Charles University in Prague

Faculty of Social Sciences
Institute of Economic Studies

BACHELOR THESIS

Ising Model in Finance

From Microscopic Rules to Macroscopic Phenomena

Author: **Pavel Dvořák**

Supervisor: **PhDr. Ladislav Krištoufek**

Academic Year: **2011/2012**

Declaration of Authorship

The author hereby declares that he compiled this thesis independently, using only the listed resources and literature. The author also declares that he has not used this thesis to acquire another academic degree.

The author grants to Charles University permission to reproduce and to distribute copies of this thesis document in whole or in part.

Prague, May 9, 2012

Signature

Acknowledgments

I am grateful especially to PhDr. Ladislav Krištoufek for his priceless comments and compelling ideas. I would like to express gratitude to PhDr. Tomáš Havránek for a great LaTeX thesis template and to RNDr. Antonín Slavík, Ph.D. for advice on acceleration of computations in Wolfram Mathematica. I also appreciate keen suggestions from RNDr. Miroslav Kotrla, CSc. and RNDr. František Konopecký on the Ising model and programming methods.

Bibliographic entry

Dvořák, P. (2012): "From Microscopic Rules to Macroscopic Phenomena: Ising Model in Finance." *(Unpublished bachelor thesis).* Charles University in Prague. Supervisor: PhDr. Ladislav Krištoufek.

Length: 59 036 characters

Abstract

The main objective of this thesis is to inspect the abilities of the Ising model to exhibit selected statistical properties, or stylized facts, that are common to a wide range of financial assets. The investigated properties are heteroskedasticity of returns, rapidly decaying linear autocorrelation, volatility clustering, heavy tails, negative skewness and non-Gaussianity of the return distribution. In the first part of the thesis, we test the presence of these stylized facts in S&P 500 daily returns over the last 30 years. The main part of the thesis is dedicated to the Ising model-based simulations and to discussion of the results. New features such as Poisson process governed lag or magnetisation dependent trading activity are incorporated in the model. We conclude that the Ising model is able to convincingly replicate most of the examined statistical properties while even more satisfactory results can be obtained with appropriate tuning.

JEL Classification	C15, D84, G17, G63,
Keywords	Ising model, ferromagnetism, econophysics, financial markets, stylized facts, volatility clustering, long-term dependency

Author's e-mail	dvorak.pavel@centrum.cz
Supervisor's e-mail	kristoufek@ies-prague.org

Abstrakt

Hlavním cílem této práce je zjistit, zdali je Isingův model schopen reprodukovat vybrané statistické vlastnosti (někdy též stylizovaná fakta), které jsou typické pro širokou škálu finančních aktiv. Zkoumanými vlastnostmi jsou heteroskedasticita výnosů, rapidně klesající autokorelace, shluky volatilit, těžké

chvosty, záporná šikmost a nenormalita rozdělení výnosů. V první části práce testujeme přítomnost těchto stylizovaných faktů na denních výnosech indexu S&P 500 za posledních 30 let. Hlavní část práce je věnována Isingovým simulacím a shrnutí souvisejících výsledků. Do modelu jsou také včleněny nové prvky jako časová prodleva určená Poissonovým procesem či aktivita obchodníků ovlivněná velikostí celkové magnetizace. Docházíme k závěru, že Isingův model je schopen spolehlivě replikovat většinu zkoumaných statistických vlastností, přičemž dalšího zlepšení lze dosáhnout vhodnou modifikací modelu.

Klasifikace JEL	C15, D84, G17, G63,
Klíčová slova	Isingův model, ferromagnetismus, ekonofyzika, finanční trhy, stylizovaná fakta, shluky volatilit, dlouhá paměť ve volatilitě

E-mail autora	`dvorak.pavel@centrum.cz`
E-mail vedoucího práce	`kristoufek@ies-prague.org`

Contents

List of Tables

List of Figures

Acronyms

ACF	Autocorrelation Function
ARCH	Autoregressive Conditional Heteroskedasticity
CDF	Cumulative Distribution Function
CI	Confidence Interval
EMH	Efficient Market Hypothesis
FOREX	Foreign Exchange
GARCH	Generalized Autoregressive Conditional Heteroskedasticity
HAM	Heterogeneous Agents Model
iid	independent identically distributed
JB	Jarque-Bera
NASDAQ	National Association of Securities Dealers Automated Quotations
NYSE	New York Stock Exchange
PDF	Probability Density Function
S&P 500	Standard and Poor's 500

Chapter 1

Introduction

Financial markets have undergone great transformation since the 1970s. With the advent of electronic trading system, the volume of transactions escalated and the geographical location of traders was no longer important. In 1973, foreign currencies began to be traded in Foreign Exchange (FOREX) market worldwide. As Mantegna & Stanley (1999) point out, the volume of FOREX transactions in 1995 was 80 times higher than in 1973. Data from financial time series are highly reliable and readily accessible to anyone which makes them attractive to researchers not only from the field of economics but also, as we shall see, from the area of physics.

Until the end of 1990s, economic systems were investigated mainly by economists and mathematicians. Since then, however, the number of physicists trying to explain economic phenomena with tools characteristic to statistical mechanics has grown rapidly. One reason is that new discoveries in statistical mechanics, phase transition and non-linear dynamics over the last decades may prove suitable for explaining and interpreting various economic processes. The book by Mantegna & Stanley (1999), labelled by its publisher as the first monograph on econophysics, provides basic introduction to financial analysis targeted at researchers with background in physics. With reference to economic processes, it accentuates the use of central limit theorem and law of large numbers, discusses various distribution families of raw and logarithmic stock returns, random walks, testing validity of Efficient Market Hypothesis (EMH), scaling, option pricing, etc.

With the onset of computer-based analysis of the high-frequency financial data, certain patterns or common statistical properties were exhibited by a lot of mutually independent financial instruments and markets. Partial discussion

of these can be found, for example, in works of Bollerslev *et al.* (1992), Brock & De Lima (1995), Campbell *et al.* (1998) and Gopikrishnan *et al.* (2000). Succinct overview of these properties, or stylized facts, with some empirical examples is presented in a widely cited paper by Cont (2001). We chose this work as a reference base for Chapter 2.

Even though the empirical stylized facts are commonplace in much of the financial data, they are generally difficult to imitate with a specifically constructed model. Nevertheless, there has been a variety of more or less successful attempts that are worth mentioning.

The Heterogeneous Agents Model (HAM), summarized for example by Brock & Hommes (1998), Hommes (2006), or Baruník *et al.* (2009) and used frequently in the works trying to model financial processes, is a behavioural agent-based model with emphasis on different expectations of each agent (we speak of *heterogeneity* of expectations, investment strategies and so on). As opposed to traditional rational expectation theory, the HAM assumes imperfect knowledge of market fundamentals, bounded rationality and high cost of information leading to erroneous and biased decisions of the agents. As Simon (1978) suggests, this type of behaviour is more realistic than assuming perfect rationality of agents.[1] The argument is that agents do not have the ability to acquire and evaluate all the information that is necessary to make a correct investment decision into the future. The agents rather adopt simple, straightforward strategies based on their not always rigorously stipulated decision rules. Last but not least, the development in computer science and mathematics has allowed for numerical simulations of the HAM that was not possible before the 1970s.

Several HAM originated models have been introduced over the last decades. Among these are agent-based models as pioneered in works of Caldarelli *et al.* (1997) and Lux & Marchesi (1998) or minority games models elaborated by Challet & Zhang (1997), Challet *et al.* (1999) and Challet *et al.* (2001). The pitfall of these otherwise well designed models is relatively high difficulty and excessive parametrisation that hinders simple analytical tractability. Another branch of literature concerns spin models that also comprise the Ising model. General spin model was described by Chowdhury & Stauffer (1998), the Ising model-based simulations with many extensions were carried out mainly by Iori (1999), Bornholdt (2001), Kaizoji *et al.* (2002), Sornette & Zhou (2005), Sieczka & Hołyst (2007) and many more.

[1] Nevertheless, rational expectations and EMH had been the basic paradigms in economic theory until relatively recent times.

There are various advantages of using the Ising model. Most prominently, all models referenced in the literature describe two counteracting forces exerted upon agents – the herding behaviour driving them to act in conformity and the tendency to stay in minority where profits are realizable. While the former force evokes setup of agent-based models, the latter is essential in all minority game simulations. The Ising model as presented in this thesis is thus a synthesis of both above mentioned models while preserving the simplicity needed for an easy-to-understand evaluation. The major contribution of this thesis lies in collection, sorting and mutual comparison of so far described variants of the Ising model, in introducing many new extensions and upgrades of the model as well as in suggestions for additional modifications.

The structure of the thesis is as follows. In Chapter 2, we investigate presence of some of the stylized facts in real data. We state reasons for choosing Standard and Poor's 500 (S&P 500) as the reference index and list the properties we are about to contemplate. Chapter 3 is dedicated solely to the Ising simulations. We describe the features of the model, define parameters and compare the results of the simulations for different types of models. The ability to mimic selected statistical properties of the real financial data is also discussed within each section.

Majority of the regressions, statistical tables and analyses are performed using Stata software while all the scripts needed for simulations of the Ising model are written in Wolfram Mathematica. Figures and charts are also produced in Wolfram Mathematica software. Finally, some of the post-estimation analyses such as Jarque-Bera (JB) or Ljung-Box tests are carried out in Gretl. All scripts, do files and Wolfram Mathematica notebooks, necessary for understanding the principles of the simulations, are available upon request.

Chapter 2

Statistical properties of financial markets

There is a wide range of statistical properties, often called *stylized facts* as proposed by Cont (2001), that are common to most financial instruments and markets. The aim of this chapter is to inspect and describe the price variation on the S&P 500 index along with testing the presence of the stylized facts on this stock market. The properties under scrutiny are absence of autocorrelation in returns, slow decline in autocorrelation of absolute returns with increasing lag, heteroskedasticity and heavy tails of returns distribution, aggregate normality of such distribution, volatility clustering, etc. We begin with statistical description of the S&P 500 index in Section 2.1 and then proceed to inspection of each stylized fact in corresponding subsections.

2.1 Standard & Poor's 500 index

The S&P 500 index consists of 500 leading companies from various industry sectors of the US economy and is traded on the two biggest stock exchanges in the world (sorted by market capitalisation), i.e. on NYSE and on NASDAQ. It is a market value weighted index which means that each stock is weighted adequately to the market value of the firm. The S&P 500 is seen as the most representative index of the US large cap companies.

We gathered daily data for the S&P 500 index over the last thirty years. The data range from January 2, 1981 to November 11, 2011. The overall development of the price P and daily returns (as defined in Equation 2.1 for

$\Delta t = 1$) are depicted in Figure 2.1.[1]

$$r_{\Delta t}(t) = \log(P_t) - \log(P_{t-\Delta t}) \tag{2.1}$$

The S&P 500 index clearly reflected all major events and crises. During Black Monday on October 17 in 1987, the price fell by more than 22%, the largest one-day drop in the history of the index. The S&P 500 index enjoyed a period of accelerating growth during the Dotcom bubble in late 1990s with subsequent long lasting descend. The financial crisis of 2007-2009 saw another significant fall in prices but since then the index has grown incessantly. The daily returns clearly show signs of clustered volatility which will be described formally in later subsections. Main characteristics of the S&P 500 index, daily log-returns and daily absolute log-returns are summarized in Table 2.1. The table suggests strongly leptokurtic distribution of the daily log-returns and negative skewness. Note that the lowest value of daily returns is slightly less than 20 standard deviations away from the mean value. If the distribution of returns were normal, the probability of occurrence of such extreme values would practically equal zero. We will now test the presence of stylized facts in these data.

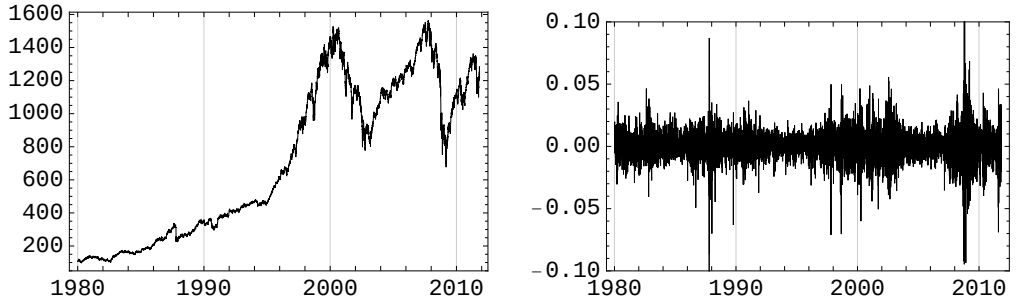

(a) S&P 500 Index 2001-2011 (b) S&P 500 daily returns 2001-2011

Figure 2.1: S&P 500 index and daily returns

[1]There are several reasons for using logarithmic returns instead of raw returns. These include time additivity and easier mathematical manipulation while showing only negligible deviation from raw returns for small Δt. For more thorough discussion of compound returns see Meucci (2010).

Variable	Mean	Std. Dev.	Min.	Max.	Skewn.	Kurt.	N
index	708.94	464.05	98.22	1565.15	0.196	1.48	8034
r_1	0	0.012	-0.229	0.11	-1.19	29.69	8033
$\|r_1\|$	0.0077	0.0086	0	0.229	4.94	73.37	8033

Table 2.1: Summary statistics for S&P 500

2.1.1 Linear autocorrelation of daily returns

First, we define the sample Autocorrelation Function (ACF) as:

$$\rho(\tau) = \frac{\sum_{t=\tau+1}^{T}(r_t - \bar{r})(r_{t-\tau} - \bar{r})}{\sum_{t=1}^{T}(r_t - \bar{r})^2} \tag{2.2}$$

where r is the return, τ is the time lag and $\bar{r} = \frac{1}{T}\sum_{t=1}^{T} r_t$ is the time average. We computed the values of the ACF for $\tau \in \{1, \ldots, 150\}$ and summarized them graphically in Figure 2.2. We also included the 95% confidence intervals around zero. Here the null hypothesis is that there is no autocorrelation for given lag on 5% significance level. If the null holds, the distribution of correlation estimates asymptotically approaches normal distribution with increasing lag.[2] The critical value of correlation for the 95% significance equals to $0 \pm 1.96/\sqrt{T}$ where T is the sample size (in our case $T = 8033$). In Figure 2.2, the confidence intervals form a dashed belt approx. 0.0435 wide. We observe that sample correlations at several lags are significantly different from zero which seems to contradict the stylized fact of no autocorrelation of returns.

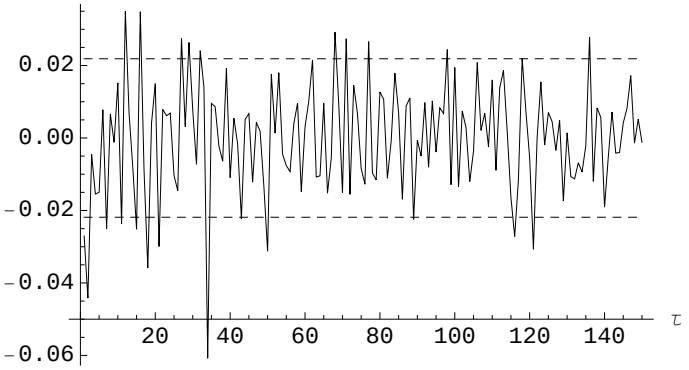

Figure 2.2: Linear autocorrelation of S&P 500 daily returns with 95% CIs

[2]To ensure such asymptotic convergence, the assumption of covariance stationarity of the time series must also hold. As we shall see later, this assumption is fulfilled only asymptotically. Treatment of non-stationarity and further issues are partly discussed in Subsection 2.1.6.

We also tested the presence of autocorrelation via the Ljung-Box test. Under the null the returns are independently distributed and no non-zero autocorrelation may occur. The p-value of the Q-statistic for any given lag practically equals to zero and the null hypothesis is rejected. See Ljung & Box (1978) for definition and application of the test.

In our data, the autocorrelation of returns proved statistically significant. Nevertheless, it is questionable whether ACF values ranging from approximately -0.04 to 0.03 carry any economic importance. Furthermore, the calculation of confidence intervals relies on the rarely met assumption of normality (discussed in Subsection 2.1.6).

Even though absence of autocorrelation is listed as one of the stylized facts, we cannot say it generally holds for any Δt in Equation 2.1. For example, Gopikrishnan & Plerou (1999) and Cont (2001) suggest that autocorrelation disappears for Δt in order of seconds or minutes.[3] For $\Delta t > 1$ day, the speed of autocorrelation decay varies from sample to sample.

2.1.2 Autocorrelation of absolute returns

Various non–linear correlation functions were proposed to capture the difference between the asset returns and white noise (Cont, 2001, p. 229). We will work with sample ACF of absolute values of returns, defined as:

$$\rho_A(\tau) = \frac{\sum_{t=\tau+1}^{T}(|r_t| - \overline{|r|})(|r_{t-\tau}| - \overline{|r|})}{\sum_{t=1}^{T}(|r_t| - \overline{|r|})^2} \tag{2.3}$$

Such definition is used as a measure of memory in volatility. To illustrate this, we decompose returns into its magnitude and sign:

$$r_t = \mathrm{sgn}(r_t)|r_t| \tag{2.4}$$

The autocorrelation plotted as a function of τ is shown in Figure 2.3.

We see that the the autocorrelation of absolute returns is fairly high for $\tau \leq 30$ and remains positive for all analysed lags. Slow decay in the functional values can be regarded as a sign of volatility clustering. This means that phases of high volatility today will be most likely pursued by phases of high volatility tomorrow, the day after tomorrow and so on. The CIs are not included

[3] It is necessary to point out that there were also studies that showed non-zero autocorrelation for very short Δt. See, for example, Andersen & Bollerslev (1997).

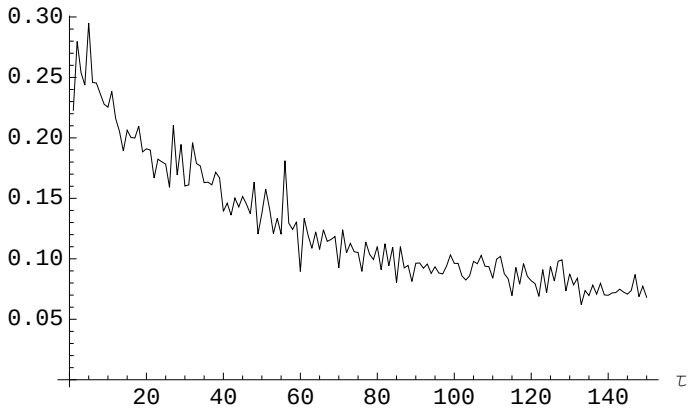

Figure 2.3: Autocorrelation of S&P 500 absolute daily returns

in Figure 2.3 since the distribution of absolute returns does not in any way approach normal distribution.

The autocorrelation functions often exhibit various repetitive patterns with increasing τ. If we plotted the values of our autocorrelation function up to $\tau = 500$, for example, the patterns would also stand out. These regular fluctuations are most likely caused through periodicity which is more pronounced for intra-day market data as seen on various examples in Andersen & Bollerslev (1997).

It was observed (for example Cont *et al.*, 1997) that the decay of autocorrelation of absolute returns is asymptotically proportional to a power-law function

$$\rho_A(\tau) = \frac{A}{\tau^\alpha} \tag{2.5}$$

where A and α are parameters with α empirically identified to lie between 0.2 and 0.4. The α parameter is characteristic for long-range dependent processes. We also compared the decay of autocorrelation to exponential function, defined as

$$\rho_A(\tau) = \frac{B}{\exp(\beta\tau)} \tag{2.6}$$

where B and β are also parameters with β typical for short-range dependent processes. For comparison of the two reference functions, see Kantelhardt (2009). For the S&P 500 returns, α and β were found to equal approximately 0.12 and 0.0096, respectively. The comparison for $A = 0.137$ and $B = 0.24$ is depicted in Figure 2.4. Since power-law is an asymptotic property, we are interested in the behaviour of the ACF for high values of τ. Henceforth the fitted power-law function is always computed with values for $\tau \geq 60$. It seems

the power-law fit is more suitable to approximate the behaviour of the ACF for high τ. However, its dominance is probably not as unambiguous as expected or generally claimed.

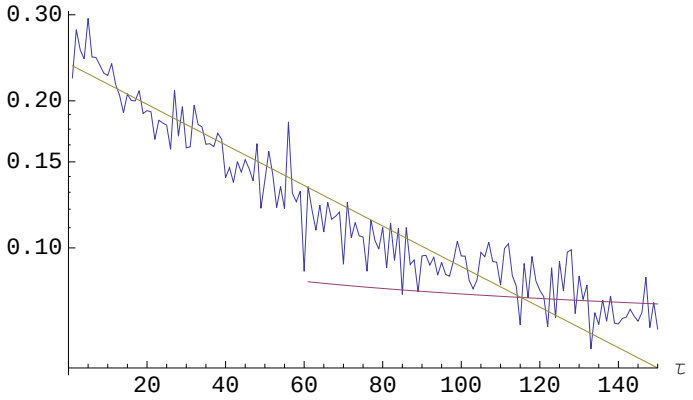

Figure 2.4: Semi-log plot of decay of autocorrelation in absolute returns (blue), including fitted power-law (purple) and exponential (purple) functions as defined in Equation 2.5 and Equation 2.6, respectively.

It is suitable to mention many other functional forms that were inspected throughout the literature. These include logarithms, trigonometric functions or powers of higher order. For further discussion of functional forms within autocorrelation see Comte & Renault (1996).

2.1.3 Heteroskedasticity of returns

Simply by looking back at Subfigure 2.1(b) we see that the variance of returns does alter with time. To detect heteroskedasticity, we choose to perform Autoregressive Conditional Heteroskedasticity (ARCH) test as proposed by Engle (1982). Though the presence of ARCH does not necessarily imply that $\mathrm{Var}(u_{it}) = \sigma_{it}^2$, it does mean that there is a certain dynamic form of heteroskedasticity that would have to be considered for further regressions or data analyses. Presence of ARCH only implies local non-stationarity whereas asymptotic stationarity of the model is not affected.

Assuming that returns have zero mean and no autocorrelation is present, the ARCH(q) model is specified as follows:

$$r_t^2 = \gamma_0 + \sum_{i=1}^{q} \gamma_i r_{t-i}^2 + u_t \tag{2.7}$$

where u_t is an independent identically distributed (iid) sequence with zero mean

and variance equal to σ_u^2. Using Stata software, we performed ARCH(q) tests for $q \in \{1, \ldots, 4\}$. All the lags proved to be significant up to $q = 4$. The results are summarized in Table 2.2. Only the results for ARCH(4) test are presented.

Table 2.2: ARCH(4) test results

	S&P 500 returns	
ARCH		
L.arch	0.0993***	(16.54)
L2.arch	0.188***	(17.77)
L3.arch	0.177***	(15.50)
L4.arch	0.204***	(17.48)
_cons	0.0000441***	(46.34)
N	8033	

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The null hypothesis for the ARCH test is $H_0 : \gamma_i = 0$ for $i \in \{1, \ldots, q\}$. As Table 2.2 suggests, we strongly reject the null in all four cases even on 0.1% significance level. The ARCH(4) is present.

We also estimated the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) test introduced by Bollerslev (1986) and well explained by Engle (2001). We use the GARCH(p,q) model described by Calvet & Fisher (2008):

$$r_t = \epsilon_t \sqrt{h_t} \tag{2.8}$$

$$h_t = \alpha + \sum_{i=1}^{p} \beta_i h_{t-i} + \sum_{j=1}^{q} \gamma_j r_{t-j}^2 \tag{2.9}$$

Here, the log returns r_t are modelled by iid standard normal random variable ϵ_t scaled by the conditional variance of r_t at time $t - 1$ labelled as h_t. This variable follows autoregressive process and is determined by past squared returns r_t^2. In our case, we computed the GARCH(1,1) model.

As seen in Table 2.3, the GARCH(1,1) effect is also very significant. More information about GARCH models can be found in Subsection 2.1.6.

2.1.4 Heavy tails and normality

About a half century ago, Mandelbrot (1963) emphasized the non-Gaussian character of empirical returns. The reason he states is that "*the empirical*

Table 2.3: GARCH(1,1) test results

	S&P 500 returns	
ARCH		
L.arch	0.0780***	(43.75)
L.garch	0.913***	(314.04)
_cons	0.00000133***	(10.93)
N	8033	

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

distributions of price changes are usually too peaked to be relative to samples from Gaussian populations." As we shall see, the same issue arises in our case.

As can be inferred from Table 2.4, the sample kurtosis index of S&P 500 daily log returns equals to 29.69. Given that the kurtosis of the normal distribution is always 3, the distribution of returns has much heavier tails than normal distribution. This means that the extreme values (far from the mean) are more likely to occur than under the normal distribution with corresponding expected value and variance.[4] We found the empirical Probability Density Function (PDF) of standardized returns using Smooth Kernel Estimation and compared it with the normal PDF with zero mean and unit variance. Both PDFs are plotted in Figure 2.5.

We also ran the JB test of normality introduced by Jarque & Bera (1980). The null hypothesis is that both skewness and excess kurtosis (i.e. kurtosis minus three) of the investigated distribution are zero. For our data the p-value of the test is practically zero and the null is rejected at any reasonable significance level.

Even though it is not too hard to approximate the distribution of returns, its precise form has been subject to many attempts throughout the literature. Suggestions include Student, hyperbolic or Wald distributions (Cont, 2001). Blattberg & Gonedes (1974) and Eberlein *et al.* (1998) provide a thorough discussion of further potential distribution families.

[4]The empirical kurtosis may actually reach far higher values. For certain daily stock returns, Pagan (1996) identified kurtosis over 200. At the same time, however, it is necessary to ponder these results with caution as they are very sensitive to data censoring. After dropping the maximum observation, for example, kurtosis may drop by more than half. Generally speaking, the smaller the Δt in Equation 2.1 the more pronounced are the non-Gaussian properties of return distributions.

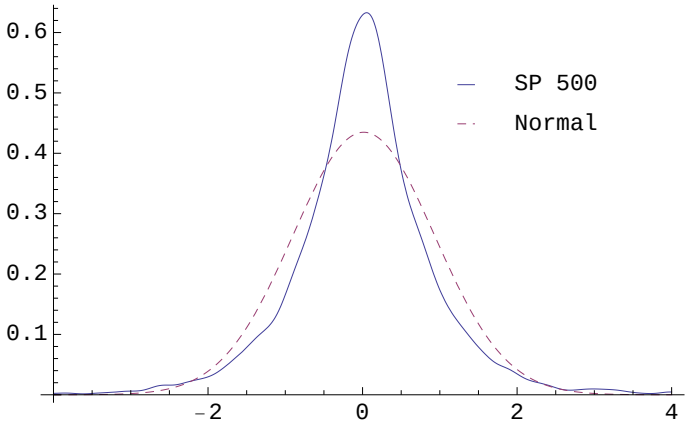

Figure 2.5: Standardized empirical distribution of the S&P 500 daily log returns

2.1.5 Aggregate normality

Even though it is common that the empirical distributions of daily returns do not resemble normal even for large samples, it has been suggested by Cont (2001) that with increasing Δt in Equation 2.1 the distribution of returns does approach normal. We computed the logarithmic returns for various lags and summarized their statistical properties in Table 2.4.

Variable	Mean	Std. Dev.	Min.	Max.	Skew.	Kurt.	N
Lag 1	0	0.012	-0.229	0.11	-1.19	29.69	8033
Lag 4	0.001	0.023	-0.336	0.165	-1.90	32.05	2007
Lag 8	0.002	0.03	-0.237	0.114	-1.13	10.06	1003
Lag 16	0.005	0.043	-0.234	0.156	-1.36	9.66	501
Lag 64	0.02	0.087	-0.367	0.291	-1.30	8.19	124
Lag 128	0.04	0.122	-0.449	0.277	-1.33	6.22	61

Table 2.4: Summary statistics for S&P 500 at different time lags.

While the mean remains very close to zero for all lags, the standard deviation σ increases with lag as expected, roughly following the rule $\sigma(\Delta t) = \sigma(1)\sqrt{\Delta t}$, where $\sigma(1)$ is the standard deviation of returns for lag 1. The skewness hovers around -1.3 which indicates that occurrence of lower negative values is more probable than occurrence of the positive ones. This is also supported by the list of minimum and maximum values. In all cases, $|\text{min}| > |\text{max}|$. The only sustaining argument for aggregate normality is thus decreasing kurtosis. By performing JB test, however, normality is still strongly rejected for any lag even though the test statistic decreases with increasing lag. We conclude more observations would be needed to accurately assess the issue of normality in

our data. The negative skewness and other properties are well illustrated by PDFs of empirical distributions of the log returns in Figure 2.6. The dashed line belongs to standard normal PDF. The distributions were found by Smooth Kernel Estimation.

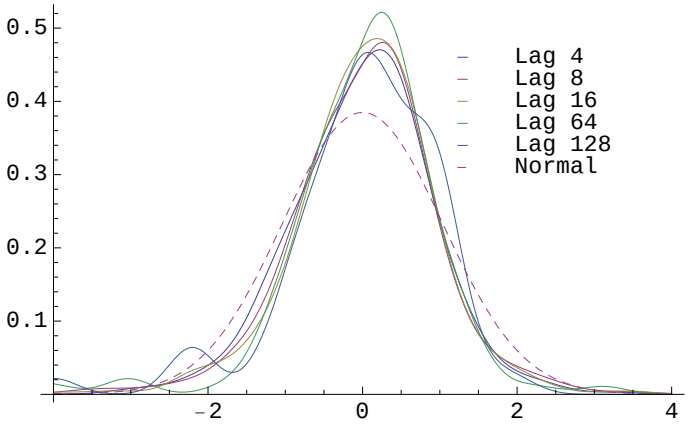

Figure 2.6: Empirical distributions of standardized S&P 500 log returns for various values of Δt. The dashed line is standard normal distribution.

2.1.6 Validity and statistical issues

To compute autocorrelations of returns, the assumption of covariance stationarity is needed. By confirming presence of ARCH and GARCH, we conclude this assumption is fulfilled only asymptotically and thus all statistical tests performed on returns might be questionable. Nevertheless, it is possible to modify the data in the following way. We introduce a new variable x_t, defined as

$$x_t = \frac{r_t}{h_t} \tag{2.10}$$

where r_t are the S&P 500 daily log returns as plotted in Figure 2.1 and h_t is the conditional variance of x_t. The variable x_t is already cleared off the GARCH effect and should be homoskedastic, covariance stationary and thus suitable for calculation of autocorrelations. However, this assumes that some specific form of ARCH or GARCH is actually the true model for conditional variance process, which is an issue beyond the scope of this work. We might also discuss the relevance of ACF of $|r_t|$ since these are likely to be non-stationary. Once again, however, this would account for a whole new paper. Thus, the main goal of the thesis is still reserved to comparing statistical properties of returns

generated by our model with the real-world observed data which are mainly obtained from raw log returns. Therefore, we only use raw returns in the rest of this work while focusing on standard stylized facts without further discussion of their statistical correctness.

Chapter 3

Ising model-based simulations

In this chapter, we present various approaches to modelling statistical properties that are characteristic to financial markets. At the beginning of each section we describe the general features of the simulated model and the dynamics it should follow. In following subsections we specify all the parameters and conditions under which the simulations were performed and decide whether the statistical properties of the simulated data are consistent with our expectations, i.e. whether they approximate the properties obtained from real data. For the sake of easier comparison, we placed tables with all relevant statistics together in Appendix A. Due to lack of space within body of the text, plots related to aggregate normality of returns are also included in Appendix A.

3.1 Simple Ising model

The Ising model is a ferromagnetic model in statistical mechanics consisting of discrete integer variables, called spins, organized in a lattice. The spins can only take values -1 and $+1$. In the economic terminology we can perceive the $+1$ spin as a buyer and -1 as a seller.

3.1.1 Description of the model

We start with the modified Ising model as first defined in Bornholdt (2001). The spins are organized in a squared lattice of $k \times k$ dimension. Let $s_i(t)$ denote the spin value on the i-th position in the lattice at time $t \in \{0, \ldots, T\}$, and $i \in \{1, \ldots, k^2\}$. The total magnetisation of the system at time t is given as

$$M(t) = \frac{1}{k^2} \sum_{j=1}^{k} s_j(t). \tag{3.1}$$

The total magnetisation can be viewed as the general mood or tendency of the whole market and is computed as a simple average of all spins. The intermediary between total magnetisation and neighbouring forces is the local field h_i. It is defined as

$$h_i(t) = \sum_{j=1}^{k} J_{ij} s_i(t) - \alpha s_i(t) \left| M(t) \right|, \tag{3.2}$$

where $\alpha > 0$ is the global coupling constant, $J_{ij} = 1$ for close neighbours of s_i (but not s_i itself) and $J_{ij} = 0$ otherwise. Whereas the first term tends to align the value of s_i with its neighbours, the second term (beginning with α) does the opposite. The higher the absolute total magnetisation is, the higher the chance it will force s_i to flip its spin is. We say the first term induces ferromagnetic order as it helps s_i imitate the behaviour of its neighbours while the second term induces anti-ferromagnetic order since it tempts s_i to change the value of its spin. The predefined constant α amplifies ($\alpha > 1$) or mitigates ($\alpha < 1$) the strength of the magnetisation.

The local field h_i summarizes the two counteracting inclinations of the agent i. The first is that the decision to buy should be made only if everybody else also decides to buy. Thus the first term in Equation 3.2 stands for the herding behaviour. The conflicting tendency (represented by the second term in the equation) is that only being in the minority of traders could possibly be beneficial for the agent. This is reminiscent of the minority game as described in Challet & Zhang (1997).

In our simulations, neighbours of $s_i(t)$ are four sites adjacent to $s_i(t)$ on both horizontal and vertical axis. The structure of neighbours on borderline sites is illustrated in Figure 3.1. This system is called Periodic boundary conditions.

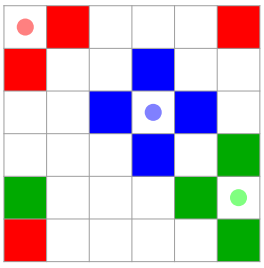

Figure 3.1: Periodic boundary conditions

The spins in the lattice are updated according to the following formulae:

$$s_i(t+1) = +1 \quad \text{with} \quad p = \frac{1}{1+\exp(-2\beta h_i(t))} \tag{3.3}$$

$$s_i(t+1) = -1 \quad \text{with} \quad 1-p \tag{3.4}$$

where p is the probability of a spin flip and β is a parameter. The probability p can be regarded as the Cumulative Distribution Function (CDF) of the independent variable h_i. The function is plotted in Figure 3.2 for various values of β parameter (unrelated to β in Equation 2.5). Higher value of β implies higher slope of CDF around zero. One can discern that for high β even a small increment in local field h_i leads to a great jump in probability of a spin flip to $+1$. In other words, β parameter partially regulates the measure of volatility during the simulation process as it controls the intensity of reaction to changes in the local field h_i.

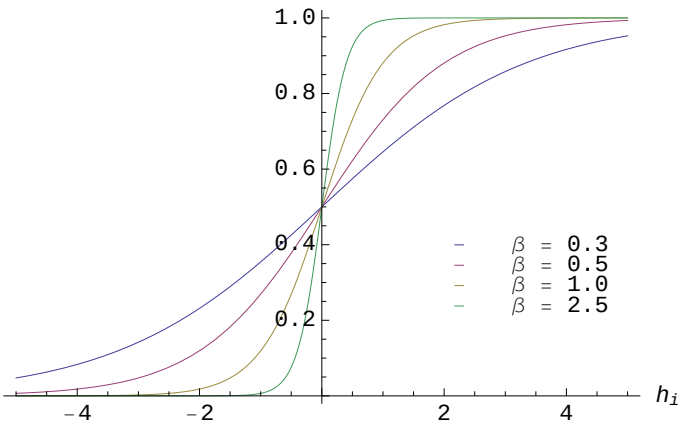

Figure 3.2: The CDF as defined in Equation 3.3 for four values of β.

To imitate the stylized facts observed on financial markets we are mainly interested in the distribution of returns. Since the total magnetisation as defined in Equation 3.1 can also take zero or negative values, we cannot use logarithmic returns as a measure of relative change. Instead the return is simply defined as

$$r_{\Delta t}(t) = M(t) - M(t-\Delta t), \tag{3.5}$$

as proposed by Sieczka & Hołyst (2007).

3.1.2 The simulation

In our first simulation we work with a 2-dimensional 32×32 lattice. In each time period $t \in \{0, \ldots, 10^6\}$, randomly chosen spin was updated according to Equation 3.3.[1] The parameters in our simulation are $\alpha = 20$ and $\beta = 1$.[2] A randomized matrix is used at time $t = 0$, i.e. $s_i(0) = 1$ with $p = 0.5$ for all i. Since this matrix is not a result of our simulation we allow for a warm-up stage ($t \leq 200\,000$).

The returns $r_{\Delta t}$ were computed in two different ways. First, we let $\Delta t = 100$. This time span is sufficient for observing patterns in return distribution. It is also more realistic to compare r_{100} to daily returns on stock exchange as there are many traders selling or buying during one specific time period. We further loosened this parameter by letting the lag Δt follow a Poisson process with mean 100. This allows for higher flexibility in the number of participants in each round. Also, letting 100 agents participate in trade during one time period allows us to refer to r_{100} or r_{100k} as "Lag 1" and "Lag k" returns, respectively. This notation will be used often throughout the rest of this work.

If the r_1 were computed, we would obtain a very limited set of information because the lattice at time t may differ from the lattice at time $t - 1$ only in one single site at most. For example, by computing r_{100} instead of r_1 we assume 100 traders participate in trade during one time period. Hence we obtain $10^6/100 = 10\,000$ returns less 2 000 dedicated to warm-up stage. These numbers may differ slightly in case of the Poisson process. In any case, all plots depicting returns always range from 2000 to approximately 10 000.

Figure 3.3 shows the standardized returns after 10^6 iterations. It does not alter fundamentally from Figure 2.1. The easily perceptible stages of stable returns are occasionally interrupted by short periods of high volatility (called intermittent phases). The main differences between the two plots is that the Poisson process governed time lag allows for even higher occasional returns.

Four snapshots taken during the simulation with r_{100} are shown in Figure 3.4. In the first and third picture, the agents are mainly influenced by their neighbours and total magnetisation plays marginal role in the decision process. In economic terms, this imitates relatively stable development on financial markets with no or weak influx of important information. The second and fourth

[1] Each spin in the lattice carries the same weight and thus can be chosen with equal probability $1/k^2$.

[2] These parameters were chosen to accent the phases of high volatility. For these values the simulation process most closely replicates the real-world data.

picture, however, show unstable intermittent phases where the role of neighbours is ambiguous and global magnetisation reaches extreme values. These are phases of high volatility that follow advent of new information into the market. After the new piece of information is embraced, stability is re-established and no extreme returns occur anymore.

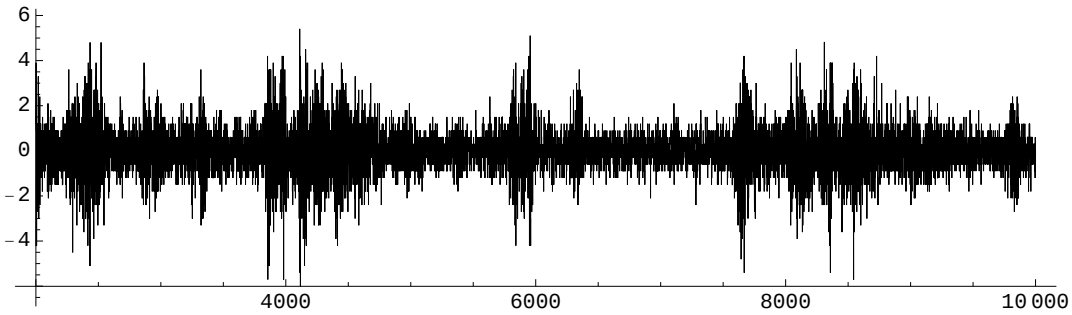

(a) Standardized returns r_{100} for t divisible by 100.

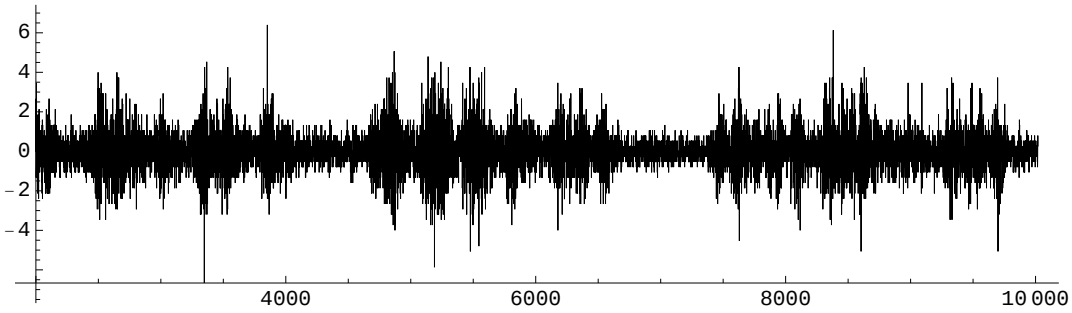

(b) Standardized returns $r_{\Delta t}$ for Δt following POI(100).

Figure 3.3: The returns for r_{100} and $r_{\Delta t}$ with $\Delta t \sim \text{POI}(100)$.

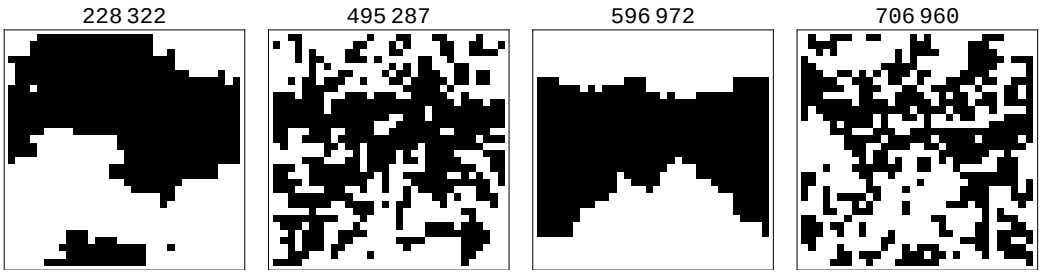

Figure 3.4: Four snapshots taken during stable (first and third picture) and intermittent phases (second and fourth picture). The number above represents the time at which the snapshots were taken.

We now investigate some of the statistical properties of returns. Following

the same procedure as in Subsection 2.1.1, we examine the form of autocorrelation function for $\tau \in \{1, \ldots, 150\}$.

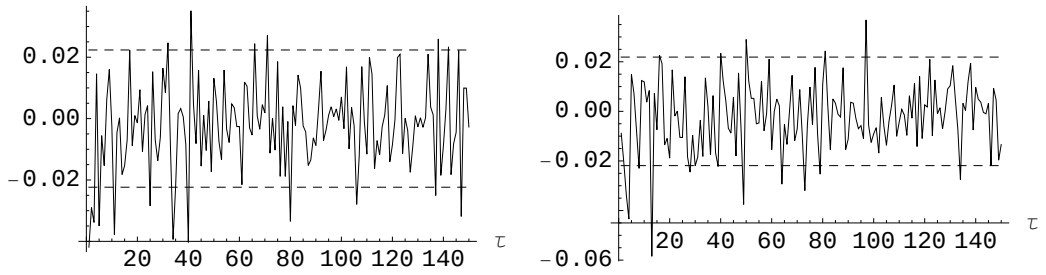

(a) ACF for r_{100} returns. (b) ACF for returns with lag $\sim$ POI(100).

Figure 3.5: Correlation of r_{100} and $r_{\Delta t}$ with $\Delta t \sim$POI(100).

As Figure 3.5 may suggest, the results do not differ much from the analysis of S&P 500 returns (Figure 2.2). The confidence interval, defined as $\pm 1.96/\sqrt{T}$ for $T = 8\,000$, equals to 0.022 for both r_{100} and Poisson process lag returns. The Ljung-Box test also helped to reject the null hypothesis that returns are independently distributed and no autocorrelation is present. The summary of all the important statistics is presented in Table A.1. Note that mean and standard deviation are not included since we work with standardized returns for easier comparison. The mean and standard deviation are always equal to zero and one, respectively. Instead we include the Ljung-Box Q-statistic for the 150th lag and the JB test statistic. The 5% critical region for the Q-statistic is $Q > \chi^2_{0.05}(150) = 179.581$. For large samples the critical region for JB statistic equals to $JB > \chi^2_{.05}(2) = 5.99$. For lag 1, the p-value of both tests is almost zero so the null hypotheses of no autocorrelation and normality are rejected. For higher lags the Q-statistic is still too high, but JB statistic gradually decreases for both regular and Poisson lags. Furthermore, for both types of returns at lag 64 the JB statistic equals to 4.108 and 2.599 with p-values of 0.103 and 0.219, respectively. In other words, the normality of returns is not rejected. This suggests that the distribution of returns approaches normal distribution with increasing lag. The plots of PDF of empirical distribution of returns are given in Figure A.1.

With the autocorrelation function for absolute returns already defined (Equation 2.3), we plot its value to see the speed of autocorrelation decay. We used semi-log scale plot to accent the difference between power-law and exponential fits for higher τ. Since autocorrelation decay is an asymptotic property, we focus on the fitting the reference functions for high values of τ. As such, the

Variable	A	α	B	β
r_{100}	0.201	0.152	0.249	0.007
POI(100)	0.193	0.193	0.251	0.009

Table 3.1: Parameters of the fitted power-law and exponential functions for the ACFs of absolute returns as shown in Figure 3.6.

power-law function seems to provide a more precise fit in the case of r_{100} as can be seen in Figure 3.6. For the Poisson lag, the correlation of absolute returns seems to decay too fast and is rather approximated by the exponential fit. The parameters for the fitted functions (described in Equation 2.5 and Equation 2.6) are given in Table 3.1. Note that values of both α parameters are fairly close to the α parameter of the S&P 500 returns (discussed in Subsection 2.1.2).

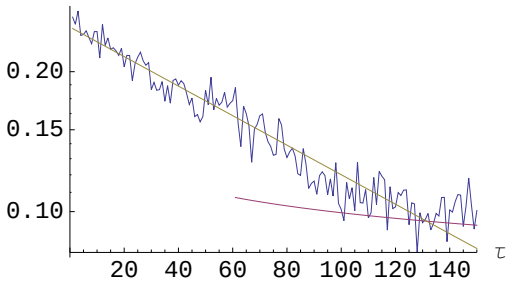

(a) ACF for r_{100} absolute returns.

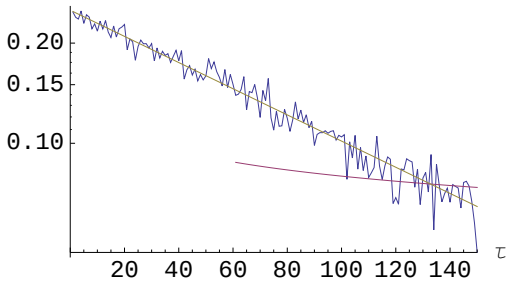

(b) ACF for absolute returns with lag following POI(100).

Figure 3.6: Correlation of absolute r_{100} and POI(100) lag returns from the simulation of simple the Ising model plotted on a semi-log scale. Power-law (purple) and exponential (green) fits are presented.

3.2 Simple model with a strategy spin

The local field as defined in Equation 3.2 is a special version of a generalized local field containing a strategy spin C_i:

$$h_i(t) = \sum_{j=1}^{k} J_{ij} s_i(t) - \alpha C_i(t) M(t). \tag{3.6}$$

The value of C_i depends on individual preferences of each agent. If $C_i(t) = +1 \ \forall t$, then the agent i always seeks to be a minority trader. For example, if $M(t)$ is positive, i.e. the majority of spins at time t are equal to 1, then the

second term in Equation 3.6 is negative and decreases the value of local field. Holding the first term of Equation 3.6 unchanged, this lowers the probability of opting for $+1$ at time $t+1$, as can be inferred from Equation 3.3.

The agents with strategy $C_i(t) = C_i = +1$ believe that the reason they struggle to be in the minority is that, unlike the majority of traders, they know the true, intrinsic value of the traded stock. These agents are called fundamentalists. Agents with an opposite strategy, i.e. $C_i(t) = C_i = -1$, are called chartists (or noise traders) and they prefer to act like the majority of traders as proposed by Bornholdt (2001). It is noteworthy that the first division of agents into fundamentalists and chartists was already common in many HAM models described in Chapter 1.

If we ran the test from Subsection 3.1.2 again with modified local field (as in Equation 3.6), we may obtain two different results. For $C_i = -1$ the results are not too interesting since both terms in Equation 3.6 induce ferromagnetic order. The returns r_k quickly fall to zero for any integer $k > 0$ and the magnetisation matrix consists only of 1's or -1's. For $C_i = 1$ the process stabilizes at certain spin structure. The returns are negligible and no sudden bursts of volatility as seen in Figure 3.3 may occur.

It would be more realistic to allow agents choose which strategy to follow in each round. It remains yet to define the transition rule between the two strategies. If we assume all agents choose the riskier strategy to expect higher potential returns, Bornholdt (2001) defines the transition rule as

$$C_i(t+1) = -C_i(t) \quad \text{if} \quad s_i(t)C_i(t)M(t) < 0. \tag{3.7}$$

Here, majority of traders always choose $C_i(t) = 1$ while minority of them always pick $C_i(t) = -1$. For example, if the majority of traders at time t decide to sell ($M(t) < 0$) and I happen to be part of that majority ($s_i(t) < 0$) but my strategy is that I should switch to minority of traders ($C_i(t) = 1$), then I have no reason to change my strategy for $t+1$ as it is already configured to do so. One can confirm that in this case $s_i(t)C_i(t)M(t) = (-1) \times 1 \times (-1) > 0$ and so $C_i(t+1) = C_i(t)$.

This process follows a similar dynamics as the one described in Subsection 3.1.2 but this time we keep track of the ratio of strategies used in each round. Figure 3.7 shows the usual r_{100} standardized returns along with the fraction of chartist strategies for the same simulation. The parameters of the simulation are the same as in Subsection 3.1.2. The returns with lag follow-

ing the Poisson process are not shown as they do not differ significantly from Figure 3.7.

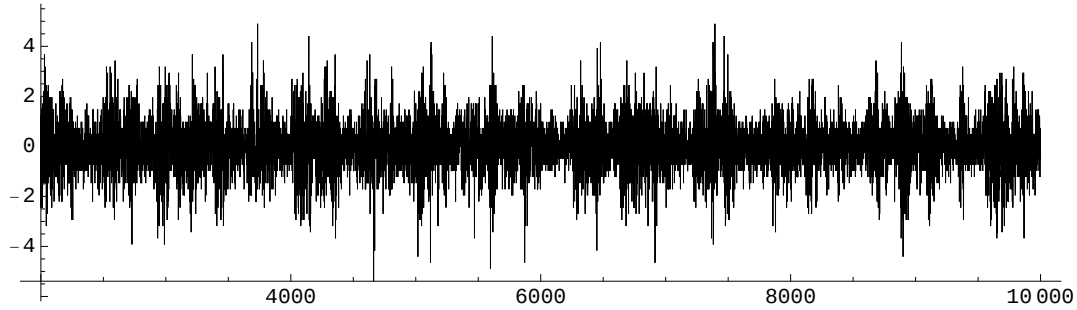

(a) The standardized r_{100} returns for the simulation with a strategy spin.

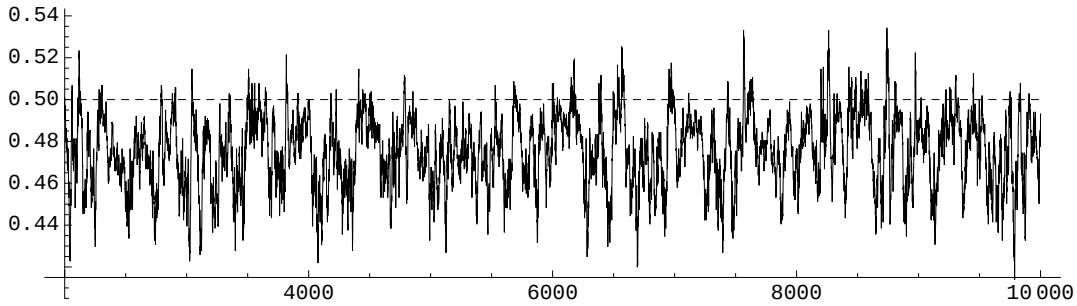

(b) Fraction of chartists in the same run.

Figure 3.7: Returns and fraction of chartists from the simple Ising simulation with a strategy spin.

Nevertheless, properties of the standardized returns from both types of simulations are summarized in Table A.2. The normality of returns is strongly rejected in both cases by the JB test for all lags but the trend in magnitudes of the test statistics is clearly decreasing. The highest p-value (i.e. the highest probability of not rejecting normality) is attained for Poisson returns of lag 16. In this case, the JB statistic equals 9.235 and p-value 0.0256 so we would not reject normality on 1% significance level. The empirical PDFs of the return distributions are plotted in Figure A.2.

The linear autocorrelation (as depicted in Figure 3.8, CIs included) is also statistically significant and, for lower lags, the returns do not resemble white noise as confirmed by the Ljung-Box Q-test. However, the magnitude of autocorrelation remains relatively low (below 0.1) as commonly observed in real financial data. For higher lags, the null of no autocorrelation is not rejected (look for $Q < \chi^2_{0.05}(150) = 179.581$).

Unlike in the previous simulation, the autocorrelation of absolute returns fades away very fast and eventually turns negative. This is captured in Fig-

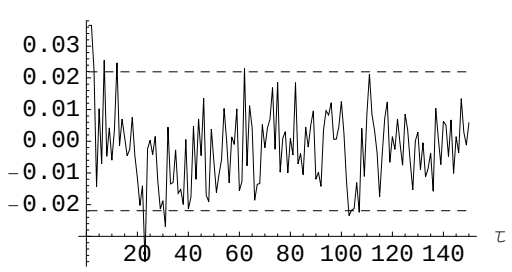

(a) ACF for r_{100} absolute returns.

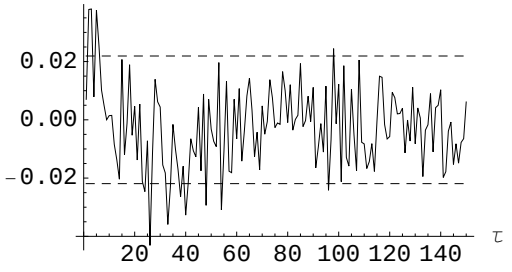

(b) ACF for returns with lag following POI(100).

Figure 3.8: Correlation of r_{100} and POI(100) lag returns from the simulation with a strategy spin.

ure 3.9. The simulation with Poisson lag generally shows higher correlation but it is still more precisely modelled by the exponential fit. Parameters of the fitted functions can be found in Table 3.2. Note that even though both α parameters are higher than in the case of S&P 500, they still lie within the commonly observed range $\langle 0.2, 0.4 \rangle$.

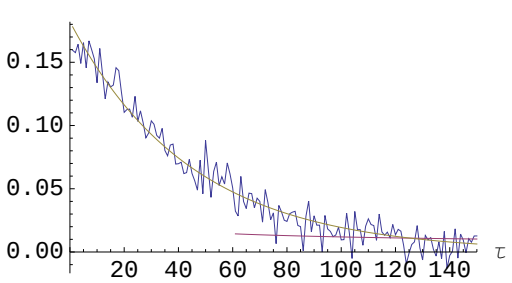

(a) ACF for r_{100} absolute returns.

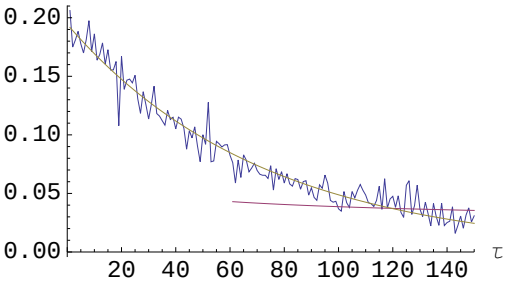

(b) ACF for absolute returns with lag following POI(100).

Figure 3.9: Correlation of absolute r_{100} and POI(100) lag returns from the simulation of the simple Ising model with a strategy spin. Power-law (purple) and exponential (green) fits are presented.

It is interesting to plot magnetisation matrices with corresponding strategy matrices side-by-side as we did in Figure 3.10. It is clear there is no regular pattern through which the strategy of an agent may determine its future action. The four snapshots show completely different combinations of the proposed strategies and actual decisions. The first snapshot shows inverse relationship. If the strategy is to sell, the decision actually made is to buy. The second picture in the row displays a stable magnetisation matrix even though the strategies of the agents are scattered arbitrarily. The first snapshot in the second row

Variable	A	α	B	β
r_{100}	0.066	0.372	0.182	0.022
POI(100)	0.103	0.212	0.194	0.014

Table 3.2: Parameters of the fitted power-law and exponential functions for the ACFs of absolute returns as shown in Figure 3.9.

exhibits perhaps the most expected connection between strategy and action, i.e. both matrices signal buy (or sell) at the same time. The last couple of snapshots captures the turbulent regime in both magnetisation and strategy matrix.

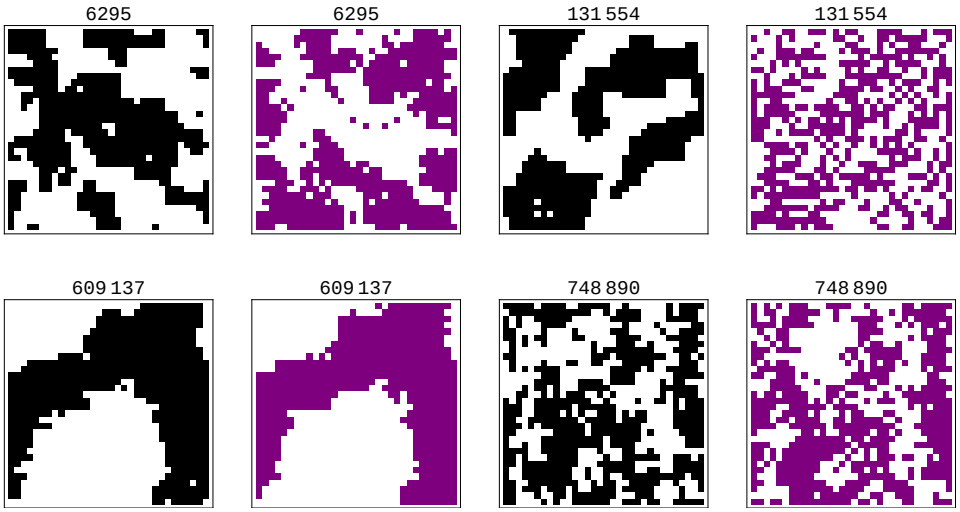

Figure 3.10: Four snapshots of magnetisation (black and white) and strategy (purple and white) matrices at different time periods.

3.3 Magnetisation dependent lag

So far, the number of trading participants in each round was independent of magnetisation, i.e. on average there were as many traders during zero magnetisation as there were during phases of high magnetisation. As the empirical data suggest (Karpoff (1987) and Jain & Joh (1988), for example), the number of traders on real stock markets increases with price. Judging from reviewed literature, no model in a hitherto published paper included this feature. In

this section, we present our own model that successfully accommodates this property.

First, we design the auxiliary function ω that inputs a value of magnetisation and returns a corresponding number of trade participants. Since the mean value of magnetisation is zero and extreme values of absolute magnetisation do not usually exceed 0.5, the function should be very sensitive to changes at very low magnetisation levels. The ω function is presented in Equation 3.8.

$$\omega(x) = \lfloor a\sqrt[b]{x^2} + c \rfloor, \tag{3.8}$$

where $a, b, c > 0$ are parameters and $\lfloor \dots \rfloor$ stands for the floor function rounding each real number to the largest previous integer. The c parameter sets the minimum number of participants when the magnetisation is zero, a regulates the range of the function and b determines the sensitivity to changes in the independent variable. Figure 3.11 shows the plot of ω function for two different domains.

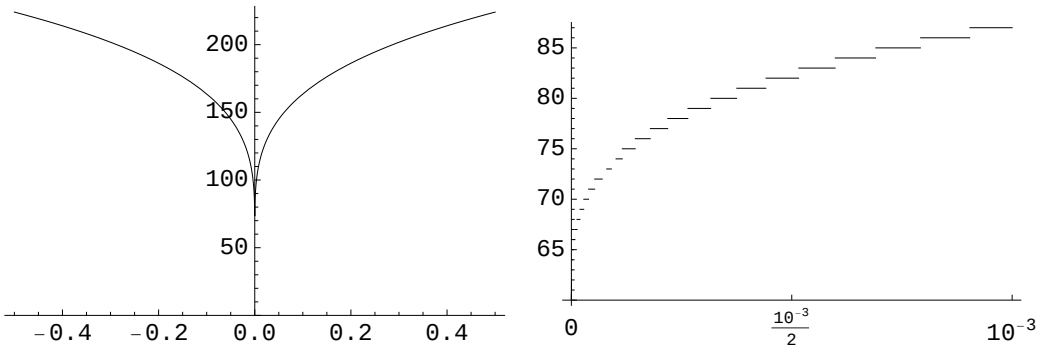

(a) The $\omega(x)$ function for $x \in (-0.5, 0.5)$

(b) Detail of the $\omega(x)$ function for $x \in (0, 10^{-3})$.

Figure 3.11: The $\omega(x)$ function for various values of x. The parameters in this case are: $a = 200$, $b = 7$, $c = 60$.

x	0	0.05	0.1	0.2	0.5	1
$\omega(x)$	60	144	163	186	224	260

Table 3.3: Value of the ω function for selected x.

We ran the simulation with the same parameters as in Subsection 3.1.2. This time, however, we set the number of iterations so that we end up with exactly 8 000 returns. At each time period when $r_{\Delta t}$ was computed the ω function found the appropriate number of participants for the next round. This number

was then used as the λ parameter for the Poisson process. In other words, we worked with $r_{\Delta t}$ where $\Delta t \sim \text{POI}(\omega(M(t)))$.

Figure 3.12 displays the $r_{\Delta t}$ returns along with plot of lags (or number of participants) that were used for calculation of these returns. Due to higher variability in lag selection the returns reach much higher values than in previous simulations. The descriptive statistics of returns and magnetisation dependent lags are included in Table A.3. Both normality and no autocorrelation null hypotheses are strongly rejected by the JB and Ljung-Box tests, respectively, for low lags. As we have already seen in the previous simulations, both JB and Ljung-Box statistics decrease with higher lags. For lag 64, the p-value of the JB test statistic equals to 0.445. In other words, normality is not rejected on any usual significance level and the stylized fact of aggregate normality is confirmed. The empirical PDFs are plotted in Figure A.3

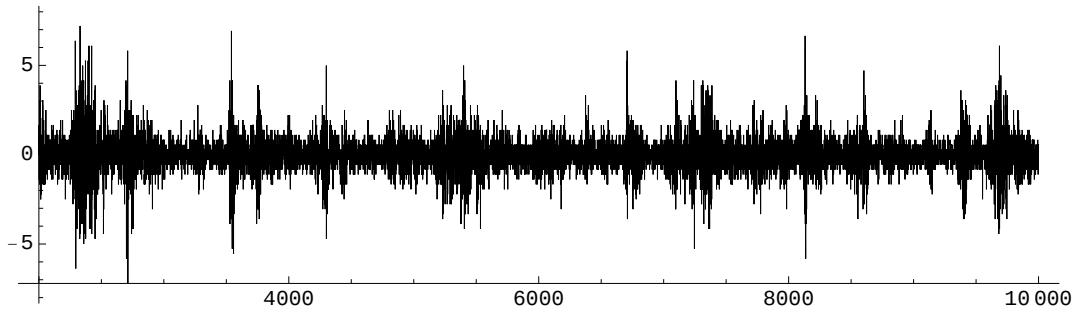

(a) The $r_{\Delta t}$ with Δt determined by ω function and Poisson process.

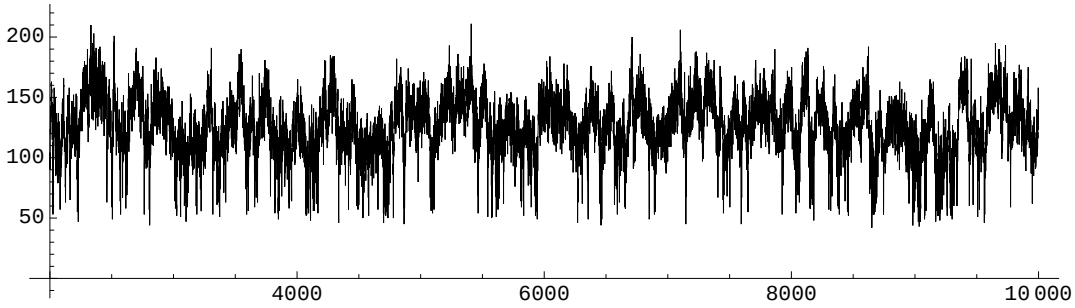

(b) Plot of lags used for calculation of the returns above.

Figure 3.12: The plots of returns and lags from the simulation with a magnetisation dependent lag.

The graphs of ACFs of regular and absolute returns are displayed in Figure 3.13. It is clear the correlation of absolute returns decays exponentially at the beginning but follows no obvious pattern for higher lags. Nevertheless, the α parameter was estimated to be equal to 0.322 which is, once again, in hand with the mentioned stylized facts.

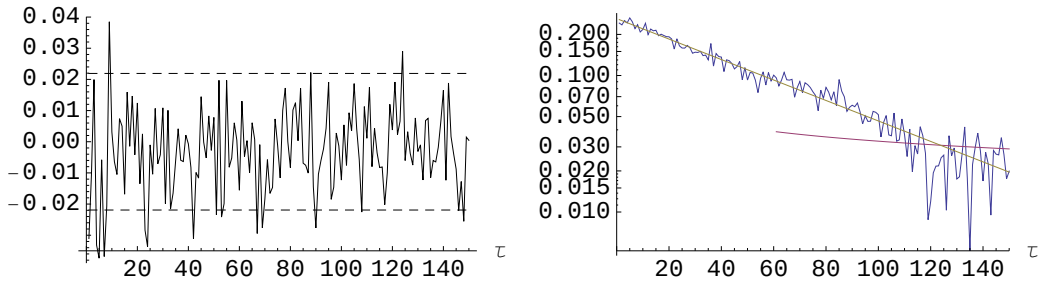

(a) ACF for $r_{\Delta t}$ returns. (b) ACF for $r_{\Delta t}$ absolute returns.

Figure 3.13: Plots of ACFs for linear and absolute returns. The power-law and exponential fits are presented.

Variable	A	α	B	β
$r_{\Delta t}$	0.145	0.322	0.261	0.017

Table 3.4: Parameters of the power-law and exponential fits of the ACFs of absolute returns as plotted in Figure 3.13.

3.4 Threshold model

3.4.1 Description of the model

Sieczka & Hołyst (2007) present their own model highly motivated by the work of Bornholdt (2001). The main difference lies in the introduction of the third possible action agents may perform. Besides *buy* and *sell*, the option *stay inactive* is now available. While the local field of each agent is defined in the same manner as in Equation 3.2, the spin update is now governed by the following equation:

$$s_i(t+1) = \mathrm{sgn}_{\lambda|M(t)|}\left(\sum_{j=1}^{T} J_{ij}s_j(t) + \sigma\nu_i(t)\right), \tag{3.9}$$

where sgn_q is the threshold function, defined as:

$$\mathrm{sgn}_q(n) = \begin{cases} -1, & \text{if } x < -q \\ 0, & \text{if } -q \leq x < q \\ 1, & \text{if } q \leq x \end{cases} \tag{3.10}$$

The $\lambda > 0$ parameter influences the width of range in which agents choose to stay inactive. The ν is normally distributed random variable with zero mean and unit variance representing unpredictable behaviour of each agent. Note that, as opposed to models in previous sections, this is the first time any kind

of stochastic term is used in the model. The σ parametrizes the strength of this variable.

3.4.2 The simulation

The simulation of the threshold model is similar to one performed in Subsection 3.1.2. Randomly chosen spin is refreshed according to Equation 3.10 over the span of 10^6 iterations. We calculated both r_{100} as well as Poisson lagged returns with mean and variance of the Poisson distribution equal to 100. At the beginning of the simulation, we repeatedly allowed for a warm-up phase of 200 000 iterations. The standardized r_{100} returns and $r_{\Delta t}$ returns for $\Delta t \sim \text{POI}(100)$, $\lambda = 15$ and $\sigma = 1$ are plotted in Figure 3.14.[3] The main difference between the simple Ising returns are the units of the vertical axis. Due to possibility of withholding from any trade activity, the threshold model returns are much lower than in the previous case. The usual descriptive statistics for different values of λ are included in Table A.4. Normality is rejected via JB test for all λ. For $\lambda = 5$ and $\lambda = 40$, the null hypothesis of no autocorrelation is not rejected because the value of Q test statistic does not exceed 179.581 (a 5% critical value for the χ^2 distribution with 150 degrees of freedom). For $\lambda = 15$ and different lags the JB test still rejects normality in all cases for regular returns but decreases rapidly with increasing lag in the case of Poisson returns. For lag 64, the JB test statistic and corresponding p-value equal to .286 and 0.858, respectively. These results suggest that the stylized fact of aggregate normality holds for threshold model returns as well.

Figure 3.15 presents four snapshots of magnetisation matrix at various time periods. The first and third snapshots confirm that agents choose no action on the borderline where the influence of sellers and buyers is in balance. The second and fourth snapshot show fairly stable phases with majority of dormant traders.

The autocorrelation function of returns for three different lags is in Figure 3.16. For lag 2, the autocorrelation of returns remains clung to zero and waver rather gently. The reason for this is that such a low lag does not allow total magnetisation to change drastically in such a short time period. With increasing lag, both linear and absolute ACFs manifest much higher variabili-

[3]For low λ the results are similar to the simple Ising model described in the previous subsection. For high λ the belt of inactivity is too wide (caused by high q in Equation 3.10) resulting in negligible yet stable returns. For $\lambda = 15$ the simulation pronounces the singular features of the model.

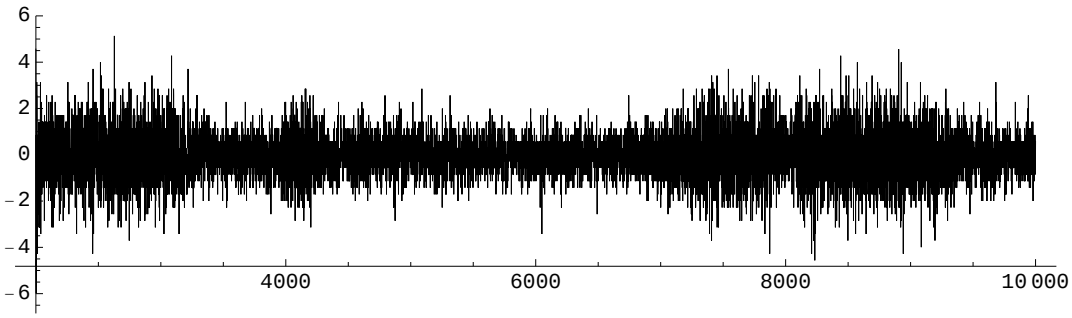

(a) Standardized threshold returns r_{100}.

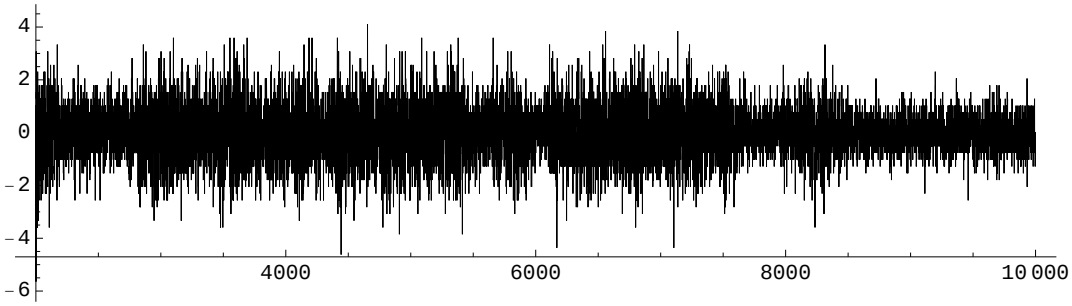

(b) Standardized threshold returns $r_{\Delta t}$ with $\Delta t \sim$ POI(100).

Figure 3.14: Threshold r_{100} returns and returns with lag following POI(100).

ty and sudden bursts. As seen from the summary Table A.4, the presence of autocorrelation is statistically confirmed by Ljung-Box test for all lags except lag 16 in r_{100}. For this lag the null hypothesis of no autocorrelation cannot be rejected. To avoid overcrowded charts, the usual power-law and exponential fits are not included in Figure 3.16. However, the usual α and β parameters for Lag 2 and Lag 4 can be read from Table 3.5. Parameters for Lag 16 are not included since the ACF of absolute returns often takes negative values.

Even though JB test strongly rejects normality for all lags in r_{100}, the test

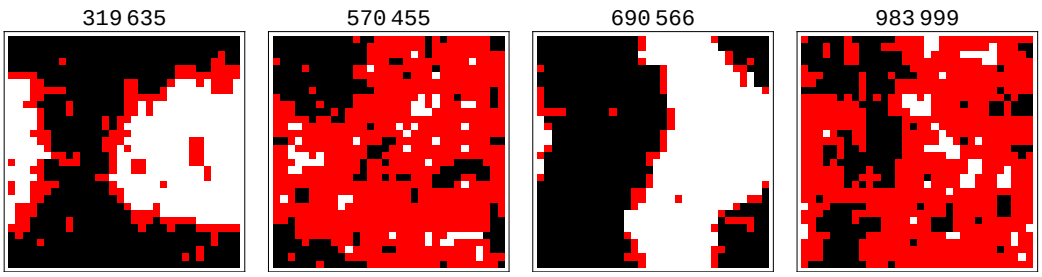

Figure 3.15: Four snapshots of the magnetisation matrix at different time periods. In addition to black (+1) and white (−1), red colour stands for inactive participants (0).

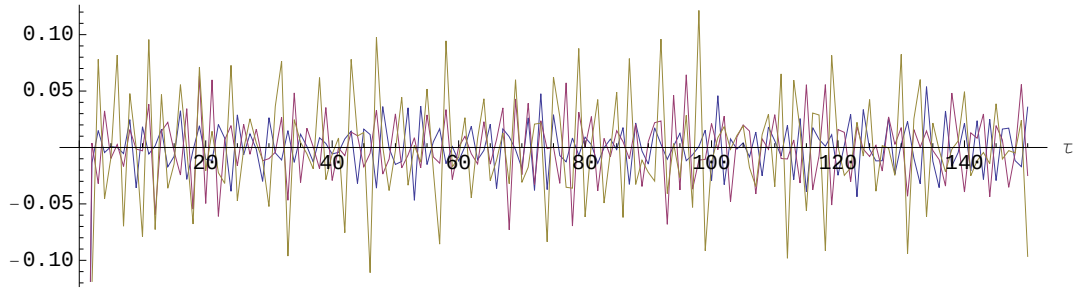

(a) The autocorrelation of threshold model returns for $\lambda = 15$ and time lags 2, 4 and 16 represented by blue, purple and green lines, respectively.

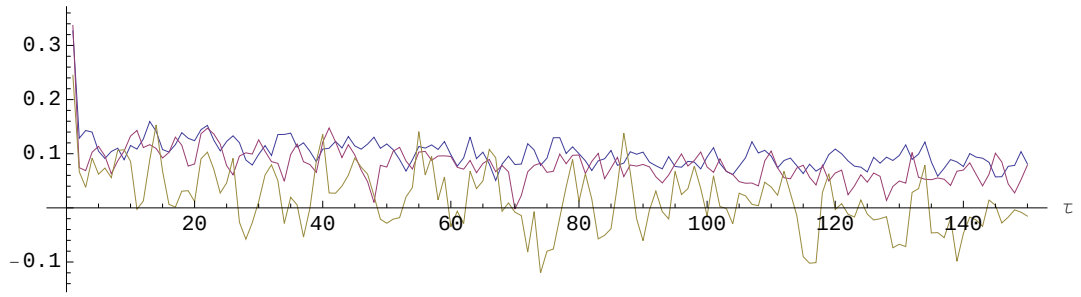

(b) The autocorrelation of threshold model absolute returns for $\lambda = 15$ and time lags 2, 4 and 16 represented by blue, purple and green lines, respectively.

Figure 3.16: The ACF for some values of λ (upper picture) and for fixed λ but diferent lags (lower picture).

Variable	A	α	B	β
Lag 2	0.101	0.0352	0.132	0.00373
Lag 4	0.089	0.0832	0.119	0.0056

Table 3.5: Parameters of the power-law and exponential fits of the ACFs of absolute returns as plotted in Figure 3.16.

statistic gradually decreases with increasing lag in case of Poisson returns and for lag 64 the p-value of the test equals to 0.858. This, again, favours validity of aggregate normality.

3.5 Serial versus parallel update

In all of the examples and simulations above, the magnetisation matrix in each time period t was updated serially which means that only one randomly chosen spin was altered according to rules given by Equation 3.3 and Equation 3.4. The alternative is to update all sites at once. This is called parallel updating. In our simulation described in previous chapters, this would mean to update 32×32 sites in one time period t.

We run a simulation similar to the one in Subsection 3.1.2 except that now parallel updating of all sites is employed. Since any magnetisation matrix may look completely different in two subsequent time periods, the warm-up phase does not take but only a few iterations. The choice of α and β parameters showed no special impact on the dynamics of the model. In fact, for any reasonable value of α and β, the model always reaches a balanced state within few iterations. The value of total magnetisation $M(t)$ quickly falls to zero and fluctuates in almost imperceptible amounts. The same holds for returns r_1. Figure 3.17 shows four snapshots taken during the simulation. In this case, the process consisted of mere 10^3 iterations. Since the process becomes repetitive after few dozens of iterations, increasing the number of iterations does not uncover any new phenomena.

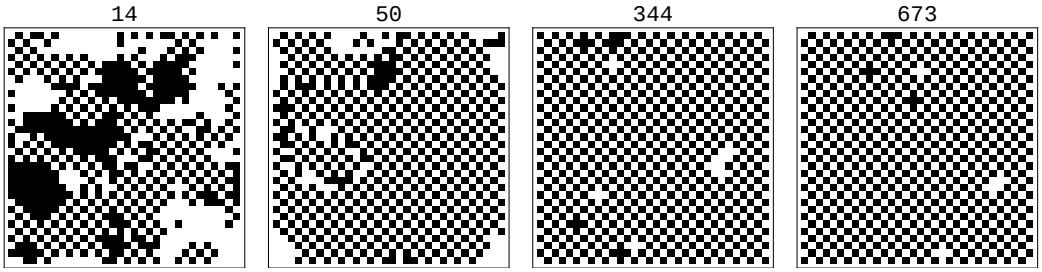

Figure 3.17: Four snapshots of magnetisation matrix using parallel updates. The parameters are: $T = 10^3$, $\alpha = 20$ and $\beta = 1$. For $t > 50$, the probability of a spin flip on every site is approximately 0.995.

The spin structure resembles a checker board that inverts its values at every time period, i.e. $s_i(t) = -s_i(t-1)$ almost for every higher t. The inversion of spin values is due to fact that the anti-ferromagnetic term in Equation 3.2 takes very low values as $M(t)$ approaches zero. Consequently, only neighbouring spins play role in the decision making process of each agent. Once all the neighbours adopt the same strategy, the agent is forced to follow them in the next round.

Since this can be said about every spin in the magnetisation matrix, alternating checker boards are the only possible outcome of the simulation. The decrease in magnetisation due this stabilisation can be seen in Figure 3.18.

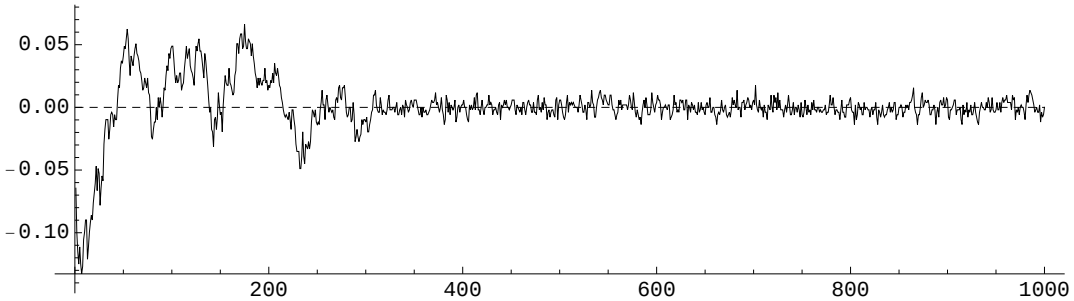

Figure 3.18: Fall and stabilisation of magnetisation (measured on the vertical axis) for $T = 1000$.

It is interesting to plot magnetisation matrices after stabilisation of the simulation for different structures of neighbours. In Figure 3.19, we present four patterns to which the simulation converges for different number and organisation of neighbours (in Figure 3.17 the final stable pattern resembles a checker board). The snapshots suggest that asymmetry in neighbour arrangement leads to distorted organisation of the magnetisation matrix.

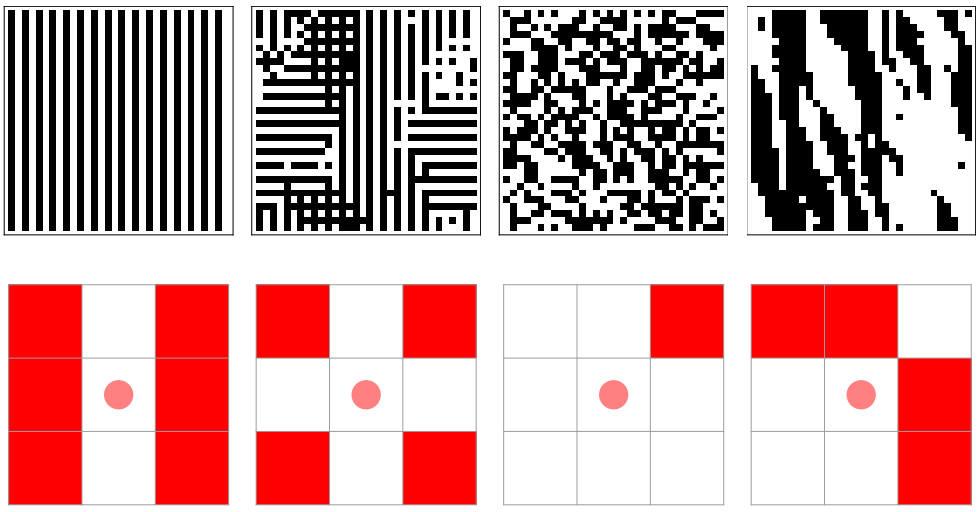

Figure 3.19: The impact of various neighbour configuration (in red) on magnetisation matrix (upper row) when using parallel updating of the sites. The parameters are $\alpha = 20$ and $\beta = 1$. All snapshots were taken after 500+ iterations.

Chapter 4

Conclusion

The goal of this thesis was to create a model that would reproduce the same statistical properties as can be observed in real financial data. The selection of the properties followed Cont (2001) and thereat listed stylized facts.

First, we investigated presence of stylized facts in S&P 500 index which serves as a proxy for the whole US market. We confirmed that S&P 500 does contain most of the common properties such as slow decay in autocorrelation of absolute returns (more likely exhibiting power-law than exponential decay), heteroskedasticity of returns, heavy tails and non-Gaussianity of the return distribution, negative skewness and positive excessive kurtosis. The test of aggregate normality did not confirm that the distribution of $r_{\Delta t}$ returns approaches normal with increasing Δt (for definition of returns see Equation 2.1), even though the JB test statistic gradually decreased with increasing lag. Strictly speaking, the stylized fact stating that there is no autocorrelation in raw returns does not hold for S&P 500 data either. In this case, the autocorrelations of daily returns showed to be statistically significant but from the economic viewpoint they are rather negligible as the value of ACF never exceeded 0.1.

In the main chapter of the thesis we ran various Ising model simulations in effort to obtain a deterministic process featuring properties listed above. We started with the simplest form of the model with adding extra features after each execution. As a result, 17 different simulations were carried out. Results of some important simulations were thoroughly discussed (e.g. the simple Ising model with a fixed lag) while others only served to illustrate certain peripheral phenomena (such as simulation with parallel updates under different structures of neighbours).

The results of the simulations are generally inconclusive. While some of

the easily replicable stylized facts were encountered in each simulation (departure from normality or positive excessive kurtosis and heavy tails), others were achieved only in some cases (power-law decay or aggregate normality). Where successful, however, these models (except threshold model) were able to replicate the stylized facts without any random, erratic noise as it is common in other techniques used to model financial markets. The ACF of absolute returns showed exponential or even faster decay in most of the simulations. Important exceptions are the simple Ising model with r_{100} returns and the magnetisation dependent lag model where number of trade participants depends on actual magnetisation in each round. Here the ACF of absolute returns showed slow, power-law decay for higher values of time lags suggesting long memory in volatility. Generally, with increasing randomness incorporated in the model (for example by letting the lag follow Poisson distribution) the ACF of absolute returns declines much faster than in other cases.

There are manifold possibilities for further research. The author originally pondered performing additional simulations of Ising three dimensional model as in Preis *et al.* (2009). However, these simulations are much more computationally demanding with only small added value to the goal of this thesis. Substantial tuning of the model can be accomplished through altering the decision rule (recall strategy spin and Section 3.2) or neighbours arrangement. To our best knowledge, none of the reviewed models in the literature was designed to yield returns with negative skewness (neither was ours, and any such results are coincidental). This would mean to break the point reflection of the CDF in Equation 3.3 and allow for more likely occurrence of negative values. Ad hoc construction of such a CDF is not a simple task. In fact, none of the most frequently used distributions allow for negative skewness but one can find inspiration in definitions of generalized hyperbolic, skew normal or skew student distributions. It would be also worthwhile to design a search engine that would input certain descriptive statistics extracted from real financial data as independent variables and return the most appropriate parameters (e.g. α and β in Equation 3.2 and Equation 3.3) for the simulated model in order to best approximate the real data. This would, once again, require much stronger computational power at hand.

Bibliography

Andersen, T. & T. Bollerslev (1997): "Intraday Periodicity and Volatility Persistence in Financial Markets." *Journal of Empirical Finance* **4**: pp. 115–158.

Baruník, J., L. Vácha, & M. Vošvrda (2009): "Smart Predictors in the Heterogeneous Agent Model." *Journal of Economic Interaction and Coordination* **4(2)**: pp. 163–172.

Blattberg, R. C. & N. J. Gonedes (1974): "A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices." *The Journal of Business* **47(2)**: pp. pp. 244–280.

Bollerslev, T. (1986): "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* **31(3)**: pp. 307–327.

Bollerslev, T., R. Y. Chou, & K. F. Kroner (1992): "ARCH Modeling in Finance : A Review of the Theory and Empirical Evidence." *Journal of Econometrics* **52(1-2)**: pp. 5–59.

Bornholdt, S. (2001): "Expectation Bubbles in a Spin Model of Markets: Intermittency from Frustration across Scales." *International Journal of Modern Physics A* **12(5)**: pp. 667–674.

Brock, W. & P. De Lima (1995): "Nonlinear Time Series, Complexity Theory, and Finance." *Working papers 9523*, Wisconsin Madison - Social Systems.

Brock, W. A. & C. H. Hommes (1998): "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model." *Journal of Economic Dynamics and Control* **22(8-9)**: pp. 1235–1274.

Caldarelli, G., M. Marsili, & Y. C. Zhang (1997): "A Prototype Model of Stock Exchange." *Quantitative Finance Papers cond-mat/9709118*, arXiv.org.

Calvet, L. & A. Fisher (2008): *Multifractal volatility: Theory, Forecasting, and Pricing*. Academic Press advanced finance series. Academic Press.

Campbell, J. Y., A. W. Lo, A. C. MacKinlay, & R. F. Whitelaw (1998): "The Econometrics Of Financial Markets." *Macroeconomic Dynamics* **2(04)**: pp. 559–562.

Challet, D., M. Marsili, & Y.-C. Zhang (1999): "Modeling Market Mechanism with Minority Game." *Quantitative Finance Papers cond-mat/9909265*, arXiv.org.

Challet, D., M. Marsili, & Y.-C. Zhang (2001): "Stylized Facts of Financial Markets and Market Crashes in Minority Games." *Quantitative Finance Papers cond-mat/0101326*, arXiv.org.

Challet, D. & Y.-C. Zhang (1997): "Emergence of Cooperation and Organization in an Evolutionary Game." *Physica A: Statistical Mechanics and its Applications* **246(3–4)**: pp. 407 – 418.

Chowdhury, D. & D. Stauffer (1998): "A Generalized Spin Model of Financial Markets." *Quantitative Finance Papers cond-mat/9810162*, arXiv.org.

Comte, F. & E. Renault (1996): "Long Memory Continuous-time Models." *Journal of Econometrics* **73**: pp. 101–150.

Cont, R. (2001): "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues." *Journal of Economic Literature* **1**: pp. 223–236.

Cont, R., M. Potters, & J.-P. Bouchaud (1997): "Scaling in Stock Market Data: Stable Laws and Beyond." *Science & Finance (CFM) working paper archive 9705087*, Science & Finance, Capital Fund Management.

Eberlein, E., U. Keller, & K. Prause (1998): "New Insights into Smile, Mispricing, and Value at Risk: The Hyperbolic Model." *The Journal of Business* **71(3)**: pp. 371–405.

Engle, R. (1982): "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica* **50(4)**: pp. 987–1007.

Engle, R. (2001): "GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics." *Journal of Economic Perspectives* **15(4)**: pp. 157–168.

Gopikrishnan, P. & V. Plerou (1999): "Scaling of the Distribution of Fluctuations of Financial Market Indices." *Physical Review* **60**: pp. 5305–5316.

Gopikrishnan, P., V. Plerou, X. Gabaix, & H. E. Stanley (2000): "Statistical Properties of Share Volume Traded in Financial Markets." *Quantitative finance papers*, arXiv.org.

Hommes, C. H. (2006): "Heterogeneous Agent Models in Economics and Finance." In L. Tesfatsion & K. L. Judd (editors), "Handbook of Computational Economics," volume 2 of *Handbook of Computational Economics*, chapter 23, pp. 1109–1186. Elsevier.

Iori, G. (1999): "Avalanche Dynamics and Trading Friction Effects on Stock Market Returns." *Int. J. Modern Phys. C* **10**: pp. 1149–1162.

Jain, P. C. & G.-H. Joh (1988): "The Dependence between Hourly Prices and Trading Volume." *Journal of Financial and Quantitative Analysis* **23(03)**: pp. 269–283.

Jarque, C. M. & A. K. Bera (1980): "Efficient Tests for Normality, Homoscedasticity and Serial Independence of Regression Residuals." *Economics Letters* **6(3)**: pp. 255–259.

Kaizoji, T., S. Bornholdt, & Y. Fujiwara (2002): "Dynamics of Price and Trading Volume in a Spin Model of Stock Markets with Heterogeneous Agents." *Physica A* **316**: pp. 441–452.

Kantelhardt, J. (2009): *Encyclopedia of Complexity and Systems Science*. Springer.

Karpoff, J. M. (1987): "The Relation between Price Changes and Trading Volume: A Survey." *Journal of Financial and Quantitative Analysis* **22(01)**: pp. 109–126.

Ljung, G. M. & G. E. P. Box (1978): "On a Measure of Lack of Fit in Time Series Models." *Biometrika* **65(2)**: pp. 297–303.

Lux, T. & M. Marchesi (1998): "Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents." *Discussion Paper Serie B 437*, University of Bonn, Germany.

Mandelbrot, B. (1963): "The Variation of Certain Speculative Prices." *The Journal of Business* **36**: p. 394.

Mantegna, R. N. & H. E. Stanley (1999): *Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge University Press, 0 edition.

Meucci, A. (2010): "Quant Nugget 2: Linear vs. Compounded Returns – Common Pitfalls in Portfolio Management." *GARP Risk Professional* **2**: pp. 49–51.

Pagan, A. R. (1996): "The Econometrics of Financial Markets." *Journal of Empirical Finance* **3(1)**: pp. 15–102.

Preis, T., P. Virnau, W. Paul, & J. J. Schneider (2009): "GPU Accelerated Monte Carlo Simulation of the 2D and 3D Ising Model." *Journal of Computational Physics* **228(12)**: pp. 4468–4477.

Sieczka, P. & J. A. Hołyst (2007): "A Threshold Model of Financial Markets." *Quantitative Finance Papers 0711.3106*, arXiv.org.

Simon, H. A. (1978): "Rational Decision-Making in Business Organizations." *Nobel Prize in Economics documents 1978-1*, Nobel Prize Committee.

Sornette, D. & W.-X. Zhou (2005): "Importance of Positive Feedbacks and Over-confidence in a Self-Fulfilling Ising Model of Financial Markets." *Quantitative Finance Papers cond-mat/0503607*, arXiv.org.

Appendix A

Tables and plots

Notation remark: As already mentioned in Subsection 3.1.2, returns for the first lag are labelled r_{100} since we allowed 100 agents to trade in one time period, returns r_{100k} are labelled as "Lag k". To avoid confusion, only notation in the form "Lag k" will be used in all summary tables.

Variable	**Min.**	**Max.**	**Skewn.**	**Kurt.**	**JB**	**Q(150)**	**N**
Lag 1	-6.092	5.485	.0112	6.361	3774.87	267.16	8001
Lag 2	-5.707	5.272	-.1341	6.041	1560.88	262.65	3999
Lag 4	-6.397	6.074	-.0230	5.969	741.76	190.64	1999
Lag 16	-4.687	3.46	-.1752	4.842	76.235	120.54	499
Lag 64	-2.917	2.937	.0670	3.782	4.108	N/A	124
POI Lag 1	-5.971	6.513	.119	5.708	2476.8	258.80	8020
POI Lag 2	-5.132	5.319	.1408	5.500	1062.43	188.69	4009
POI Lag 4	-3.729	4.967	.2512	4.830	303.103	154.24	2004
POI Lag 16	-3.95	3.867	-.0706	4.652	59.873	218.32	500
POI Lag 64	-2.504	3.291	.1717	3.531	2.599	N/A	124

Table A.1: Summary statistics for the simple Ising simulation for different lags. The upper part of the table describes the r_{100k}, $k \in \{1, 2, 4, 16, 64\}$, the lower part summarizes results from the simple Ising model with Poisson lag.

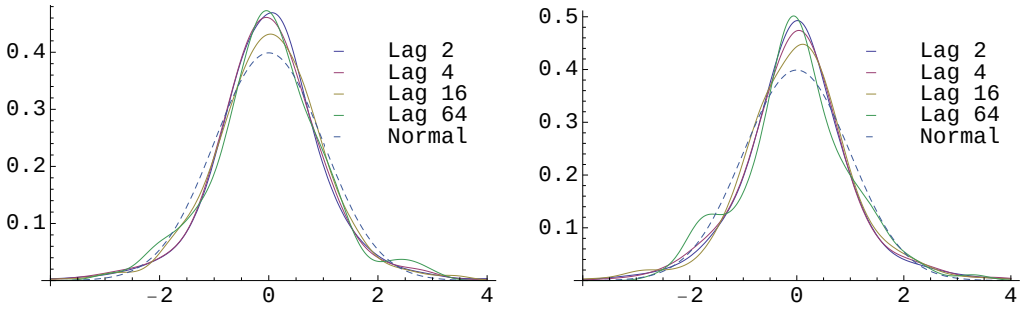

(a) Distributions for r_{100k}, where k is the value of the lag.

(b) Distributions for lag following Poisson distribution.

Figure A.1: Standardized empirical distributions of simple Ising model returns for different lags. The dashed curve belongs to PDF of standard normal distribution.

Variable	Min.	Max.	Skewn.	Kurt.	JB	Q(150)	N
Lag 1	-5.391	4.902	-.029	4.686	951.59	182.51	8001
Lag 2	-4.175	4.00	-.0299	4.142	219.234	211.924	3999
Lag 4	-4.603	3.54	-.1510	4.061	102.573	163.444	1999
Lag 16	-3.92	3.1	.0365	4.040	23.822	144.111	499
Lag 64	-2.951	3.963	.334	4.606	18.373	N/A	124
POI Lag 1	-5.105	5.103	-.031	4.713	983.24	246.84	8004
POI Lag 2	-5.268	4.889	-.1437	4.773	540.507	240.66	4001
POI Lag 4	-4.043	5.555	-.1316	4.846	292.733	161.54	2000
POI Lag 16	-3.132	3.427	.01288	3.644	9.235	195.37	499
POI Lag 64	-4.284	2.236	-.5357	5.038	31.569	N/A	124

Table A.2: Summary statistics for simple Ising model with a strategy spin.

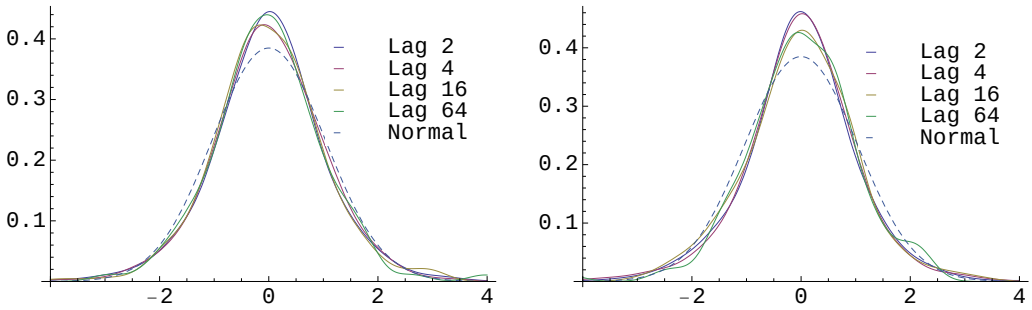

(a) Distributions for r_{100k}, where k is the value of the lag.

(b) Distributions for lag following Poisson distribution.

Figure A.2: Standardized empirical distributions of simple Ising model returns with a strategy spin for different lags. Dashed curve belongs to standard normal distribution.

Variable	Min.	Max.	Skewn.	Kurt.	JB	Q(150)	N
Lag 1	-6.304	7.206	.1149	6.843	4953.71	243.28	8000
Lag 2	-8.201	7.771	.1324	7.368	3204.74	233.549	3999
Lag 4	-6.387	5.478	.0342	6.654	1122.92	192.862	1999
Lag 16	-4.811	4.65	-.0747	5.280	113.051	135.708	499
Lag 64	-2.642	2.887	.1782	3.303	1.406	N/A	124

Variable	Mean	Median	Min.	Max.
Δt	125	126	42	211

Table A.3: Summary statistics for returns from the simulation with a magnetisation dependant lag. The lower table provides basic descriptive statistics for lags Δt produced by the ω function.

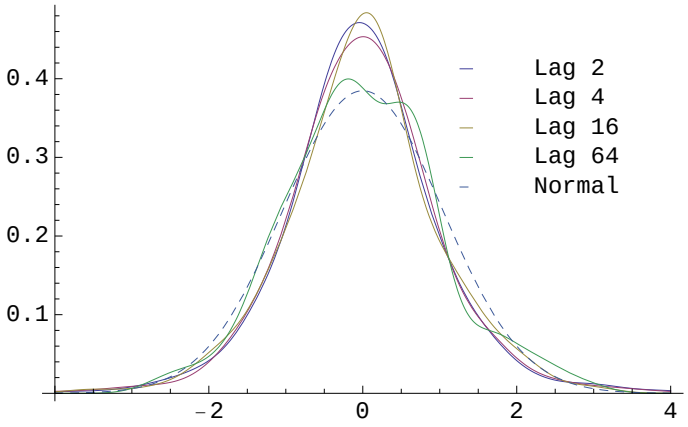

Figure A.3: Empirical distributions of returns from simulation with magnetisation dependant lag.

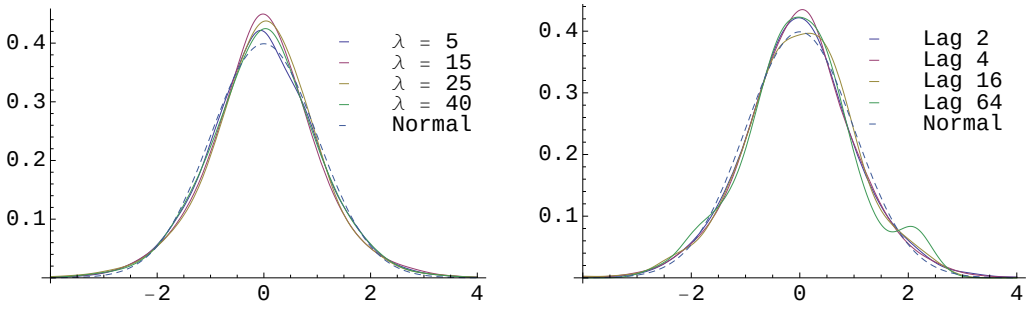

(a) Standardized smoothed empirical distributions of the threshold model returns for multiple values of λ. The dashed line belongs to PDF of standard normal distribution. All empirical distributions show signs of fat tails.

(b) Standardized smoothed empirical distributions of threshold model returns with $\lambda = 15$ and Δt ranging from 4 to 64 (derived from simulation with Poisson lag). The dashed line represents PDF of standard normal distribution.

Figure A.4: The empirical distributions of threshold model returns.

Variable	Min.	Max.	Skewn.	Kurt.	JB	Q(150)	N
$\lambda = 5$	-4.327	4.333	0.017	3.556	103.959	145.51	7993
$\lambda = 15$	-4.704	4.19	.0278	3.812	221.913	229.93	8001
$\lambda = 25$	-5.516	4.693	-0.203	4.538	844.24	231.24	7987
$\lambda = 40$	-4.667	4.89	-0.0455	3.761	197.049	176.86	8003

Variable	Min.	Max.	Skewn.	Kurt.	JB	Q(150)	N
Lag 1	-4.817	4.817	-.0094	4.116	416.424	183.91	8001
Lag 2	-4.084	4.299	-.00938	4.021	174.99	225.73	3999
Lag 4	-5.544	4.465	.0289	4.341	151.791	197.473	1999
Lag 16	-2.688	4.199	.2847	4.009	29.159	143.53	499
Lag 64	-4.165	4.885	.6469	8.839	211.067	N/A	124
POI Lag 1	-4.704	4.19	.0278	3.812	221.913	229.93	8001
POI Lag 2	-4.16	4.173	.0437	3.663	79.4953	242.26	3999
POI Lag 4	-4.115	4.549	.0294	3.829	58.46	241.31	1999
POI Lag 16	-4.093	2.622	-.2118	3.593	11.617	244.15	499
POI Lag 64	-2.664	2.408	.0593	3.142	.286	N/A	124

Table A.4: The upper table shows summary statistics for the threshold simulations with POI(100) lagged returns and different value of λ. The other table summarizes usual regular and Poisson returns for $\lambda = 15$ and different lags.

Bachelor Thesis Proposal

Author	Pavel Dvořák
Supervisor	PhDr. Ladislav Krištoufek
Proposed topic	From Microscopic Rules to Macroscopic Phenomena: Ising Model in Finance

Topic characteristics In this bachelor thesis, we focus on ability of the Ising model to mimic basic stylized facts of financial markets. Ising model is a physics model of ferromagnetism with local and global magnetic forces. Applied to finance, magnetized particles represent the market participants. These are influenced by behaviour (magnetization) of the neighbouring particles (market participants with similar preferences, trading strategies, investment horizons, etc.) as well as the magnetization of the whole field (market). From this very simply described system, we examine whether its behaviour is qualitatively similar to the behaviour of the real financial markets. We test whether this artificial system exhibits standard stylized facts of the financial markets - rapidly decaying autocorrelations, volatility clustering, long-range dependence in volatility, fat tails and non-zero skewness.

Charakteristika práce v češtině V této bakalářské práci se zaměřujeme na Isingův model a jeho schopnost vysvětlit základní stylizovaná fakta o finančních trzích. Isingův model je fyzikální koncept feromagnetismu s lokálními a globálními magnetickými silami. V řeči ekonomie, magnetizované částice představují aktéry na finančních trzích. Ti jsou ovlivněni hlavně chováním (magnetizací) sousedních částic (neboli aktéry s podobnými preferencemi, obchodní a investiční strategií, atd.) a navíc také magnetizací celého silového pole (trhu). V rámci tohoto zjednodušeně popsaného modelu zjišťujeme, zda vykazuje kvalitativně shodné vlastnosti jako skutečné finanční trhy. Dále testujeme, zda tento uměle vytvořený systém potvrzuje obecná stylizovaná fakta finančních

trhů jako např. rapidně klesající autokorelace, shluky volatilit, dlouhá paměť ve volatilitě, těžké chvosty a nenulová šikmost rozdělení výnosů.

Outline

1. Introduction
2. Stylized facts in empirical data
3. The Ising Model simulations
4. Conclusion

Core bibliography

1. Bornholdt, S. (2001): "Expectation Bubbles in a Spin Model of Markets: Intermittency from Frustration across Scales." *International Journal of Modern Physics A* **12(5)**: pp. 667–674.
2. Cont, R. (2001): "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues." *Journal of Economic Literature* **1**: pp. 223–236.
3. Herrero, C. (2004): "Ising Model in Scale-free Networks: A Monte Carlo Simulation." *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics* **69(6)**: p. 067109.
4. Kaizoji, T., S. Bornholdt, & Y. Fujiwara (2002): "Dynamics of Price and Trading Volume in a Spin Model of Stock Markets with Heterogeneous Agents." *Physica A* **316**: pp. 441–452.
5. Krawiecki, A. (2009): "Microscopic Spin Model for the Stock Market with Attractor Bubbling on Scale-free Networks." *Journal of Economic Interaction and Coordination* **4**: p. 213–220.
6. Queiros, S., E. Curado, & F. Nobre (2007): "A Multi-Interacting-Agent Model for Financial Markets." *Physica A* **374**: pp. 715–729.
7. Sieczka, P. & J. A. Hołyst (2007): "A Threshold Model of Financial Markets." *Quantitative Finance Papers 0711.3106*, arXiv.org.
8. Sornette, D. & W.-X. Zhou (2005): "Importance of Positive Feedbacks and Overconfidence in a Self-Fulfilling Ising Model of Financial Markets." *Quantitative Finance Papers cond-mat/0503607*, arXiv.org.