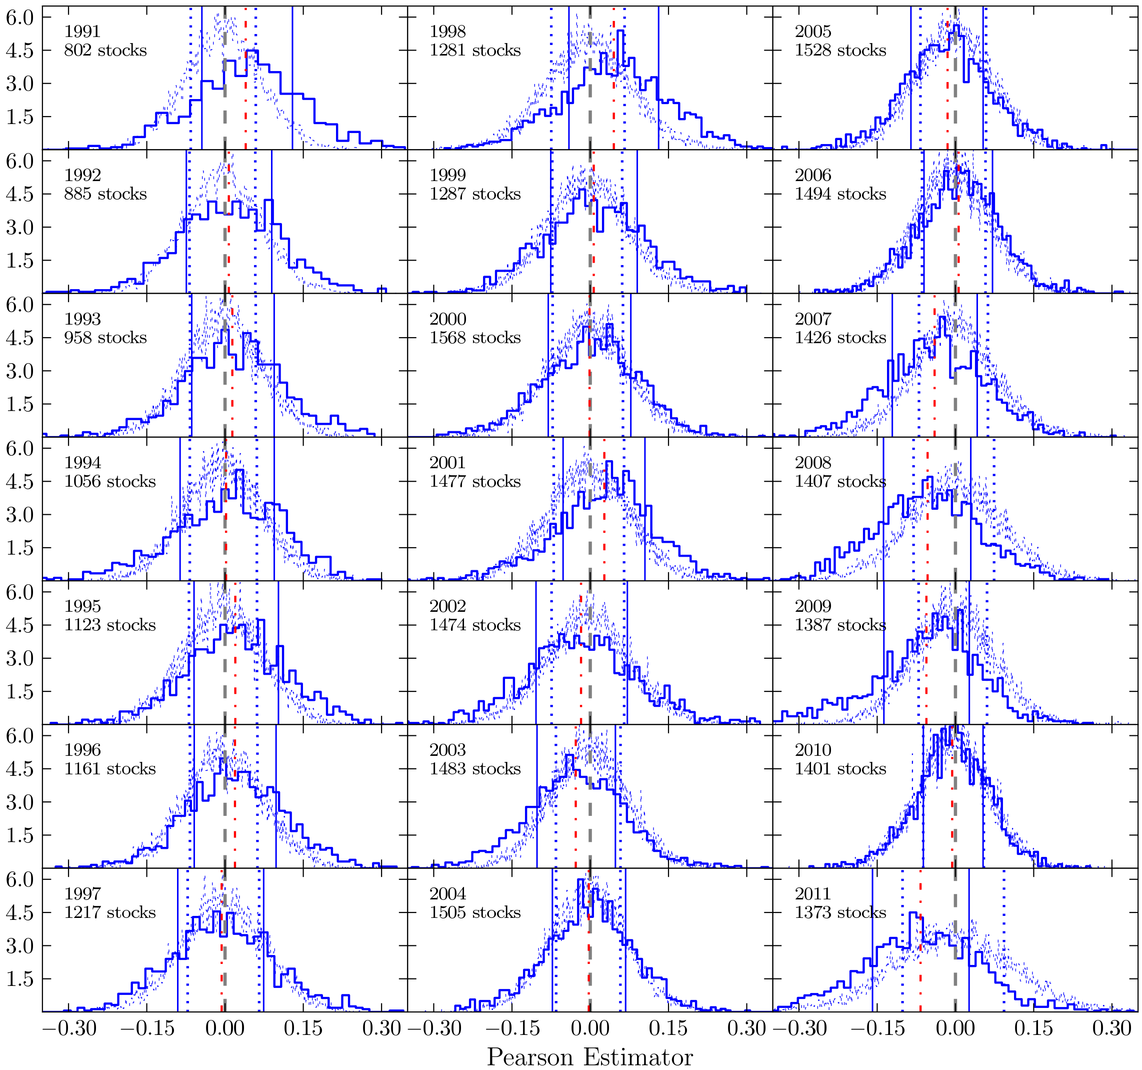

Figure 17: Same as Fig. 1, with the only difference being that the autocorrelation is estimated using Pearson's autocorrelation.

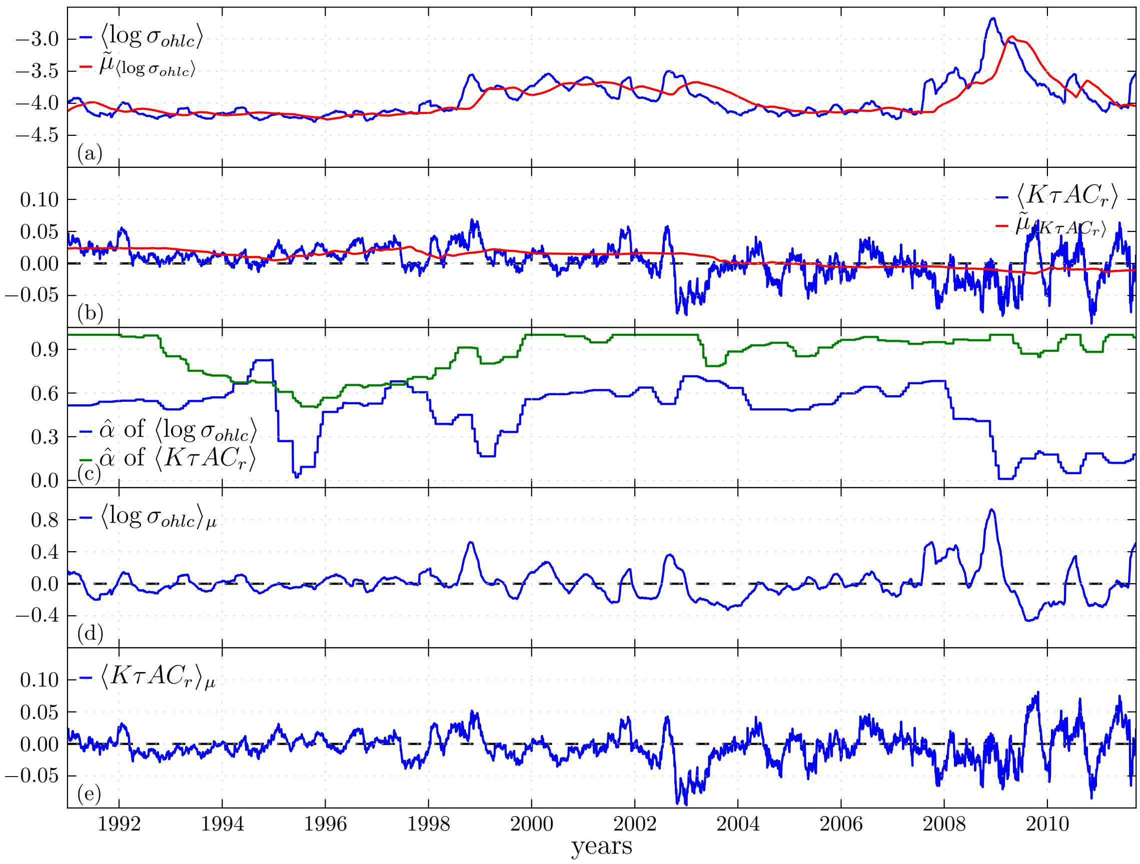

Figure 18: Time-series of the volatility and local return autocorrelation, estimated with $\Delta = 55$. The non-stationary detrending is performed with $\Delta_{\hat{\alpha}} = 32\Delta \sim 1750$ working days, $\delta_{\hat{\alpha}} = 2\Delta \sim 100$ working days, $\tau_{\hat{\alpha}} = 5$. Panel (a) (resp. (b)) show the cross-sectional average of the volatility (resp. of the daily return autocorrelation) and its local average. Panel (c) displays their respective optimal α and Panel (d) and (e) display the detrended volatility and return autocorrelation, which are used to compute the cross-correlations of Fig 2.

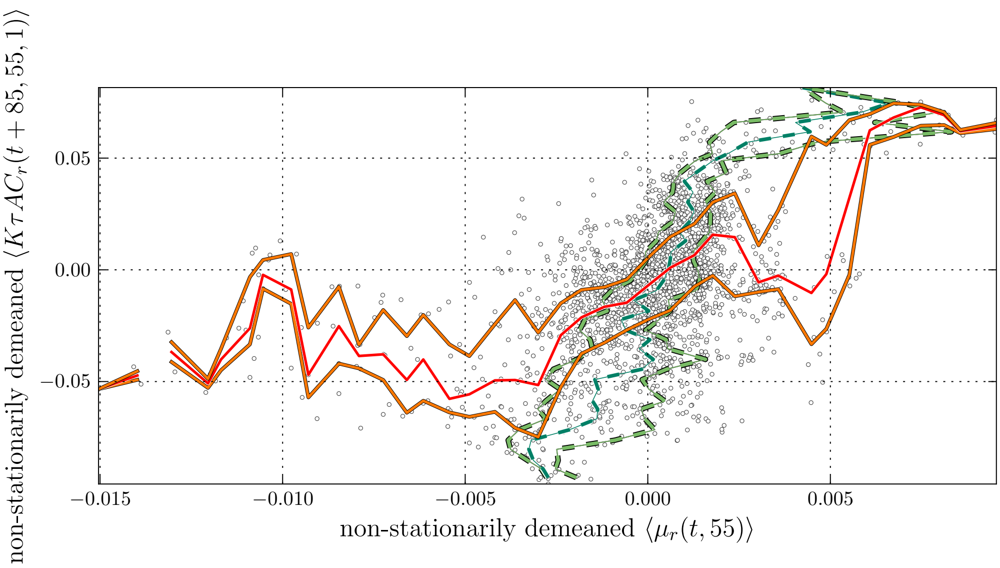

Figure 19: Similar to Figure 6 with the difference that the data plotted here ranged from January 1 2002 to September 9 2011.

References

Amihud, Y., Jan. 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5 (1), 31–56.

Avramov, D., Chordia, T., Goyal, A., 2006. Liquidity and Autocorrelations in Individual Stock Returns. The Journal of Finance 61 (5).

Ball, R., Kothari, S. P., Wasley, C. E., 1995. Can We Implement Research on Stock Trading Rules?: The Journal of Portfolio Management.

Bianco, S., Corsi, F., Renò, R., Jul. 2009. Intraday LeBaron effects. Proceedings of the National Academy of Sciences 106 (28), 11439–11443.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. The Journal of Political Economy 100 (5), 992–1026.

Black, F., 1976. Studies of stock price volatility changes. In: Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section. pp. 177–181.

Bohl, M. T., Siklos, P. L., Sep. 2008. Empirical evidence on feedback trading in mature and emerging stock markets. Applied Financial Economics 18 (17), 1379–1389.

Bondt, W. F., Thaler, R., 1985. Does the Stock Market Overreact? The Journal of Finance 40 (3), 793–805.

Booth, G. G., Koutmos, G., 1998. Volatility and autocorrelation in major European stock markets. European Journal of Finance 4 (1), 61–74.

Campbell, J. Y., Grossman, S. J., Wang, J., Nov. 1993. Trading Volume and Serial Correlation in Stock Returns. The Quarterly Journal of Economics 108 (4), 905–939.

Campbell, J. Y., Lo, A. W., MacKinlay, A. C., Lo, A. Y., Dec. 1997. The econometrics of financial markets. Princeton University Press.

Chopra, N., Lakonishok, J., Ritter, J. R., Apr. 1992. Measuring abnormal performance. Journal of Financial Economics 31 (2), 235–268.

Chordia, T., Roll, R., Subrahmanyam, A., 2001. Market Liquidity and Trading Activity. The Journal of Finance 56 (2).

Chordia, T., Roll, R., Subrahmanyam, A., Feb. 2008. Liquidity and market efficiency. Journal of Financial Economics 87 (2), 249–268.

Conrad, J., Gultekin, M. N., Kaul, G., 1997. Profitability of Short-Term Contrarian Strategies: Implications for Market Efficiency. Journal of Business & Economic Statistics 15 (3).

Conrad, J. S., Hameed, A., Niden, C., 1994. Volume and Autocovariances in Short-Horizon Individual Security Returns. The Journal of Finance 49 (4).

Croux, C., Rousseeuw, P. J., 1992. Time-efficient algorithms for two highly robust estimators of scale. Computational Statistics 1.

Cutler, D. M., Poterba, J. M., Summers, L. H., 1990. Speculative Dynamics and the Role of Feedback Traders. The American Economic Review 80 (2), 63–68.

Daniel, K., Hirshleifer, D., Subrahmanyam, A., 1998. Investor Psychology and Security Market Under- and Overreactions. The Journal of Finance 53 (6), 1839–1885.

Ding, Z., Granger, C. W. J., Engle, R. F., Jun. 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1 (1), 83–106.

Fama, E. F., May 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. The Journal of Finance 25 (2), 383–417.

Fama, E. F., French, K. R., 1988. Permanent and Temporary Components of Stock Prices. Journal of Political Economy 96 (2).

Ferguson, T. S., Genest, C., Hallin, M., 2000. Kendall's Tau for Serial Dependence. The Canadian Journal of Statistics / La Revue Canadienne de Statistique 28 (3), 587–604.

French, K. R., Roll, R., Sep. 1986. Stock return variances: The arrival of information and the reaction of traders. Journal of Financial Economics 17 (1), 5–26.

Gallant, A. R., Rossi, P. E., Tauchen, G., Apr. 1992. Stock prices and volume. Review of Financial Studies 5 (2), 199–242.

Granger, C. W. J., Newbold, P., Jul. 1974. Spurious regressions in econometrics. Journal of Econometrics 2 (2), 111–120.

Grossman, S. J., Miller, M. H., 1988. Liquidity and Market Structure. The Journal of Finance 43 (3), 617–633.

Gutierrez, R. C., Kelley, E. K., Feb. 2008. The Long-Lasting Momentum in Weekly Returns. The Journal of Finance 63 (1), 415–447.

Hamilton, J. D., Jan. 1994. Time Series Analysis, 1st Edition. Princeton University Press.

Hasanhodzic, J., Lo, A. W., 2011. Black's Leverage Effect is not Due to Leverage. SSRN eLibrary.

Hens, T., Steude, S. C., 2009. The leverage effect without leverage. Finance Research Letters 6 (2), 83–94.

Hong, H., Stein, J. C., 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. The Journal of Finance 54 (6), 2143–2184.

Jegadeesh, N., 1990. Evidence of Predictable Behavior of Security Returns. The Journal of Finance 45 (3), 881–898.

Jegadeesh, N., Titman, S., Mar. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. The Journal of Finance 48 (1), 65–91.

Jegadeesh, N., Titman, S., Apr. 1995. Short-Horizon Return Reversals and the Bid-Ask Spread. Journal of Financial Intermediation 4 (2), 116–132.

Jones, C. M., Kaul, G., Lipson, M. L., Oct. 1994. Transactions, volume, and volatility. Review of Financial Studies 7 (4), 631–651.

Karpoff, J. M., 1987. The Relation Between Price Changes and Trading Volume: A Survey. The Journal of Financial and Quantitative Analysis 22 (1).

Kaul, G., Nimalendran, M., Nov. 1990. Price reversals. Journal of Financial Economics 28 (1-2), 67–93.

Kendall, M. G., 1938. A New Measure of Rank Correlation. Biometrika 30 (1/2).

Koutmos, G., Aug. 1997. Feedback trading and the autocorrelation pattern of stock returns: further empirical evidence. Journal of International Money and Finance 16 (4), 625–636.

Kruskal, W. H., 1958. Ordinal Measures of Association. Journal of the American Statistical Association 53 (284).

LeBaron, B., 1992. Some Relations Between Volatility and Serial Correlations in Stock Market Returns. The Journal of Business 65 (2).

Lehmann, B. N., Apr. 1990. Fads, Martingales, and Market Efficiency. National Bureau of Economic Research Working Paper Series, 2533+.

Liu, Y., Gopikrishnan, P., Cizeau, Meyer, Peng, Stanley, E. H., 1999. Statistical properties of the volatility of price fluctuations. Physical Review E 60 (2), 1390+.

Lo, A. W., MacKinlay, A. C., Jan. 1988. Stock market prices do not follow random walks: evidence from a simple specification test. Review of Financial Studies 1 (1), 41–66.

Lo, A. W., MacKinlay, A. C., 1990. When are Contrarian Profits Due to Stock Market Overreaction? The Review of Financial Studies 3 (2).

Nam, K., Kim, S.-W., Arize, A., Nam, K., Kim, S.-W., Arize, A., Mar. 2006. Mean Reversion of Short-Horizon Stock Returns: Asymmetry Property. Review of Quantitative Finance and Accounting 26 (2), 137–163.

Poterba, J. M., Summers, L. H., Jul. 1989. Mean Reversion in Stock Prices: Evidence and Implications. National Bureau of Economic Research Working Paper Series, 2343+.

Richards, A. J., 1997. Winner-Loser Reversals in National Stock Market Indices: Can They be Explained? The Journal of Finance 52 (5), 2129–2144.

Samuelson, P. A., 1965. Proof that Properly Anticipated Prices Fluctuate Randomly. Industrial Management Review 6, 41–49.

Sentana, E., Wadhwani, S., 1992. Feedback Traders and Stock Return Autocorrelations: Evidence from a Century of Daily Data. The Economic Journal 102 (411), 415–425.

Venetis, I. A., Peel, D., Jan. 2005. Non-linearity in stock index returns: the volatility and serial correlation relationship. Economic Modelling 22 (1), 1–19.

Yang, D., Zhang, Q., 2000. Drift-Independent Volatility Estimation Based on High, Low, Open, and Close Prices. Journal of Business 73 (3), 477–91.

Yule, G. U., 1926. Why do we Sometimes get Nonsense-Correlations between Time-Series?–A Study in Sampling and the Nature of Time-Series. Journal of the Royal Statistical Society 89 (1).

6.3 Finalizing Comments

Besides the puzzling long lag of 3-4 month between the volatility and the return autocorrelation changes, a most interesting fact about the phenomena reported in the previous section, is that they have not been found or published before, given their persistence over more than the last two decades and the easy availability of the data[2]. One explanation might be the non-stationary character to the time-series, which makes the analysis less straight forward with general econometrics software. Another one might be the long lag, which is unlikely to be predicted from first principles by a theoretical model on investor behavior.

As is stated in the paper, these findings open the door to many further analyses. The most direct ones being based on, either constructing trading strategies, which exploit the predictably of the return autocorrelation, or investigating intra-day data to confirm or reject the proposed mechanism behind the overreaction.

Note that the mechanism responsible for the intra-day overreaction, described in Section 6.2, is closely related to the volatility generating mechanism studies in the previous chapter, specifically in Section V.B, where the NIV-phenomenon is proposed as a model of excess volatility. As in the present study, it is hypothesized that it is the perceived need for more information, triggered by uncertainty, that initiates investors' behavior of either imitating the actions of their peers or trend-following trading strategies. Due to its general setup, the NIV-model is unable to generate the negative autocorrelation, which is observed as a reaction to the increased volatility and is due to features, special to investors and financial markets. The daily price reversals originates from the fact that, after the closing of the market, investor reevaluate the valuations of their assets, observe possible deviations from their fair price estimates and react accordingly the next day, pushing the price in the other direction. The reason for which such information, concerning the fundamental value of an asset, does not impact as strongly during trading hours can be found in the diminished correlation between macroscopic dynamics and driving force, one of the characteristics of the noise- or in this case, interaction-induced volatility.

In addition to the negative relation between lagged changes of volatility and changes of the return autocorrelation, we also document a positive relation between lagged market trends and changes of the return autocorrelation. Whereas our results focus on the autocorrelation of daily returns, similar results, but for monthly returns, have been reported by Cooper et al. (2004). They study the relation between lagged three-year market trends (indices going up or down) and momentum in individual stocks. Significant profits are reported for momentum strategies in the following six months only after up-markets, not after down-markets. The momentum is then followed by a reversal in the long-run. These results are perfectly consistent with our findings, as positive autocorrelation on daily returns can result in positive autocorrelation on monthly returns and negative autocorrelation on daily returns, which points towards mean reverting prices, can result in statistical insignificant results.

It is worth mentioning that the return autocorrelation, which is the main subject of this study, is a two-point statistic (comparing the returns at two different time-steps) and that there exist high dimensional dependence measures,

[2]The Yahoo! Finance data is provided for free.

which are not covered here. An example of such a dependence is studies by Sornette and Pisarenko (2008), who analyses a process with zero autocorrelation, but a non-zero three-point dependence. Other examples of non-linear dependencies are found in technical analysis (Brock et al., 1992), or simple pattern analysis as performed by Zhang (1999) and Vandewalle et al. (2000).

Besides the novel method introduced to address the long memories of the analyzed time-series, another method for data-analysis is used, which I never saw in any other study, published or not, and which is of a striking simplicity. I am referring to the scatter-plots of Figure 5, 9 and 19, with the horizontally and vertically percentile lines, which allow to clearly distinguish between the *explanatory* and the *predictive* power of the two lagged time-series. In the case of Figure 6, the horizontal lines reveal information about the question: 'Given an change of x in volatility now, how will be the autocorrelation change in the future?', informing the research about the predictive power of the observable. The vertical lines reveal information about the question: 'Given an change of x in the autocorrelation now, how well can it be explained by past volatility changes?', analyzing the ex-post explanatory power of the observable. In a similar vein, it should also be noted not to equate correlation with regression as documented by Warren (1971), as the correlation between two time-series investigates their resemblance, whereas a well done regression from one onto the other time-series investigates how much one changes in respect to the other one.

Chapter 7

Conclusion

This thesis addresses three important aspects of price dynamics of publicly traded assets. For one, the emergence of a trending price is studied, resulting from the myopic optimization of socially influenceable investors, leading to the destabilization of the price and the growth of a bubble. My second contribution focuses on the volatility of price dynamics and, more generally, on the volatility of dynamic macroscopic observables, governed by a large number of interconnected units under the influence of a rapidly varying external forcing. Whereas the first two contributions explore the first and second moment of collective/price dynamics via theoretical studies, my third contribution is an empirical study investigating the autocorrelation of daily price returns and its dependencies on other macroscopic variables such as volatility, long-term price movements and illiquidity.

For both, the theoretical models as well as the interpretation of the phenomenon observed in the empirical study, it is assumed that the environment of the individuals influences their behavior. This environment can either be their social environment, like colleagues and neighbors, or the past price movements, which represent an aggregation of the opinions of all active investors. As it is stated in Section 3.4 and 3.5, this assumption is well confirmed by empirical studies, which show that investors are imitating each other and indulge in positive feedback trading. More generally, there is a wealth of experiments showing that humans are very susceptible to their surrounding and are likely to change their opinion, either consciously or unconsciously, to conform with the opinion of the group.

A financial market model is presented in Chapter 4, where the price dynamics and the behavior of trading agents, interconnected by their social network, are investigated. Agents invest according to their opinion on future price movements, which is based on three sources of information, (i) public information, i.e. news, (ii) the aforementioned information from their social network and (iii) private information. In order to form the best predictor of future price movements, agents are continuously adapting their trading strategy to the current market regime by weighting the news and information from their peers according to their recent predicting performance. Paradoxically, it is their myopic adaptation to the current market regime which leads to a dramatic amplification of the price volatility and the occurrence of bubbles. The origin of these large deviations from the equilibrium price are found to nucleate from the news.

A random occurrence of a sequence of same signed news pushes the price in one direction and starts the coordination process of the agents, developing into a transient collective herding regime. The positive feedback loop is created by the two dominating mechanisms (adaptation and imitation) which, by reinforcing each other, result in a transient over-valuation of the price, up to unsustainable levels. As such, crashes can be identified as a rapid price correction after an inflation of prices via a socially induced irrational exuberance. The model offers a simple reconciliation of the two opposite (herding versus fundamental) proposals for the origin of crashes within a single framework and shows that a crash is not a reaction to an extreme negative news event but a sudden correction of an unsustainable high price. More general, this model shows that even with rational and adapting agents, bubbles and crashes emerge naturally.

By reducing the complexity of the previous model, but keeping the same three basic influences, it is possible to apply this model to a very wide range of systems, generalizing the interpretation of the individual agent from an investor to any bistable entity, susceptible to its surrounding, a common and varying driving force and independent noise sources. This model, which is based on the dynamic Ising model, is a priori is a physical model but can easily be related to social systems via the equivalence between the Ising model and a discrete choice model with social interactions, as is shown in Chapter 2. Due to its general setup, an analytical treatment of the model is possible and is presented, together with numerical simulations, in Chapter 5. It is found that, independently of the shape of the driving force, strong fluctuations of the macroscopic dynamics are found for intermediate levels of noise (or of coupling, depending on the setup), a phenomenon which can be traced back to the presence of the unavoidable phase transition in such systems.

For a periodic forcing, this peak corresponds to a pronounced amplification of the signal, with a strong correlation between the macroscopic dynamics and the driving force at the optimal lag, the paradigmatic signatures of stochastic resonance. On the other hand, when the driving force is aperiodic, a similar peak appears but here the amplitude of the fluctuations exceeds by far those observed for periodic signals. Coincidental with the increase of fluctuations, the correlation between the driving force and the system dynamics is completely destroyed. This shows that even though these fluctuations are induced by the common forcing, the macroscopic dynamics has an endogenous origin. This phenomenon of noise-induced volatility contrasts with that of stochastic resonance, with the major difference being that it is not the signal, but the fluctuations that are amplified.

Moreover, this phenomenon of noise-induced volatility also constitutes a new indicator for the approaching of a phase transition, and it applies to a broader range of real-world systems due to the common setup given of a coupled system driving by an aperiodic forcing and its robustness with respect to changes in the underlying network of interactions. As an example of a system where this phenomenon can be observed, we have proposed the social system of stock markets, in which we have been able to not only explain the excess of volatility observed in stock prices, but also the apparent absence of correlation between news and price changes and the persistence of volatility during times of crises.

The last part of this thesis contains an empirical study, motivated by the question of whether investors behave differently in different market regimes. To investigate this question, the impact of volatility changes and market trends on daily return autocorrelation is investigated for a large number of individual stocks traded on the New York Stock Exchange. A very strong negative relation between market-wide volatility changes and market-wide daily return autocorrelation 3-4 months later is found. The robustness of the phenomenon is phenomenal, as it is found in every analyzed sub-period and across all kinds of stocks. The phenomenon is also not sensitive to estimation parameters or statistical methods, and it is even found for autocorrelations based on open-to-close returns. For the dependency of the market-wide price trend on the market-wide daily return autocorrelation, a slightly weaker but still very strong positive relation is observed, also with a lag of 3-4 months. The weaker correlation is explained by the smaller robustness of that observation, which is persistent over the last 12 years.

As neither illiquidity nor volume are strongly related to the return autocorrelation dynamics, we can rule out a mechanical origin of the phenomenon and propose a behavioral explanation of this effect. We state that long periods of increased volatility and falling prices diminish investors' confidence and lead to insecurities regarding their trading decisions. As a result, investors try to gather more information, either from direct interaction with their peers, leading to herding, or by looking at recent market moves, which represent the opinion of all active investors, thus leading to trend-following behavior. During high volatility regimes, this information-gathering process is performed through-out their daily trading activity, such that investors pursue intra-day trend-following strategies, which leave the stock price over- or under-valued at the end of the day. This over- or under-valuation increases the probability for a price correction on the following day, which due to the same positive feedback mechanisms will under- or over-shoot the stocks' "fundamental" value again. The fact, that the same relations are found for the autocorrelation of open-to-close return, is a strong indication, that the origin of the volatility-return-autocorrelation dependency can be found in the intra-day behavior of traders.

It is found that the effect is approximately symmetric, such that for a decrease of volatility (an increase of prices) an increase of momentum in daily price dynamics is observed. This means that information takes more time to get absorbed by the price and that a trend following mechanism over longer time-scales is at work. This points to an exaggerated increase of investors confidence, which clouds their view and leaves them unaware of a company's overvaluation.

Last but not least, as the majority of the time-series involved have long memories, bordering on non-stationarity, I have developed a powerful and yet intuitive method to address this memory, such that meaningful correlations between time-series can be computed.

Bibliography

Abreu, D., Brunnermeier, M. K., 2003. Bubbles and Crashes. Econometrica 71 (1), 173–204.

Allen, F., Morris, S., Postlewaite, A., 1993. Finite bubbles with short sale constraints and asymmetric information. Journal of Economic Theory 61 (2), 206–229.

Almazan, A., Brown, K. C., Carlson, M., Chapman, D. A., Aug. 2004. Why constrain your mutual fund manager? Journal of Financial Economics 73 (2), 289–321.

Andreassen, P. B., Kraus, S. J., 1990. Judgmental extrapolation and the salience of change. J. Forecast. 9 (4), 347–372.

Asch, S. E., 1952. Social psychology. Prentice-Hall.

Badrinath, S. G., Wahal, S., 2002. Momentum Trading by Institutions. The Journal of Finance 57 (6), 2449–2478.

Banerjee, A. V., Aug. 1992. A Simple Model of Herd Behavior. The Quarterly Journal of Economics 107 (3), 797–817.

Barberis, N., Shleifer, A., Vishny, R., Sep. 1998. A model of investor sentiment. Journal of Financial Economics 49 (3), 307–343.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. Journal of Political Economy 100 (5), 992–1026.

Black, F., 1986. Noise. The Journal of Finance 41 (3), 529–543.

Black, F., 1988. An Equilibrium Model of the Crash. In: NBER Macroeconomics Annual 1988, Volume 3. NBER Chapters. National Bureau of Economic Research, Inc, pp. 269–276.

Blackmore, S., May 2000. The Meme Machine. Oxford University Press, USA.

Blanchard, O. J., Watson, M. W., 1982. Bubbles, Rational Expectations, and Financial Markets.

Blume, L. E., 1993. The statistical mechanics of strategic interaction. Games and economic behavior 5 (3), 387–424.

Bondt, W. F., Thaler, R., 1985. Does the Stock Market Overreact? The Journal of Finance 40 (3), 793–805.

Brandenburger, A., Polak, B., 1996. When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See. The RAND Journal of Economics 27 (3), 523–541.

Brennan, M. J., 1990. Latent Assets. The Journal of Finance 45 (3), 709–730.

Brock, W., Lakonishok, J., LeBaron, B., Dec. 1992. Simple Technical Trading Rules and the Stochastic Properties of Stock Returns. The Journal of Finance 47 (5), 1731–1764.

Brock, W. A., 1993. Pathways to Randomness in the Economy: Emergent Nonlinearity and Chaos in Economics and Finance. Estudios Económicos 8 (1 (15)), 3–55.

Brock, W. A., Durlauf, S. N., 2001. Discrete Choice with Social Interactions. The Review of Economic Studies 68 (2), 235–260.

Brown, N. C., Wermers, R. R., Wei, K. D., Mar. 2009. Analyst Recommendations, Mutual Fund Herding, and Overreaction in Stock Prices. Social Science Research Network Working Paper Series.

Brunnermeier, Nagel, S., 2004. Hedge Funds and the Technology Bubble. The Journal of Finance 59 (5), 2013–2040.

Brunnermeier, M. K., 2008. Bubbles: Entry in New Palgrave Dictionary of Economics. In: Durlauf, S. N., Blume, L. E. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke.

Brush, S. G., Oct. 1967. History of the Lenz-Ising Model. Reviews of Modern Physics 39 (4), 883–893.

Cai, F., Zheng, L., Sep. 2004. Institutional trading and stock returns. Finance Research Letters 1 (3), 178–189.

Chen, J., Hong, H., Stein, J. C., Nov. 2002. Breadth of ownership and stock returns. Journal of Financial Economics 66 (2-3), 171–205.

Chen, Y.-F., Wang, C.-Y., Lin, F.-L., Jul. 2008. Do Qualified Foreign Institutional Investors Herd in Taiwan's Securities Market? Emerging Markets Finance and Trade 44 (4), 62–74.

Chevalier, J., Ellison, G., 1999. Career Concerns of Mutual Fund Managers. The Quarterly Journal of Economics 114 (2), 389–432.

Clement, M. B., Tse, S. Y., 2005. Financial Analyst Characteristics and Herding Behavior in Forecasting. The Journal of Finance 60 (1), 307–341.

Cochrane, J. H., Jun. 2002. Stocks as Money: Convenience Yield and the Tech-Stock Bubble. National Bureau of Economic Research Working Paper Series, 8987+.

Cooper, M. J., Gutierrez, R. C., Hameed, A., 2004. Market States and Momentum. The Journal of Finance 59 (3).

Cutler, D. M., Poterba, J. M., Summers, L. H., 1989. What Moves Stock Prices? Journal of Portfolio Management 15 (2).

Cutler, D. M., Poterba, J. M., Summers, L. H., 1990. Speculative Dynamics and the Role of Feedback Traders. The American Economic Review 80 (2), 63–68.

Daniel, K., Hirshleifer, D., Subrahmanyam, A., 1998. Investor Psychology and Security Market Under- and Overreactions. The Journal of Finance 53 (6), 1839–1885.

De Bondt, W. F. M., Forbes, W. P., 1999. Herding in analyst earnings forecasts: evidence from the United Kingdom. European Financial Management 5 (2), 143–163.

de Long, J. B., Shleifer, A., Summers, L. H., Waldmann, R. J., Jun. 1990. Positive Feedback Investment Strategies and Destabilizing Rational Speculation. The Journal of Finance 45 (2), 379–395.

Debreu, G., 1960. Review of R. D. Luce, Individual choice behavior: A theoretical analysis. American Economic Review 50, 186–188.

Demirer, R., Kutan, A. M., Chen, C.-D., Nov. 2010. Do investors herd in emerging stock markets?: Evidence from the Taiwanese market. Journal of Economic Behavior & Organization 76 (2), 283–295.

Dennis, P. J., Strickland, D., 2002. Who Blinks in Volatile Markets, Individuals or Institutions? The Journal of Finance 57 (5), 1923–1949.

Diba, B. T., Grossman, H. I., 1988a. Explosive Rational Bubbles in Stock Prices? The American Economic Review 78 (3), 520–530.

Diba, B. T., Grossman, H. I., 1988b. The Theory of Rational Bubbles in Stock Prices. The Economic Journal 98 (392), 746–754.

Duffie, D., Gârleanu, N., Pedersen, L. H., Nov. 2002. Securities lending, shorting, and pricing. Journal of Financial Economics 66 (2-3), 307–339.

Dugatkin, L. A., 1992. Sexual selection and imitation: Females copy the mate choice of others. American Naturalist 139, 1384–1389.

Dugatkin, L. A., Guy, J., 1992. Reversal of Female Mate Choice by Copying in the Guppy (Poecilia reticulata). Proceedings: Biological Sciences 249 (1325), 179–184.

Dunbar, R. I. M., 1998. The social brain hypothesis. Evol. Anthropol. 6 (5), 178–190.

Durlauf, S. N., Sep. 1999. How can statistical mechanics contribute to social science? Proceedings of the National Academy of Sciences 96 (19), 10582–10584.

Edelson, M., Sharot, T., Dolan, R. J., Dudai, Y., Jul. 2011. Following the Crowd: Brain Substrates of Long-Term Memory Conformity. Science 333 (6038), 108–111.

Evans, G. W., 1991. Pitfalls in Testing for Explosive Bubbles in Asset Prices. The American Economic Review 81 (4), 922–930.

Fama, E., Nov. 2007. Interview with Eugene Fama. The Federal Reserve Bank of Minneapolis Newsletter - The Region.

Fama, E. F., 1965a. Random Walks in Stock Market Prices. Financial Analysts Journal 21 (5), 55–59.

Fama, E. F., 1965b. The Behavior of Stock-Market Prices. The Journal of Business 38 (1), 34–105.

Flood, R. P., Hodrick, R. J., 1990. On Testing for Speculative Bubbles. The Journal of Economic Perspectives 4 (2), 85–101.

Föllmer, H., Mar. 1974. Random economies with many interacting agents. Journal of Mathematical Economics 1 (1), 51–62.

Foucault, T., Sraer, D., Thesmar, D. J., 2011. Individual Investors and Volatility. The Journal of Finance 66 (4), 1369–1406.

French, K. R., Roll, R., Sep. 1986. Stock return variances: The arrival of information and the reaction of traders. Journal of Financial Economics 17 (1), 5–26.

Froot, K. A., Scharfstein, D. S., Stein, J. C., 1992. Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation. The Journal of Finance 47 (4), 1461–1484.

Galam, S., Moscovici, S., 1991. Towards a theory of collective phenomena: Consensus and attitude changes in groups. Eur. J. Soc. Psychol. 21 (1), 49–74.

Galbraith, J. K., Sep. 1954. The Great Crash 1929, 1st Edition. Boston: Houghton Mifflin.

Gibson, R. M., Höglund, J., Jul. 1992. Copying and sexual selection. Trends in Ecology & Evolution 7 (7), 229–232.

Giraldeau, L. A., Feb. 1997. The ecology of information use, 4th Edition. Wiley-Blackwell.

Graham, J. R., 1999. Herding among Investment Newsletters: Theory and Evidence. The Journal of Finance 54 (1), 237–268.

Greenwood, R., Nagel, S., Aug. 2009. Inexperienced investors and bubbles. Journal of Financial Economics 93 (2), 239–258.

Greiner, W., Neise, L., Stöcker, H., May 1995. Thermodynamics and Statistical Mechanics. Springer.

Griffin, J. M., Harris, J. H., Shu, T., Topaloglu, S., 2011. Who Drove and Burst the Tech Bubble? The Journal of Finance 66 (4), 1251–1290.

Griffin, J. M., Harris, J. H., Topaloglu, S., 2003. The Dynamics of Institutional and Individual Trading. The Journal of Finance 58 (6), 2285–2320.

Grinblatt, M., Titman, S., Wermers, R., 1995. Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior. The American Economic Review, 1088–1105.

Guedj, O., Bouchaud, J.-P., Oct. 2004. Experts' earning forecasts: bias, herding and gossamer information.

Guedj, O., Bouchaud, J.-P., 2005. Experts' earning forecasts: bias, herding and gossamer information. International Journal of Theoretical and Applied Finance (IJTAF) 08 (07), 933–946.

Gumbel, A., Apr. 2005. Herding in delegated portfolio management: When is comparative performance information desirable? European Economic Review 49 (3), 599–626.

Harras, G., Sornette, D., Sep. 2011. How to grow a bubble: A model of myopic adapting agents. Journal of Economic Behavior & Organization 80 (1), 137–152.

Hausmann, J., Ruján, P., Nov. 1997. Stationary Properties of a Randomly Driven Ising Ferromagnet. Physical Review Letters 79 (18), 3339–3342.

Hirshleifer, D., Subrahmanyam, A., Titman, S., 1994. Security Analysis and Trading Patterns when Some Investors Receive Information Before Others. The Journal of Finance 49 (5), 1665–1698.

Holden, C. W., Subrahmanyam, A., Apr. 1996. Risk aversion, liquidity, and endogenous short horizons. Review of Financial Studies 9 (2), 691–722.

Hommes, C., 2008. interacting agents in finance. In: Durlauf, S. N., Blume, L. E. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke.

Hommes, C. H., 2006. Chapter 23 Heterogeneous Agent Models in Economics and Finance. Vol. 2. Elsevier, pp. 1109–1186.

Hong, H., Kubik, J. D., Solomon, A., FebJan 2000. Security Analysts' Career Concerns and Herding of Earnings Forecasts. The RAND Journal of Economics 31 (1), 121+.

Hong, H., Kubik, J. D., Stein, J. C., 2005. Thy Neighbor's Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers. The Journal of Finance 60 (6), 2801–2824.

Hong, H., Scheinkman, J., Xiong, W., Jun. 2006. Asset Float and Speculative Bubbles. The Journal of Finance 61 (3), 1073–1117.

Hong, H., Stein, J. C., 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. The Journal of Finance 54 (6), 2143–2184.

Jarrow, R., 1980. Heterogeneous Expectations, Restrictions on Short Sales, and Equilibrium Asset Prices. The Journal of Finance 35 (5), 1105–1113.

Jegadeesh, N., 1990. Evidence of Predictable Behavior of Security Returns. The Journal of Finance 45 (3), 881–898.

Jegadeesh, N., Titman, S., Mar. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. The Journal of Finance 48 (1), 65–91.

Joulin, A., Lefevre, A., Grunberg, D., Bouchaud, J.-P., Mar. 2008. Stock price jumps: news and volume play a minor role. Wilmott Magazine Sep/Oct.

Kadanoff, L. P., Jun. 2000. Statistical Physics: Statics, Dynamics and Remormalization. World Scientific Pub Co Inc.

Kaizoji, T., Dec. 2000. Speculative bubbles and crashes in stock markets: an interacting-agent model of speculative activity. Physica A: Statistical Mechanics and its Applications 287 (3-4), 493–506.

Kaizoji, T., Bornholdt, S., Fujiwara, Y., Dec. 2002. Dynamics of price and trading volume in a spin model of stock markets with heterogeneous agents. Physica A: Statistical Mechanics and its Applications 316 (1-4), 441–452.

Kaizoji, T., Saichev, A., Sornette, D., Sep. 2011. Equilibrium model of rational and noise traders: bifurcations to endogenous bubbles. arXiv:1109.4726v1.

Kaizoji, T., Sornette, D., 2010. Bubbles and Crashes. In: Encyclopedia of Quantitative Finance. John Wiley & Sons, Ltd (extended version at arXiv:0812.2449v1).

Keynes, J. M., 1936. The General Theory of Employment, Interest and Money. Harcourt Brace and Co.

Kim, C., Pantzalis, C., 2003. Global/Industrial Diversification and Analyst Herding. Financial Analysts Journal 59 (2), 69–79.

Kindleberger, C. P., Oct. 1978. Manias, Panics, and Crashes: A History of Financial Crises, 1st Edition. Basic Books.

Kindleberger, C. P., Aliber, R., Oct. 2005. Manias, Panics, and Crashes: A History of Financial Crises (Wiley Investment Classics), 5th Edition. Wiley.

Kirman, A., Feb. 1993. Ants, Rationality, and Recruitment. The Quarterly Journal of Economics 108 (1), 137–156.

Klemkosky, R. C., Mar. 1977. The Impact and Efficiency of Institutional Net Trading Imbalances. The Journal of Finance 32 (1), 79+.

Koski, J. L., Pontiff, J., 1999. How Are Derivatives Used? Evidence from the Mutual Fund Industry. The Journal of Finance 54 (2), 791–816.

Kyle, A. S., 1985. Continuous Auctions and Insider Trading. Econometrica 53 (6), 1315–1335.

La Porta, R., Lakonishok, J., Shleifer, A., Vishny, R., 1997. Good News for Value Stocks: Further Evidence on Market Efficiency. The Journal of Finance 52 (2), 859–874.

Lakonishok, J., Shleifer, A., Vishny, R. W., Aug. 1992. The impact of institutional trading on stock prices. Journal of Financial Economics 32 (1), 23–43.

Lakonishok, J., Shleifer, A., Vishny, R. W., 1994. Contrarian Investment, Extrapolation, and Risk. The Journal of Finance 49 (5), 1541–1578.

Langer, E. J., Roth, J., 1975. Heads I win, tails it's chance: The illusion of control as a function of the sequence of outcomes in a purely chance task. Journal of Personality and Social Psychology 34, 191–198.

Lei, V., Noussair, C. N., Plott, C. R., 2001. Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality. Econometrica 69 (4), 831–859.

Leland, H. E., Rubinstein, M., 1988. Comments on the Market Crash: Six Months After. Journal of Economic Perspectives 2 (3), 45–50.

Lintner, J., 1969. The Aggregation of Investor's Diverse Judgments and Preferences in Purely Competitive Security Markets. The Journal of Financial and Quantitative Analysis 4 (4), 347–400.

Lorenz, J., Rauhut, H., Schweitzer, F., Helbing, D., May 2011. How social influence can undermine the wisdom of crowd effect. Proceedings of the National Academy of Sciences 108 (22), 9020–9025.

Luce, R. D., 1959. Individual Choice Behavior: A theoretical analysis. Wiley.

Luce, R. D., Suppes, P., 1965. Preference, Utility, and Subjective Probability. Handbook of mathematical psychology 3 (171), 249–410.

Lux, T., Jul. 1995. Herd Behaviour, Bubbles and Crashes. The Economic Journal 105 (431), 881+.

Lux, T., Sornette, D., 2002. On Rational Bubbles and Fat Tails. Journal of Money, Credit and Banking 34 (3), 589–610.

Marschak, J., 1959. Binary Choice Constraints on Random Utility Indicators (74).

Maug, E., Naik, N., 2011. Herding and Delegated Portfolio Management: The Impact of Relative Performance Evaluation on Asset Allocation. The Quarterly Journal of Finance 1, 265–292.

McFadden, D., 1974. Conditional Logit Analysis of Qualitative Choice Behavior. In: Zarembka, P. (Ed.), Frontiers in econometrics. Academic Press, New York, pp. 105–142.

McKelvey, R. D., Palfrey, T. R., 1992. An Experimental Study of the Centipede Game. Econometrica 60 (4), 803–836.

Michard, Q., Bouchaud, J. P., Sep. 2005. Theory of collective opinion shifts: from smooth trends to abrupt swings. The European Physical Journal B 47 (1), 151–159.

Miller, E. M., 1977. Risk, Uncertainty, and Divergence of Opinion. The Journal of Finance 32 (4), 1151–1168.

Nofsinger, J. R., Sias, R. W., 1999. Herding and Feedback Trading by Institutional and Individual Investors. The Journal of Finance 54 (6), 2263–2295.

Ofek, E., Richardson, M., 2003. DotCom Mania: The Rise and Fall of Internet Stock Prices. The Journal of Finance 58 (3), 1113–1138.

Offerman, T., Sonnemans, J., 2004. What's Causing Overreaction? An Experimental Investigation of Recency and the Hot-hand Effect. Scandinavian Journal of Economics 106 (3), 533–554.

Olsen, R. A., 1996. Implications of Herding Behavior for Earnings Estimation, Risk Assessment, and Stock Returns. Financial Analysts Journal 52 (4), 37–41.

Onsager, L., Feb. 1944. Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition. Physical Review Online Archive (Prola) 65, 117–149.

Orléan, A., 1995. Bayesian interactions and collective dynamics of opinion - Herd behavior and mimetic contagion. Journal of Economic Behavior & Organization 28, 257–274.

Parr, C. S., 1985. Changing Demographics in the Community Colleges. Hispania 68 (1).

Pennisi, E., Apr. 2010. Conquering by Copying. Science 328 (5975), 165–167.

Prendergast, C., 1993. A Theory of "Yes Men". The American Economic Review 83 (4), 757–770.

Reif, F., Jun. 1965. Fundamentals of Statistical and Thermal Physics, 1st Edition. McGraw-Hill.

Rosenthal, R., Aug. 1981. Games of perfect information, predatory pricing and the chain-store paradox. Journal of Economic Theory 25 (1), 92–100.

Rouwenhorst, K. G., 1998. International Momentum Strategies. The Journal of Finance 53 (1), 267–284.

Rouwenhorst, K. G., 1999. Local Return Factors and Turnover in Emerging Stock Markets. The Journal of Finance 54 (4), 1439–1464.

Samuelson, P. A., 1965. Proof that Properly Anticipated Prices Fluctuate Randomly. Industrial Management Review 6, 41–49.

Scharfstein, D. S., Stein, J. C., 1990. Herd Behavior and Investment. The American Economic Review 80 (3), 465–479.

Schelling, T. C., Jul. 1971. Dynamic models of segregation†. The Journal of Mathematical Sociology 1 (2), 143–186.

Schelling, T. C., 1973. Hockey Helmets, Concealed Weapons, and Daylight Saving: A Study of Binary Choices with Externalities. The Journal of Conflict Resolution 17 (3), 381–428.

Sherif, M., Jul. 1937. An Experimental Approach to the Study of Attitudes. Sociometry 1 (1/2), 90+.

Shiller, R. J., 1981. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? The American Economic Review 71 (3), 421–436.

Shiller, R. J., Nov. 1987. Investor Behavior in the October 1987 Stock Market Crash: Survey Evidence. Tech. Rep. 2446, National Bureau of Economic Research, Inc.

Shiller, R. J., Fischer, S., Friedman, B. M., 1984. Stock Prices and Social Dynamics. Brookings Papers on Economic Activity 1984 (2), 457–510.

Shiller, R. J., Pound, J., Aug. 1989. Survey evidence on diffusion of interest and information among investors. Journal of Economic Behavior & Organization 12 (1), 47–66.

Shleifer, A., Vishny, R. W., Mar. 1997. The Limits of Arbitrage. The Journal of Finance 52 (1), 35–55.

Sias, R. W., Jan. 2004. Institutional Herding. Review of Financial Studies 17 (1), 165–206.

Sias, R. W., Starks, L. T., Oct. 1997. Return autocorrelation and institutional investors. Journal of Financial Economics 46 (1), 103–131.

Sias, R. W., Starks, L. T., Titman, S., Sep. 2001. The Price Impact Of Institutional Trading. Social Science Research Network Working Paper Series.

Smith, V. L., Suchanek, G. L., Williams, A. W., 1988. Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. Econometrica 56 (5), 1119–1151.

Sornette, D., 2003. Why Stock Markets Crash. Princeton University Press. Princeton, NJ.

Sornette, D., Pisarenko, V. F., Apr. 2008. Properties of a simple bilinear stochastic model: Estimation and predictability. Physica D: Nonlinear Phenomena 237 (4), 429–445.

Thurstone, L. L., 1927. A law of comparative judgement. Psychological Review 34, 273–286.

Tirole, J., Sep. 1982. On the Possibility of Speculation under Rational Expectations. Econometrica 50 (5), 1163–1181.

Train, K. E., Jan. 2003. Discrete Choice Methods with Simulation. Cambridge University Press.

Trichet, J.-C., 2001. Preserving Financial Stability in an Increasing Globalized World. BIS Review 54, 1–4.

Trueman, B., Jan. 1994. Analyst forecasts and herding behavior. Review of Financial Studies 7 (1), 97–124.

Vandewalle, N., Boveroux, P., Brisbois, F., May 2000. Domino effect for world market fluctuations. The European Physical Journal B - Condensed Matter and Complex Systems 15 (3), 547–549.

Warren, W. G., 1971. Correlation or Regression: Bias or Precision. Journal of the Royal Statistical Society. Series C (Applied Statistics) 20 (2).

Welch, I., Jun. 1992. Sequential Sales, Learning, and Cascades. The Journal of Finance 47 (2), 695–732.

Welch, I., Dec. 2000. Herding among security analysts. Journal of Financial Economics 58 (3), 369–396.

Wermers, R., 1999. Mutual Fund Herding and the Impact on Stock Prices. The Journal of Finance 54 (2), 581–622.

Wermers, R. R., Dec. 1998. Herding, Trade Reversals, and Cascading by Institutional Investors. Social Science Research Network Working Paper Series.

West, K. D., 1987. A Specification Test for Speculative Bubbles. The Quarterly Journal of Economics 102 (3), 553–580.

West, K. D., 1988. Bubbles, Fads and Stock Price Volatility Tests: A Partial Evaluation. The Journal of Finance 43 (3), 639–656.

Zaki, J., Schirmer, J., Mitchell, J. P., Jun. 2011. Social Influence Modulates the Neural Computation of Value. Psychological Science 22 (7), 894–900.

Zhang, Y.-C., Jul. 1999. Toward a theory of marginally efficient markets. Physica A: Statistical Mechanics and its Applications 269 (1), 30–44.

Zwiebel, J., Jan. 1995. Corporate Conservatism and Relative Compensation. Journal of Political Economy 103 (1), 1.

Curriculum Vitae

1980	Born in Luxembourg, Luxembourg
1993 – 2000	Diplôme de fin d'études secondaires classiques (Secondary School leaving examination), Lycée Robert Schuman, Luxembourg. Major: Mathematics.
2000 – 2004	Bachelor of Science in Physics, with a Minor in Mathematics, University of Fribourg, Switzerland
2004 – 2006	Master of Science in Physics, University of Fribourg, Switzerland.
2006 – 2012	PhD-Student at the chair of "Entrepreneurial Risks", Department for Management Technology and Economics, ETH Zürich.