INTEGRATE-AND-FIRE MODELS OF FINANCIAL MARKETS

PhD Thesis

presented by

DAVID CHRISTIAN AUREL MEINE
Imperial College London
Department of Physics

supervised by

Prof. Dr. Dimitri Vvedenksy

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Physics of Imperial College London and the Diploma of Imperial College London

2025

My life seemed to be a series of events and accidents. Yet when I look back, I see a pattern.

Benoît B. Mandelbrot

Declaration

I herewith certify that all material in this dissertation which is not my own work has been properly acknowledged.

David Meine

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Licence (CC BY-NC-ND). Under this licence, you may copy and redistribute the material in any medium or format, provided that you credit the author and do not use the material for commercial purposes.

You may not remix, transform, or build upon the material, and you may not distribute modified versions of the work.

When reusing or sharing this work, you must clearly indicate the licence terms by naming the licence and linking to the licence text.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

PUBLICATIONS

Articles in peer-reviewed journals:

Meine, D. C. & Vvedensky, D. D. Emergence of speculation in a hierarchical agent-based model. *Physica A: Statistical Mechanics and its Applications* **620**, 128619 (2023)

Applies a hierarchical structure to the Lux–Marchesi model and demonstrates how such hierarchies lead to the emergence of speculation; see Chap. 1.

Hussain, S. A., Meine, D. C. A. & Vvedensky, D. D. Integrate-and-fire model of disease transmission. *Phys. Rev. E* **110**, 014305 (1 July 2024)

Introduces an integrate-and-fire approach to disease transmission; forms the basis of Chap. 6.

Manuscripts under review:

Meine, D. C. A. & Vvedensky, D. D. *Network Structures in an Integrate-and-Fire Model of Financial Traders* Manuscript under review. 2025

Currently under review; describes the core integrate-and-fire trader model, covers parts of Chaps. 1 and 2.

Conference contributions:

Meine, D. C. A. & Waring, P. *Open Source Research and Network Analysis* in *IAEA International Symposium on Safeguards: Re-inforcing Trust and Verification in an Evolving World (IAEA-CN-303)* (Vienna, Austria, 2022)

Presentation + paper (refereed) at the 2022 IAEA Safeguards Conference resulting from my work for Ridgeway Information.

Meine, D., Fair, K. R. & Guerrero, O. *Agent-Based Model of Endogenous Labour Market Flows* in *Data Natives 2023* Presented during the 'Society' session (City, University of London, Apr. 2023)

Joint work undertaken for the Alan Turing Institute and the UK Department for Business & Trade.

Fair, K. R., Meine, D. & Guerrero, O. *Endogenous labour flow networks and the skills evolution of the UK labour market* in *CCS'24 – Conference on Complex Systems 2024* (Exeter, United Kingdom, 2024)

Same project; preliminary results, work not yet published when this thesis was submitted.

Meine, D. & Kumar, N. *Modeling wealth-driven mobility dynamics* in *fineconets@CCS2024: 11th Workshop on Complexity in Economics & Finance, Satellite of CCS'24* (Exeter, United Kingdom, Sept. 2024)

Model presented; not yet published when this thesis was submitted.

ABSTRACT

This thesis introduces a novel theoretical and computational framework that bridges concepts from physics, complexity science, neuroscience, and behavioural economics to describe financial market dynamics using an integrate-and-fire paradigm. Motivated by the limitations inherent in classical economic theories – especially the efficient market hypothesis and the oversimplification of homogeneous, representative agents – the work explicitly models financial markets as emergent phenomena resulting from the interactions of heterogeneous agents. Traders in this framework integrate information, including market signals and social influences, and decide when decision thresholds are reached, analogous to neuronal firing events.

Through computational modelling and systematic analysis, the thesis reveals several significant findings. It demonstrates that neuronal-inspired decision mechanisms, such as inhibition and information decay (reflecting irrationality and limited memory), are critical for replicating widely observed financial phenomena, including speculative bubbles, sudden market crashes, fat-tailed return distributions, and volatility clustering. Further investigations establish clear relationships between hierarchical information structures (information silos), and mispricing, highlighting how network topology profoundly influences market stability and price dynamics.

Expanding beyond purely financial applications, the thesis generalises the integrate-and-fire framework into a versatile approach capable of describing broader social dynamics. With this approach, we manage to reproduce wealth distributions of traders in financial markets. Remarkably, this generalised framework also successfully captures complex epidemiological patterns, accurately reproducing real-world data from infectious disease outbreaks, including the COVID-19 pandemic, and capturing features such as secondary infection peaks and intervention effects.

Overall, the results underscore the importance of micro-level heterogeneity and network topology in generating complex, non-equilibrium market phenomena, providing new insights into the interplay between individual decision-making and aggregate market behaviour.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my heartfelt gratitude to my supervisor, Prof. Dimitri Vvedensky, for his guidance, mentorship, and unwavering support over the past years. It was through his influence that I was first introduced to agent-based modelling. My Bachelor's project with him during my Erasmus convinced me that a career in science would nurture my curiosity and creativity. Later, during the Covid period, when I reached out about a Master's project, his balanced mix of insightful advice and creative freedom enabled me to develop the initial model that now forms the foundation of this thesis. Our weekly meetings during my PhD always left me encouraged - he was always receptive to my ideas and used his knowledge and experience to guide me to the right path. Truly, without his influence, I would not have embarked on this PhD journey.

I am also deeply thankful to the remarkable community at the Complexity Centre at Imperial College. I never imagined working in such an interdisciplinary environment, where fresh ideas are constantly exchanged. Collaborating with and learning from individuals from around the world has not only broadened my perspective but also gifted me lifelong friendships. Special thanks to my friends Santiago Musalem Pinto, Nishanth Kumar, and Gezhi Xiu. Every coffee break chat, every memorable moment at the Union bar, and our inspiring office discussions have not only brightened my days but also deeply enriched my journey. Your friendship and insight mean the world to me. I also wish to extend my gratitude to my friends from Hamburg, with particular appreciation for Caroline Nowoczyn and Marvin Leusch, whose invaluable assistance in reviewing my thesis has been much appreciated.

Finally, and most importantly, I want to convey my deep gratitude to my family. To my grandmother, Christa, whom I should call more often; to my brother, Simeon, who is always there to intellectually challenge or cheer me up – it turns out the best time for both is after midnight; and especially to my parents, whose boundless support, love, and patience have been the cornerstone of my success. The values you have instilled in me, along with your unwavering encouragement, have made all of this possible and have always made me feel at home.

CONTENTS

1 Introduction 1
1.1 Physics and Economics have historically inspired one another 1
1.2 The digital Revolution . . . 7
1.3 Challenging classical economics . . . 9
1.3.1 Fat-tailed returns . . . 11
1.3.2 Rationality of agents . . . 11
1.3.3 Excess returns . . . 14
1.4 The advent of agent-based modelling . . . 17
1.5 Structure in Financial Markets . . . 20
1.5.1 Ising-like models . . . 21
1.5.2 Lux-Marchesi model . . . 22
1.5.3 Hierarchical Models . . . 26
1.5.4 Networks in financial markets . . . 30
1.5.5 Network Applications in Economics and Finance . . . 36
1.6 Behavioural Science and Economics . . . 38
1.7 A new paradigm: neurological basis for decision making . . . 41
1.8 Integrate-and Fire models . . . 43
1.9 Outline of the thesis . . . 46
2 First model 49
2.1 Adapting models to the integrate-and-fire approach . . . 49
2.1.1 The firing event . . . 52
2.1.2 Market dynamics . . . 53
2.1.3 Model analysis . . . 56
2.2 The role of Parameters . . . 60
2.2.1 The effect of market price information ω_1 . . . 60
2.2.2 The effect of network information . . . 67
2.2.3 The effect of the decay variable τ . . . 70
2.3 Network Analysis . . . 78
2.3.1 Network structure and trader behavior . . . 79
2.3.2 Centrality measures . . . 79
2.3.3 Current flow closeness centrality and cluster size distribution . . . 79
2.3.4 Energy . . . 82
2.3.5 Connectivity and criticality . . . 84

2.3.6 Price change scaling . . . 85
2.3.7 Hurst exponent . . . 86
2.4 Visualisation and Calibration . . . 87
2.4.1 Visualisation . . . 87
2.4.2 Calibration . . . 87
2.5 Simulation Results . . . 88
2.5.1 Network structure and community formation . . . 89
2.5.2 Spontaneous rise of chartists . . . 92
2.5.3 Extreme volatility clustering . . . 95
2.5.4 Infectious Behavior . . . 97
2.5.5 Unifying opinion . . . 98
2.5.6 Sudden shifts in behavior . . . 99
2.6 Conclusion . . . 101
3 Inhibitory Effects - Irrationality drives bubbles 103
3.1 Introduction . . . 103
3.2 Updated model . . . 106
3.3 Simulation setup and results . . . 107
3.4 Agent potentials: precursors for market phase transitions . . . 112
3.5 Entropy . . . 115
3.5.1 Bethe approximation . . . 115
3.5.2 Entropy and inhibition . . . 118
3.6 Return distributions and other measures . . . 120
3.6.1 Average potentials and trader distributions . . . 122
3.6.2 The relationship between inhibition and decay . . . 123
3.7 Discussion and Conclusion . . . 125
4 Hierarchical Influences 129
4.1 Hierarichal Network model definition . . . 132
4.1.1 Cluster Formation and Network Structure . . . 132
4.1.2 Dynamics of the Hierarchical Model . . . 133
4.1.3 Agent Potentials and States . . . 133
4.1.4 Currents and Information Integration . . . 133
4.1.5 Updating Potentials and States . . . 134
4.1.6 Market Dynamics . . . 135
4.2 Simulation and Results . . . 136
4.2.1 Simulations . . . 136
4.2.2 Phase diagrams . . . 141
4.2.3 Entropy . . . 143
4.3 Discussion and Conclusion . . . 144
5 A new framework for social interaction 149

5.1 Towards a generalized framework . . . 150
5.1.1 The role of system . . . 151
5.1.2 Agents . . . 151
5.1.3 Information . . . 153
5.1.4 Simulation Steps . . . 154
5.2 Application: Extended Model . . . 154
5.2.1 Model basics . . . 155
5.2.2 General: Multi-Stock . . . 155
5.2.3 System: Order book . . . 157
5.2.4 Agents: Portfolio . . . 159
5.2.5 Agents: Order Price . . . 162
5.2.6 Information: Network and Cross-Influences . . . 165
5.3 Results . . . 167
5.4 Conclusion . . . 174
5.4.1 Order book microstructure . . . 175
5.4.2 Network behaviour . . . 175
6 Practical Applications 177
6.1 Integrate-and-Fire Model of Disease Transmission . . . 177
6.1.1 Introduction . . . 178
6.1.2 Applying the integrate-and-fire model to infections . . . 180
6.2 Agent States and Variables . . . 180
6.3 Model Dynamics . . . 181
6.3.1 Transmission of Infectious Matter . . . 181
6.3.2 Reception and Accumulation of Infectious Matter . . . 181
6.3.3 Infection Event (Fire) . . . 182
6.3.4 Recovery and Immunity Update . . . 182
6.3.5 Immunity Decay . . . 183
6.4 Parameter Definitions . . . 183
6.5 Initial Conditions . . . 184
6.6 Model Summary . . . 184
6.7 Key Features and Effects of Parameters . . . 184
6.8 Calibration and Results . . . 185
6.8.1 UK Covid Data . . . 185
6.8.2 Simulation Results and Analysis . . . 188
6.8.3 Impact of Interventions . . . 192
6.9 Modelling Multiple Strains with the Integrate-and-Fire Model 195
6.9.1 Introduction . . . 195
6.9.2 Model Implementation . . . 195
6.9.3 Application to Influenza Spread in Canada . . . 196

6.9.4 Results and Discussion 199
6.10 Conclusion and Future Work 202
7 Summary, Conclusion and Outlook 205
7.1 Future work 208
7.1.1 Fundamentals 208
7.1.2 Calibration 209
7.1.3 Additonal applications 210
7.1.4 Final comments 217

Bibliography 219

NOMENCLATURE

β_{ij}	Inhibitory effect
ΔE	Excess Demand
Γ	Excitement
ω_1	influence of market price
ω_2	influence of network
ω_3	influence of random information
ω_4	influence of hierarchy
ρ_{ij}	Cross-influence matrix
$\mathbf{h}_i$	Asset holdings of stock i
$\mathbf{I}_+$	Optimist current
$\mathbf{I}_-$	Pessimist current
$\mathbf{I}_{\text{net}}$	Network current
$\mathbf{I}_{\text{price}}$	Price current
$\mathbf{I}_{\text{rand}}$	Random current
$\mathbf{I}_f$	Fundamentalist current
$\mathbf{I}_{t,sys}$	System Current
ϑ	Firing threshold
$A(n,t)$	Action of agent n at time t
I	Current (carrying information)

I_{sys}	System current
K	Connection matrix
M	Number of stocks
N	Number of total agents
n_+	Number of optimists
n_f	Number of fundamentalists
n_f	Number of pessimists
n_f	Time constant
p	Market price
p_f	Fundamental stock price
p_{ask}	Ask order price
p_{bid}	Bid order price
R	Aeverage real returns
r	Nominal dividens
S_t	System state at time t
T	Trigger
t_f	Firing time
u	Action potential
u_r	Resting potential

1

INTRODUCTION

1.1 PHYSICS AND ECONOMICS HAVE HISTORICALLY INSPIRED ONE ANOTHER

From its very inception, economics was inspired by physics. Adam Smith, by many regarded as the founder of the field of modern economics, published his revolutionary work *An Inquiry into the Nature and Causes of the Wealth of Nations* in 1776 [1]. This work is considered to be the first modern work of economics. In it, he lays out many of the core principles that drive economics today: The division of labour, free markets, wage theory, and the conclusion that the wealth of a nation is driven by the labour and goods produced by its people. However, to gain a deeper understanding of the philosophy of Smith, it is worthwhile to look at another previously written, posthumously published work of his, *The History of Astronomy* [2]. While both works seem to tackle very different disciplines, economics and astronomy, they share not just a similar methodology, but *The History of Astronomy* essentially outlines the scientific philosophy behind his work on the wealth of Nations as well.

While Smith examines the development of Astronomy "from ancient times to Copernicus", he pays special attention to the principles behind these discoveries. He explains how the theories evolved, their philosophical principles, and their reasoning. One overarching element he notes is how they were shaped by the ambition of researchers to impose structure and coherence onto observed phenomena. He touches upon many scholars, including Ptolemy, Copernicus and Brahe and how astronomical belief systems moved from imaginative conjecture to more logical and evidence-based theories. He also discusses another scholar who had a similar impact on his field to the one Smith had on economics: Issac Newton.

Smith has been known to have been an admirer of Newton as well as having taken great inspiration from him [3, 4]. He was especially inspired by Isaac Newton's seminal work Philosophiae Naturalis Principia Mathe-

matica [5], in which Newton lays out the foundations of classical mechanics, describing his three laws of motion and universal gravitation. The influence of Newton's work is not limited to Smith's essay on astronomy but is also visible in *The Wealth of Nations*. Both share an emphasis on mathematical and deductive reasoning: While Newton clearly laid out the laws of motion and gravity in a mathematical formulation, Smith also deduced his economic theories from basic principles and described mathematical relationships. Both Newton and Smith based their theories on empirical observations, which were not as prevalent in economic approaches before Smith.

Newton describes the Universe as operating mechanistically and following the laws of classical mechanics. Similarly, Smith viewed the economy as a mechanistic system that follows the natural market laws. In fact, the systemization approach – manifested in Smith's metaphor of the invisible hand – may best demonstrate an early inspiration that economics has taken from physics.

An even earlier instance of the natural sciences inspiring economics can be found in the concept of utility-based preferences, which is central to many economic, as well as game theoretic models where were pioneered in the St. Petersburg paradox by mathematician Nicholas Bernoulli [6]. It was crucial in inspiring Ramsey and Savage two hundred years later in formulating the expected utility hypothesis [7, 8].

In the late 19th century, Maxwell [9] and Boltzmann [10] developed the theory of gases and statistical mechanics. By taking a statistical perspective on the molecules that make up a gas, they showed how one can derive macroscopic properties like temperature and pressure, thus showing how order and equilibrium can emerge from the random movement of molecules. Francis Edgeworth saw the analogy to actors in markets and conceived the concept of the economy as a system of exchanging particles in which every particle tries to maximise utility [11]. All of these actions would lead to economic equilibrium. Similarly, Alfred Marshall [12], when examining demand and supply in markets, decided to use demand and supply curves to describe the aggregate behaviour of the market. Just like molecules in a gas, the behaviour of actors in markets leads to an equilibrium price, Marshall found. As such, the thermodynamic insights of physics helped shape another critical concept of economics. Fig. 1.1 shows equilibrium processes both for a thermodynamic system and a fundamental economics

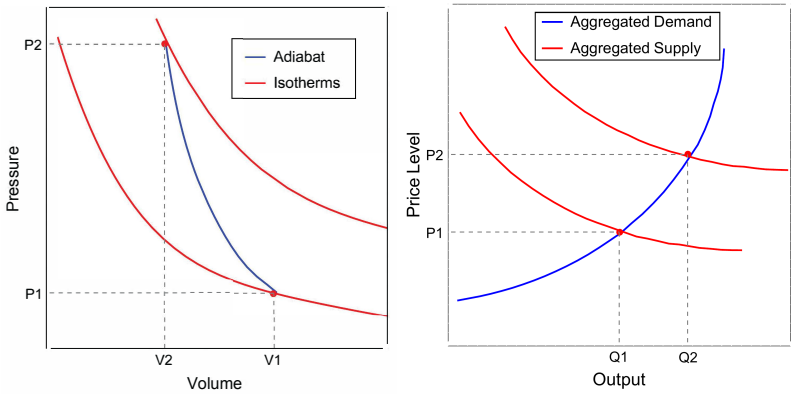


FIGURE 1.1: Left: Pressure-volume diagram of an adiabatic process. Right: In the AS-AD model, a shift in aggregated demand leads to a new equilibrium, which lies at the intersection of aggregrated demand and supply.

model, the AS-AD model. Both describe the aggregated behaviour of their respective subjects.

However, one should not make the mistake of assuming that economics simply adapted the best ideas from physics. On the contrary, the symbiotic relationship between economics and physics was indeed a two-way street. The two disciplines were constantly influencing and benefiting from one another. An early example of this can be found in the work of Vincent Pareto. In 1906, while studying both wealth and income distributions, he discovered that goods were distributed according to something known as the 80/20 rule [13]. The 80/20 rule states that 20 percent of people own 80 percent of income and wealth.

This rule applies to a broad range of phenomena, where a minority of people are responsible for a majority of the impact. Pareto went even further in his investigation: He discovered that the distribution of wealth and income are not random occurrences but follow a mathematical pattern, the power law. Power laws are often used to model the frequency of events as a function of their magnitude. They describe the probability p of an event of magnitude x, as a function of a normalisation constant C and the exponent α.

$$p(x) = Cx^{-\alpha} \tag{1.1}$$

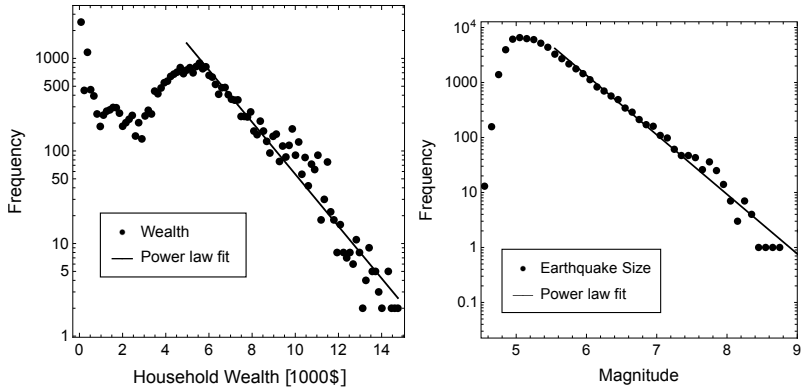


FIGURE 1.2: Left: Wealth distribution of US households, data taken from [14]. Right: Histogram of earthquake magnitudes from 1976 to 2020 [15, 16]. It is striking that both distributions follow a Pareto distribution.

Pareto himself realised that their application goes far beyond wealth and describes many patterns found in nature and society, like the population distribution of cities. Today, power laws are used to describe a broad range of phenomena, from earthquakes [17], meteorites [18], biological extinctions [19], and network connectivity [20] to Bose-Einstein condensates [21]. See Fig. 1.2 for a comparison of earthquakes and US household wealth.

In 1905, Albert Einstein finally found a mathematical description for the random movement of particles in a fluid [22], which was discovered by Robert Brown almost 80 years before in 1827 [23]. By deriving the diffusion equation, which describes the probability distribution of the location of a particle over time, he was able to solve this problem and relate it to a random walk. This concept has been applied to a wide area of physics, from neutron scattering to cosmology theories and even open quantum systems subject to dissipative processes. However, 5 years before Einstein, Louis Bachelier made a similar observation when examining stock prices [24]. In his doctoral thesis, he described the fluctuations of stock prices as a stochastic process governed by a random walk. An illustration of the similarity can be found in Fig. 1.3. By assuming a diffusion process, he derived an options pricing formula given the fair price for a european call option, essentially creating the foundation for the field of mathematical finance. The Black–Scholes model, derived in the early 1970s, builds directly on this analysis [25]. Developed by Fischer Black, Myron Scholes and seperately by

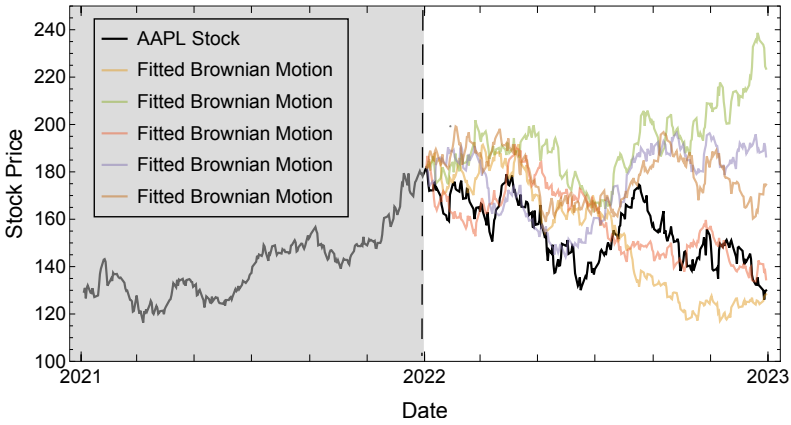


FIGURE 1.3: Stock price for Apple between 2021 and 2023. By generating a diffusion model with the stochastical features of the first year, one can generate possible stock movements for the second year. These are completely stochastic, but they can (but should not) be used for risk modelling.

Robert C. Merton [26], the Black–Scholes model is the go-to approach for option pricing today.

Conventional economic theory often assumes the general existence of an equilibrium. However in the 20th century, some scholars began to describe the market as a more complex system. In a chapter in *Coping with the Complexity of Economics*, Paul Ormerod examines how the contributions of both Keynes and Hayek are connected to complexity science [27].

Complexity science is an interdisciplinary field that studies complex systems and their behaviours [28]. It draws on insights and methods from various disciplines, such as mathematics, physics, biology, computer science, and social sciences, to understand the patterns and dynamics of complex systems. Among the systems studied in this context are biological ecosystems, the brain, ant colonies, the global economy, cities, and the internet. A core philosophy in complexity science is that a system is "more than the sum of its parts". To understand system dynamics, it is necessary to investigate the parts of a system and their relationship instead of just looking at an aggregate. Simple interactions within a system can lead to the emergence of much more complex behaviour on a larger scale. This is why

complex systems are often studied using a bottom-up approach, in which interaction in the microstructure of a system leads to an emergent order in the macroscopic structure of the system. Complex systems often operate on the edge of chaos, incorporating order and structure while still being susceptible to significant changes [29].

Conventional economic theories share some features of complexity science, like the concept of emergence. In general equilibrium theory, the macroscopic outcome of a market emerges from all of the individual preferences. Other similarities are non-unique equilibrium prices that can solve a system and non-linearity in variables [30], meaning that small changes can have disproportionate effects on the whole market. Nevertheless, the approaches of both Keynes and Hayek went beyond that.

Keynes understood the economy as usually in a transitory state between multiple possible equilibrium states, meaning that it is actually out of equilibrium most of the time, thus challenging the idea that economic systems always move towards a steady state. In addition to that, he described business expectations spreading as waves across the system, a concept that is often found in complexity science. In his proposed framework, firms can spontaneously change sentiment, thus introducing a statistical element into the system. Both Keynes and Hayek realise that firms are facing extreme uncertainty when trying to form expectations and inherent cognitive limits, thus introducing bounded rationality. Hayek used the differing expectations among different firms to explain the emergence of business cycles since incompatible expectations force firms to diverge from their plans [27].

Hayek, in general, was very critical of his profession, as is clearly visible in his 1974 Nobel lecture, which has the title *The Pretence of Knowledge* [31]. After claiming that "We have indeed at the moment little cause for pride: as a profession, we have made a mess of things", he continued to contrast the natural sciences, which allow for clear data collection and hypothesis testing, with economics, which, as he stated, was inherently plagued by the complexity of the studied systems.

Indeed, to this day, the interplay of countless heterogeneous actors makes economic and financial systems a very challenging problem to tackle. However, with recent technological evolution, there are much more powerful tools at our disposal in order to try and find solutions. This is, of course,

not limited to economics but applies to the study of for example biological and physical systems as well.

1.2 THE DIGITAL REVOLUTION

Ever since the invention of the transistor at Bell Labs in 1947, and the development of integrated circuits in 1958, technological capabilities in the processing of data have been increasing at an astonishing rate [32]. While first processor designs were relatively simple, with instructions executed serially, the invention of the integrated circuits allowed a transition from vacuum tubes to much more densely packed computing power. This greatly increased processing speeds, and with the invention of the microprocessor in the 1970s, even more elaborate computations became possible. From the development of the first supercomputers in the 1960s, the launch of the World Wide Web in 1991, to unlocking the holy grail of genetics research, sequencing the human genome in 2003, the increase in computational power has radically transformed technologies, economies, and societies starting in the latter 20th century, with further changes on the horizon with the rise of Quantum Computing [33] and Artificial Intelligence in the form of Large-Language Models [34]. With the digital revolution, both economics and physics have undergone dramatic change. In the natural sciences, especially physics, the technical transformation had dramatic implications both for how researchers perform experiments and theory. The increased computing capabilities allow for comprehensive simulations of systems, ranging from the simulation of electrons in a spin-lattice to fusion processes in stars to the behaviour of galaxies on an astronomical scale. In addition to that, experiments have seen a massive change in scale, with new sensor designs and better data management and analysis tools [35].

Some of the most ambitious projects in Physics, like the Large Hadron Collider (LHC), would not have been feasible without these improvements. Research at the LHC relies on digital sensors, which create an almost unfathomable amount of data (1 petabyte a second), which is processed and analysed with an incredible amount of processing power. Theoretical physicists then try to match the data to theories, which is done by comparing actual recorded data with simulation data. So, for a project like the LHC, these technological advancements really accompany almost every step of the way. However, the use of Big Data in research is, of course, not limited to the study of physical systems.

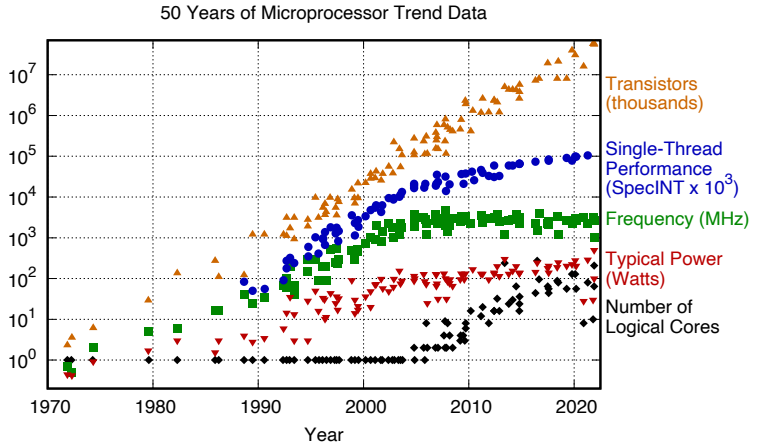


FIGURE 1.4: Various indicators of the exponential growth in computational power over the span of the last 50 years. Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten New plot and data collected for 2010-2021 by K. Rupp [32].

In *Economics in the Age of Big Data*, Einav and Levin investigate the effect that the digital revolution has had on the study of economic systems [36]. They note that within the last decades, there has been a shift into more empirical and data-driven research in economics. This is due to increased availability from both the private and the public sectors.

While private sector data contains consumer behaviour, transactions and communications, public sector data provides almost universal coverage. Especially in some Scandinavian countries, the vast data coverage of, for example, the labour market enables very detailed descriptions of economic microstructure [37]. Similar points are made in a working paper by Einav and Levin [38], where they outline opportunities for economic research, which include the possibility of new research designs and identification strategies, the use of causal inference as well as the ability to create highly customised policy tools [37]. However, these approaches are not without

potential fallacies. Athey [39] describes the issue of conflation of machine learning models with true causal inference, which requires additional assumptions and analysis. This can lead to issues for policy problems, which is why the proper use of data requires not just machine learning and data analysis but also the right analysis framework. Athey advocates for an interdisciplinary approach, in which machine learning should be combined with more traditional approaches. Increasing data availability brings many opportunities, but researchers should tread carefully in their analysis.

The newly available data sources in economics and finance require new approaches that often deviate from conventional approaches. These tools often make use of machine learning and data science like regression trees, lasso, cross-validation and others. These newly available data sources and computing capabilities have, at the fringe between economics and physics, led to a strand of researchers re-examining many of the assumptions and models of classical economics.

1.3 CHALLENGING CLASSICAL ECONOMICS

As discussed in Sec. 1.1, many of the assumptions of classical economics have been challenged in the last century. However, to this day, many policymakers still subscribe to classical economics, which has profound effects on the modelling of risk in financial markets.

Many risk models to this day rely on the efficient market hypothesis. Introduced in 1965, it relied on two observations: On the emperical side, Fama [40], after analysing stock prices, concluded that they follow a random walk. On the theoretical side, Samuelson provided an economic argument by describing market prices as a martingale process [41]. This means that the expected value of tomorrows price is similar to todays price and independent of all prior price movements. Fittingly, he titled his work *Proof that properly anticipated prices fluctuate randomly*, thus providing an economic argument for efficient markets [42]. In 1970, Fama described the market as "a market with great number of rational, profit-maximizers actively competing, with each trying to predict future market values of individual securities, and where current important information is almost freely available to all participants". There are three formulations of the efficient market hypothesis [43] :

- Weak form: The current price of an asset incorporates all of the historical financial information.
- Semi-strong form: The current price of an asset incorporates all of the historical financial information **as well as all other historical and current information. Prices change rapidly due to the arrival of new information** .
- Strong form: The current price of an asset incorporates **all** information about it, including historical financial information (weak form), all new public information (semi-strong form) and all **private information** regarding a financial asset [44].

All of the formulations have been subject to intense debate from both economists and physicists. While the strong and the semi-strong forms have been rejected by many, the weak form is still actively debated [45]. The weak form of the efficient market hypothesis can essentially be stated as the Random walk hypothesis, meaning that if the price includes all of the historical financial information, changes must follow a random walk. It would also render any financial analysis futile, as investors would not have any chance to "outprice" the market. This has been studied immensely (Sewell [42] gives a good overview).

The position taken in this thesis is that the evidence seems to suggest that the efficient market hypothesis should be rejected even in its weak form. This position is based on the following observations:

- The fact that stock returns on short time scales are fat-tailed;
- The assumption of rationality of actors in financial markets is dubious at best;
- Out-of-equilibrium phenomena like bubbles and crashes exist.
- The evidence that excess return can be made in markets, as visible by some funds and investors outperforming the market over the span of decades. This is a hotly debated topic, with many studies suggesting that it is impossible to outperform the market. However, others have shown how simple models would beat the market consistently over decades [46].
- Information asymmetries in markets seem to suggest that markets are indeed not efficient in distributing information.

Let's unpack them one by one:

1.3.1 *Fat-tailed returns*

Even before Fama first formulated his hypothesis, Mandelbrot investigated price distributions and found them to be fat-tailed [47]. However, other studies seem to suggest that stock market prices are indeed normally distributed, which is what we would expect with the efficient markets.

In 1999, Stanley *et al.*, when trying to make an argument for physicists meddling with Economics, reviewed some of this analysis. He concluded that "In summary, previous proposals for the histogram of index changes have included (i) a Gaussian distribution, (ii) a Levy distribution, and (iii) a truncated Levy distribution, where the tails become "approximately exponential" [48]. However, a key finding of this analysis was that indices can lose their fat-tailed nature because they average out a lot of the risk of individual stocks. A study of Dow Jones Returns reveals that also the time scale of measurement is important: Returns can be fat-tailed on short timescales but tend to be normally distributed over longer timescales [49].

1.3.2 *Rationality of agents*

In 1992, Tversky and Kahnemann presented cumulative prospect theory [50], which employs a cumulative functional of outcomes. They examine both value and weighting functions, which describe the transformation humans apply subconsciously to both the expected payoff and the outcome of specific options. The value function assigns subjective value to decision outcomes based on the preferences of the decider. Options with higher values are preferred to those with lower subjective values, so they can be used to compare different options. Risk preferences define the shape of these functions. While concave shapes indicate risk aversion, risk-seeking behaviour is described by a convex shape. The weighting functions transform probabilities into decision weights. They can be estimated from data by fitting curves to match probabilities to certainty equivalents. Weighting functions represent how people subjectively assess likelihood.

In a fully rational world, this behaviour would be accurately described by rational choice theory [51]. Both the value and the weighting function would exactly represent the probability-weighted payoffs of options. How-

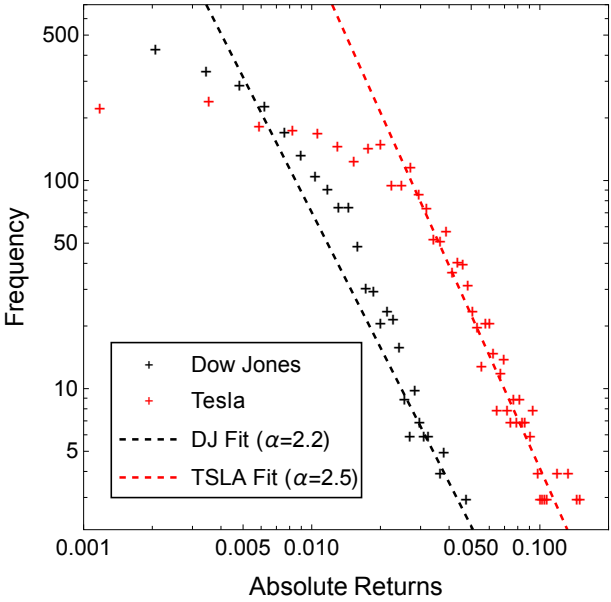


FIGURE 1.5: Histograms of absolute log returns for the Dow Jones index and for the Tesla stock. The Tesla stock has a larger tail exponent than the aggregated Dow Jones Index.

ever, experiments suggest a very different behaviour [52], thus deviating from rational choice theory. There are four main components explaining behaviour:

- Loss Aversion – Losses loom larger than gains of the same amount. The value function in prospect theory is steeper for losses than gains, representing how losses have a greater impact.
- Probability Weighting – People overweigh small probabilities and underweigh large probabilities. The weighting function in prospect theory captures this nonlinear transformation of probabilities into decision weights. See Fig. 1.6 for a visualization.
- Reference Dependence – People evaluate options based on change from a reference point like current wealth rather than comparing absolute values. Prospect theory reflects this via the utility function changing sign at the origin.

- Diminishing Sensitivity – The impact of gains and losses decreases with distance from the reference point. The prospect theory value function exhibits diminishing returns for both gains and losses.

Phenomena like loss aversion, nonlinear weighting of probabilities, reference dependence and diminishing sensitivity demonstrate that people use heuristic shortcuts and biases rather than purely rational calculus. The fact that people overweigh small probabilities and underweight high probabilities speaks to the limited ability of humans to accurately measure risk, instead relying on shortcuts and biases to make decisions. Factors like emotion can cloud judgment and lead to miscalculations.

The overvaluation of small probabilities has actual implications in markets, since there is evidence that the implied volatility of options tends to always be higher than the realized (so actually observed) volatility. This phenomenon is referred to as the volatility risk premium [53].

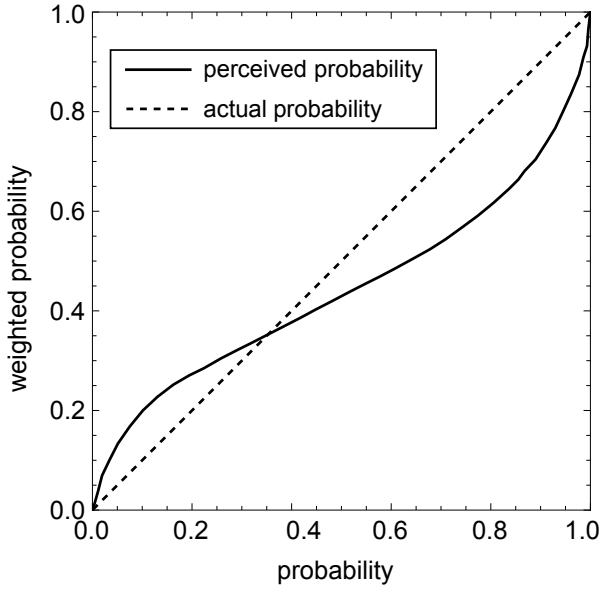


FIGURE 1.6: Probability weighting function adapted from [54]. Humans overestimate small probabilities and underestimate large probabilities.

1.3.3 *Excess returns*

In a world where the market incorporates all of the relevant financial data about an asset, no one should be able to consistently generate excess returns. Of course, the weak form of the efficient market hypothesis still allows for profits to be made through insider trading since privately held information is not necessarily incorporated. However, even without insider trading, some investors have been able to beat the market for decades by using patterns found in the market. Pesaran and Timmermann investigated stock returns from 1952-1992 and found that investors could have predicted stock prices based on factor models [46]. By creating a strategy of recursive forecasts, they show that a simple trading strategy would have outperformed the market. They observe that excess profits could mostly be made in high-volatility times like the 1970s. Similarly, Andersen and Sornette found "pockets of predictability" in markets [55].

As Colinsk laid out in 1996 [56], plenty of patterns can be exploited to gather excess returns. By utilizing these patterns and creating trading algorithms, a process often described as systematic trading, generating excess returns is indeed possible. The list Colinsk lays out is long and includes patterns like mean reversion, holiday effects, excess fluctuation in prices compared to fundamentals, the formation of bubbles, excess risk premia relative to bonds, deviation of mutual fund prices and their underlying securities, excess trading volume in shares that have risen in price, predictability from lagged insider trading, and more. These patterns are often called "stylized facts" of financial markets.

While most of them have been mentioned before, let's note the most important ones and their significance:

VOLATILITY CLUSTERING Volatility tends to cluster together – periods of high volatility in asset prices are followed by more periods of high volatility, while periods of low volatility are followed by more periods of low volatility. This can be seen in financial time series data (see Fig. 1.7), where volatility appears to come in clusters and waves rather than being constant over time.

FAT-TAILED PRICE RETURNS The distribution of asset price returns tends to have "fat tails" – there is a higher probability of extreme positive

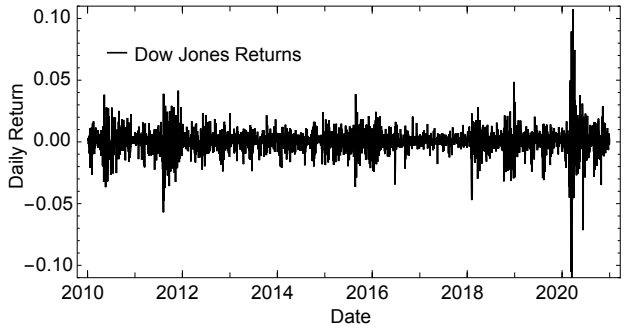


FIGURE 1.7: For the Dow Jones, large returns are clustered together

and negative returns compared to a normal distribution. This indicates a higher likelihood of extreme price movements in either direction.

AUTOCORRELATION OF RETURNS Asset price returns themselves tend to show little to no autocorrelation – the correlation between returns on consecutive days is close to zero. However, as depcited in Fig. 1.8, absolute and squared returns tend to display positive autocorrelation, indicating volatility clustering.

EXCESS VOLATILITY Asset prices demonstrate excess volatility – large price movements without apparent fundamental news or changes in dividends to justify the magnitude of the changes. Prices tend to be more volatile than would be expected based on fundamentals alone.

HIGH TRADING VOLUME Periods of high market volatility are associated with surges in trading volume as uncertainty leads to increased trading activity. High volume indicates significant disagreement among market participants about asset value.

TEMPORARY BUBBLES Asset prices occasionally detach from fundamental values, resulting in a bubble with prices exceeding true valuations. These bubbles inevitably burst as prices revert back down, such as the dot-com bubble in 2000.

TREND FOLLOWING Momentum and trend following are common market behaviours, where past price trends tend to continue for some period of time, resulting in trends like bull and bear markets. Traders attempt to

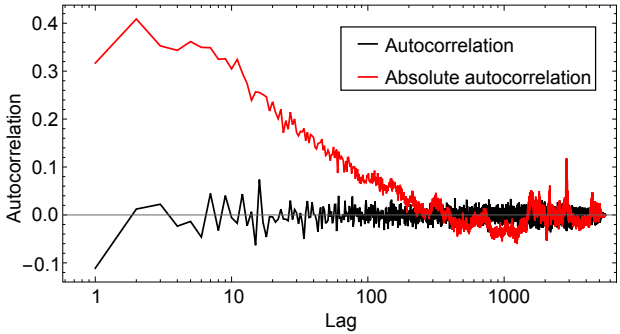


FIGURE 1.8: For the Dow Jones, autocorrelation is effectively zero. Only the first lag, which is negatively correlated, is significant. Absolute returns are correlated. This means large changes in market price tend to cluster together, but the directon of these changes does not.

identify and follow these trends. Figure 1.9 shows the Dow Jones index from 2018 to 2021. The moving averages denoted in red and green show long-term trends that are often utilized by traders betting on these trends.

SUDDEN CRASHES Markets occasionally witness sudden crashes or corrections where prices rapidly decline by large magnitudes in a short period of time, such as the 1987 Black Monday crash. These are exacerbated by leverage and panic selling. One of these crashes, which can be attributed to the Covid-19 pandemic in 2020 is shown in Fig. 1.9.

MEAN REVERSION While momentum can drive asset prices away from fair value in the short run, over longer periods, asset prices tend to revert back towards their mean or average level. Prices do not drift indefinitely away from fundamental values. The moving averages in Fig. 1.9 can be used by traders to bet on mean reversion: If the stock price is significantly over the 30 or 90-day average, chances are it will fall in the future.

INFORMATION ASYMMETRIES This point deserves its own section, which is why it is handled in Sec. 1.5.

The fact that the efficient market hypothesis has to be challenged has profound impacts on many facets of economics and finance. A prominent example of this is the Black-Scholes model [25], the go-to model for option pricing. The model is closely tied to the efficient market hypothesis since

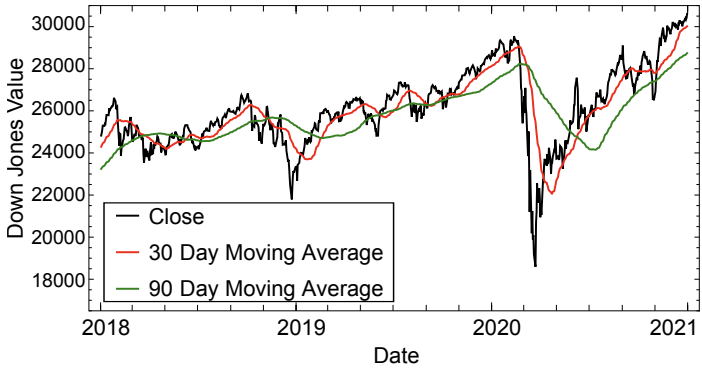


FIGURE 1.9: Daily close values and 30- and 90-day moving averages of the Dow Jones Industrial Index. From 2018 to 2020, the daily values fluctuated about the long-term averages, displaying mean reversion.

it assumes stock prices follow a random walk. Since returns are actually fat-tailed, a crucial part of this model follows erroneous assumptions, thus underestimating risk [57]. Out-of-equilibrium phenomena such as bubbles and subsequent crashes can lead to real problems when investors rely on classical risk profiles, as became evident in the 2007 financial crises [58]. This is why new tools are needed that can account for the stylized facts of financial markets and the corresponding changes in risk. One of these tools is agent-based modelling.

1.4 THE ADVENT OF AGENT-BASED MODELLING

Many of the models of classical economics rely on aggregate descriptions of human behaviour. This includes concepts like the economic equilibrium, demand and supply curves as well as the efficient market hypothesis. The main assumption behind this is that the macroeconomic behaviour of a system can be accurately modelled using aggregates and averages. Concepts like homogeneity, so the assumption that economic actors behave similarly given the same situation, and the use of representative agents allowed economists to derive comprehensive models of macroscopic economic systems from microeconomic assumptions. This is often referred to as a top-down approach. Top-down approaches usually start with an overarching theory, which is often based on the assumptions mentioned before, like market efficiency, rational expectations, and homogeneity. Following

deductive reasoning, the goal of such a model is then to provide prescriptive or normative models of how the economy should work. However, these assumptions can lead to issues when trying to model real markets and economies, as was outlined in the last section [59].

The question, however, remains how one can tackle these assumptions. The inherent problem is that deviations from simplifications like homogeneity and the averaging of microscopic behaviour result more often than not in analytically intractable systems. While dealing with these systems used to be unfeasible, the computational advancements of the last decades have allowed researchers to bring much more detail, microstructure and heterogeneity into models.

Among the solutions to tackle this problem is a method known as agent-based modelling. Contrary to macroscopic descriptions of systems, *agent-based models* (ABMs) model the individual entities within a system and study how their interactions then, in turn, lead to macroscopic dynamics. This is also known as a "bottom-up" approach. Deviating from established top-down approaches, bottom-up approaches start from individual interactions and the behaviour of agents. They allow for heterogeneity of behaviour and interactions. This is why they are handy for incorporating behavioural factors like bounded rationality, learning and social influence.

Agent-based models rely on the concept of "emergence" to model system behaviour. Emergence refers to the appearance of novel and coherent structures, patterns, or properties in complex systems on larger scales that are not inherent in or reduceable to the individual parts of the systems [60]. Interactions at the local, decentralised level can lead to self-organisation and the emergence of complex large-scale behaviour. Bottom-up approaches rely on the concept of emergence to describe the behaviour of complex systems by modelling individual entities and their relationships with relatively simple rules. An often-used example is the flocking behaviour in birds or schools of fish [61]. Simple local interactions between individual birds or fish, which are not subject to central control, can lead to fascinating patterns. For these models, the whole is truly greater than the sum of its parts. An example of this is shown in Fig. 1.10.

For economics and finance, agent-based models can implement bottom-up approaches by incorporating the behaviour of individual firms, traders,

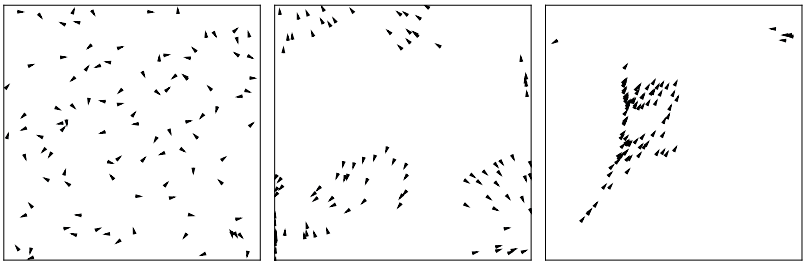


FIGURE 1.10: Boids model of flocking behaviour [62]. The three panels show a system of particles going from randomness (left) to swarm behaviour (right).

workers, citizens, etc. The concept first arose in the late 1950s when researchers began using computational models to describe both consumer and firm behaviour. First microsimulation attempts were made by Guy Orcutt, who modelled households and individuals with heterogeneity, with decision-making under constraints [63]. While still relying on neoclassical assumptions, the model shares many features with modern ABMs. At the same time, Herbert Simon built computation models of rule-following in firms [64].

In 1969-1971, Schelling presented his landmark model on residential segregation, showing that even mild individual preferences for same-group neighbours can lead to a high degree of segregation [65]. The model shows how two different groups of agents arrange their housing. Agents will move when the number of same-group agents in their neighbourhood falls below a "tolerance" threshold. Even though none of the agents has a high desire for segregation, the individual choices of the agents lead to highly segregated outcomes, as the results of Fig. 1.11 emphasize. This makes this model a great early example of how global patterns can emerge from local interactions without any macroscopic premeditation of outcome. At the same time, *cellular automata models*, inspired by Ulam and von Neumanns mathematical work became popular [66, 67]. These models feature populations of agents on lattices, with local interactions and were used to model urban dynamics, economic development and similar systems.

In the 1990s, Kirman illustrated yet again, how simple interactions can lead to complex structures with his "ant model" of recruitment dynamics [68]. The Santa Fe Institute Artificial Stock Market Model, one of the

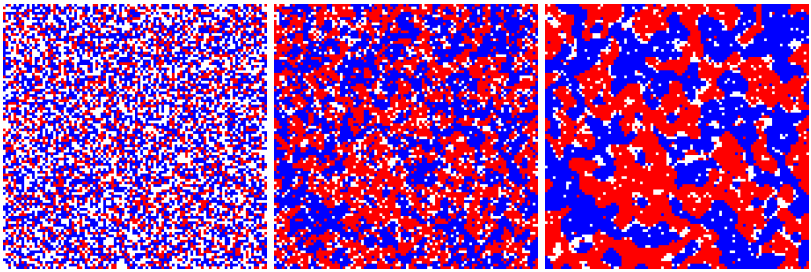


FIGURE 1.11: Emergence of spatial segregation between two groups (red and blue) of a Schelling model over time.

most influential economics ABMs, demonstrated that an artificial market does not necessarily converge to a rational expectations equilibrium [69]. On the contrary, the model produces bubbles. Thomas Lux and Michelle Marchesi introduced an agent-based model that could replicate many of the stylized facts of financial markets [70]. We will revisit this model in Chapter 2. Arthur's El Farol Bar Problem demonstrated how heterogenous expectations allow agents to arrive at a mixed strategy Nash equilibrium without explicitly computing it [71].

Models like the Santa Fe Stock market and the Lux-Marches model suggest that ABMs might be better suited to capture complex financial system behaviours. However, top-down models remain influential to this day [59]. Still, the nature of ABMs makes them particularly useful for describing non-equilibrium dynamics like bubbles and crashes.

The bottom-up nature allows ABMs to describe and investigate a plethora of market features – behaviour can be driven by the nature of the interaction as well as the structure of the system. The next sections will investigate the effects that both can have on agent behaviour.

1.5 STRUCTURE IN FINANCIAL MARKETS

Due to its ability to link microscopic and macroscopic behaviour, statistical mechanics has emerged as a viable method to describe financial markets [72]. Classical statistical descriptions of physical systems usually rely on placing interacting entities on lattices or assume homogenous links between agents. This approach is not too dissimilar from the aggregate descriptions found

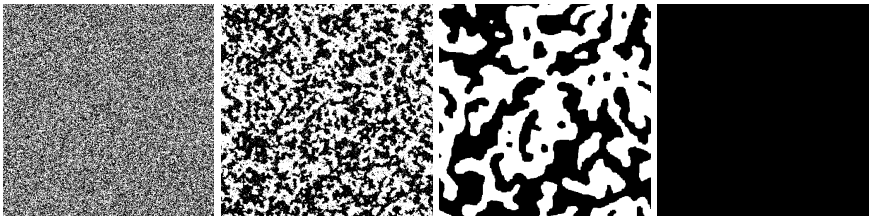


FIGURE 1.12: Ising model for high, medium, critical and zero temperature. When approaching the critical temperature, clear clusters emerge. At $T = 0$, all spins align.

in classical economics. These approaches are particularly attractive because they make systems analytically tractable and allow descriptions using mean-field approaches.

1.5.1 *Ising-like models*

A famous starting point for applying physical models to interacting humans is using the Ising model to describe opinion dynamics. The Ising model was originally used to describe a ferromagnetic spin system on a lattice. In the absence of an external field, it can be described with a Hamiltonian

$$H = -J \sum_{\langle i,j \rangle} s_i s_j \tag{1.2}$$

where the sum runs on the pairs of nearest-neighbours spins. It is convenient in the context of social dynamics because it describes the emergence of order in an initially disordered system. In the context of opinion dynamics, it has been used as a baseline model to describe the alignment of opinions within a system of agents. While the original Ising model works with spin-up $\uparrow$ and spin-down $\downarrow$, these two states can be interpreted to describe two opposing opinions. The dynamics of such a system can be seen in Fig. 1.12.

More elaborate agent-based models have utilized mean-field approaches by employing the assumption of homogeneous mixing. Mean-field approaches were originally introduced to solve analytically intractable physical systems. In Econophysics, the idea is that each agent is equally likely to interact with any other agent, washing away any topology or localization of a system. Agents essentially see the "average behaviour" of the other agents,

meaning they can be treated as if they are in a mean-field. An example of an agent-based model that successfully utilised this dynamic is the Lux-Marchesi model, which relies on a mean-field approach to describe traders in financial markets [70].

1.5.2 *Lux-Marchesi model*

In 1999, Thomas Lux and Michele Marchesi proposed an agent-based model featuring heterogeneous trader groups to introduce trader diversity into statistical models of financial markets [70]. Their model utilizes a mean-field approach and generates return distributions with fat tails and volatility clustering.

Description of the model

The model consists of two trader types – fundamentalists, who believe prices reflect fundamental values, and noise traders (or chartists), who extrapolate trends and herd. Fundamentalists base trading decisions on the market's fundamental or intrinsic value. They aim to profit by buying when the market price is below the fundamental value and selling when the market price exceeds it. In contrast, chartists disregard fundamental values and instead make inferences from past price patterns and trends. They comprise optimists, anticipating rising prices, and pessimists, expecting falls. Let n_f, n_+, and n_- denote the respective sub-populations of fundamentalists, optimists and pessimists, with total population $N = n_c + n_f$, where n_c is the sum of the number of optimists and pessimists.

Trader transitions

The model allows traders to change their behaviour, meaning they can move from one group to another. A key innovation was that transitions between groups are governed by driving forces U_i linking transitions to price patterns and majority opinions. This reflects adaptive behaviour observed among real traders [73]. The mean-field approximation factors into these dynamics in a couple of ways. Agents are assumed to randomly meet other agents, with whom they then interact. Chartists will meet other chartists with a probability $\frac{n_c}{N}$, while they will meet fundamentalists with a

probability of $\frac{n_f}{N}$. For interaction between chartists, they define the "opinion index", which is given as

$$x = \frac{n_+ - n_-}{n_+ + n_-}, x \in [-1, 1]. \tag{1.3}$$

The assumption is that chartists value all opinions equally and have the same chance of interacting with every other chartist, which is why they are driven by average opinion. x is a crucial determinant for the transitions between optimists and pessimists. The other key component governing transitions is sensitivity to price changes $\frac{dp}{dt}\frac{1}{p}$ of market price p. From this follows the driving force

$$U_1 = \alpha_1 x + \frac{\alpha_2}{v_1}\frac{dp}{dt}\frac{1}{p}. \tag{1.4}$$

Here, α_1 defines the weight placed on majority opinion, while α_2 defines the importance of the changes in the market price. v_1 represents the frequency of agents' strategies update, reflecting cognitive and information processing constraints. The driving force combines two effects: the majority opinion, denoted by the first term, and the price trend, denoted by the second term.

The other kind of interaction between chartists and fundamentalists is driven by a comparison between expected profits. When chartists meet fundamentalists with a probability of $\frac{n_f}{N}$, they then compare expected payoffs. Expected payoffs are given as $(r + dp/dt)/p - R$ for optimists, and $R - (r + dp/dt)/p$ for pessimists. Here, r describes the nominal dividends of the asset, and R denotes average real returns from other investments. Fundamentalist payoffs are given as $s|(p - p_f)/p|$, and they consist of the percentage deviation between p_f and p and a discount factor s.

This is why for the transitions between fundamentalists and chartists, there are two driving forces, $U_{2,1}$ for the transitions between optimists and fundamentalists and $U_{2,2}$ for the transitions between pessimists and fundamentalists, i.e.,

$$U_{2,1} = \alpha_3 \left\{ \frac{r + \frac{1}{v_2}\frac{dp}{dt}}{p} - R - s \cdot |\frac{p_f - p}{p}| \right\}, \tag{1.5}$$

$$U_{2,2} = \alpha_3 \left\{ -\frac{r + \frac{1}{v_2}\frac{dp}{dt}}{p} + R - s \cdot \left| \frac{p_f - p}{p} \right| \right\}, \tag{1.6}$$

where v_2 is a measure of frequency again, while α_3 is a measure of the importance of profit deviations. In addition to the market price p, this driving force also includes the fundamental price p_f. The probability of a trader switching into group b from group a in a time segment Δt is denoted by $\pi_{ba}\Delta t$. These rates are given by

$$\pi_{+-} = v_1 \frac{n_c}{N} \exp\left(U_1\right), \tag{1.7}$$

$$\pi_{-+} = v_1 \frac{n_c}{N} \exp\left(-U_1\right), \tag{1.8}$$

$$\pi_{+f} = v_2 \frac{n_+}{N} \exp\left(U_{2,1}\right), \tag{1.9}$$

$$\pi_{f+} = v_2 \frac{n_f}{N} \exp\left(-U_{2,1}\right), \tag{1.10}$$

$$\pi_{-f} = v_2 \frac{n_-}{N} \exp\left(U_{2,2}\right), \tag{1.11}$$

$$\pi_{f-} = v_2 \frac{n_f}{N} \exp\left(-U_{2,2}\right). \tag{1.12}$$

Note that the rates are proportinal to the probabilities of interacting with other traders, thus encapsulating the mean-field nature of the model.

Price changes

A key component of agent-based financial market models is the price adjustment mechanism in response to aggregate market forces. In the Lux-Marchesi framework, realized excess demand ΔE determines price pressure. ΔE has two components – chartist demand D_c and fundamentalist demand D_f:

$$\Delta E = D_c + D_f. \tag{1.13}$$

The chartist demand D_c depends on the imbalance between optimists and pessimists, representing non-fundamental momentum pressures:

$$D_c = \left(n_+ - n_-\right) t_c. \tag{1.14}$$

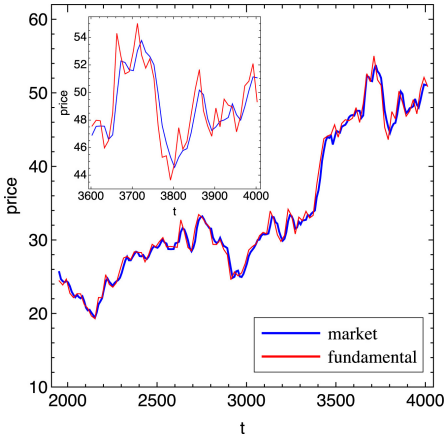


FIGURE 1.13: Market and fundamental price for the Lux-Marchesi model. The inset shows a higher time resolution excerpt of the market price following the fundamental price [74].

Here, t_c denotes the mean per capita chartist demand. The fundamentalist demand D_f arises from deviations between market prices p and perceived fundamental value p_f:

$$D_f = \gamma\, n_f \frac{p_f - p}{p}, \tag{1.15}$$

where γ scales the responsiveness to mispricing. Through this decomposition, the model aims to capture both traditional fundamental drivers and more behavioural phenomena like herding, trends, and investor sentiment. The interplay and coevolution between D_c and D_f allows complex dynamics like volatility clustering to emerge.

The probability for a price change of $\Delta p = \pm 0.001 p$ is then given by

$$\pi_{\uparrow p} = \max[0, \beta\,(ED + \mu)], \tag{1.16}$$

$$\pi_{\downarrow p} = -\min[\beta\,(ED + \mu)\,, 0]. \tag{1.17}$$

Price changes of the fundamental prices are described as a lognormal distribution.

$$p_{f,t} = p_{f,t-1} e^{\epsilon_t}, \tag{1.18}$$

with $\epsilon_t \sim N\,(0, \sigma_e)$.

Parameter	Description	Value
N	Number of agents	500
t	Mean trading volume	0.02
γ	Responsiveness to mispricing	0.01
β	Excess demand influence	6
ν_1	Frequency of strategy updates (chartism)	3
ν_2	Frequency of strategy updates (fundamentalism)	2
α_1	Weight of majority opinion	0.6
α_2	Weight of market price changes	0.2
α_3	Weight of market and fundamental value deviation	0.5
r	Nominal dividends	0.004
R	Real returns	0.75

TABLE 1.1: Summary of the parameters and example values used in the Lux and Marchesi model. Adapted from [75]

Please see Tab. 1.1 for an overview and example set of parameters. With the right set of parameters, the Lux-Marchesi model can reproduce some of the stylized facts observed in financial markets. However, the assumption of homogenous mixing fails to account for the rich microstructure of markets, which is essential to describing real-world phenomena [76]. For example, the hierarchical structure of financial markets has been investigated through numerous studies.

1.5.3 *Hierarchical Models*

In an attempt to describe crashes as critical points in markets, Johanson and Sornette proposed a buy-only model of traders that are organized in a hierarchy like the one in Fig. 1.14 as way of accounting for imitative and herding behaviour among traders. They held the view that market crashes reflect the hierarchical structure of stock markets and examined stock market crashes as critical points in the sense of statistical mechanics [77]. At this critical point, the macroscopic properties of the system dramatically change in a way which cannot be understood from the laws of the microscopic regime alone [78] because cooperative effects are of paramount importance.

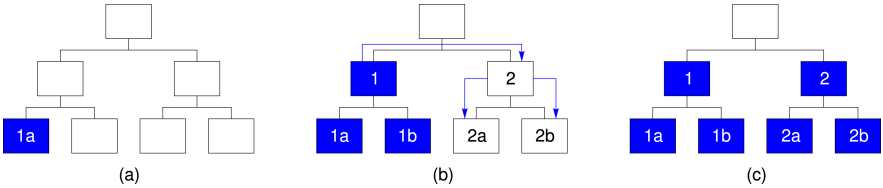


FIGURE 1.14: Illustration of how influence travels within the hierarchy. At time (a), only one trader in the hierarchy (1a) has bought. This influences the other trader in his group (1b) to buy sooner. At (b), (1b) has finally bought, resulting in the aggregate (1) to switch state as well. This influences (2) who will now buy sooner. Since (2) is an aggregate, (2a) and (2b) are influenced and will buy sooner. In (c), all individual traders switched states, leading the top trader to switch states [74].

An essential aspect of a hierarchical model is that the action of an individual trader influences only the traders in this trader's group and in lower levels of the hierarchy. Higher levels are affected by cascade effects. To produce log-periodic oscillations, this model reflects a hierarchical stock market. In this proposed model, individual traders are organized in groups, acting as effective traders on their hierarchy level. Thus, there is a single effective trader at the top of the hierarchy, which resembles the global economy. The hierarchy is preset for simplicity, even though the self-organization of hierarchies in complex systems can be observed under certain conditions.

The individual traders are called traders of level 0. These are organized in groups of m, which comprise the "traders" of level 1. These traders are also organized in groups and at the top of the hierarchy, there is a single effective "trader", the aggregate of all individual traders in their respective groups. All traders' actions primarily influence the "traders" on the same level and below, thus introducing an asymmetry in the flow of information (Fig. 1.14). Higher levels are influenced by cascade effects. This approach reflects a distribution of information that is incompatible with the efficient market hypothesis, as we have explained in Sec. 1.3. In the Sornette-Johansen hierarchical model, there are different levels of information, thus pre-empting the efficient market hypothesis. A trader of group m buys the stock, as soon as all the traders that are in the group of level $(m-1)$ of that trader have bought. The actions of a trader are determined by a utility function σ. For simplicity, all traders in this model have the same utility function. However,

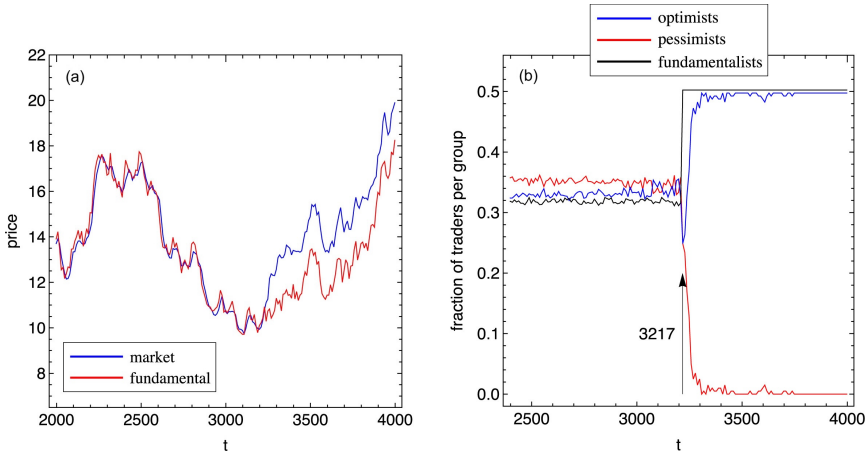


FIGURE 1.15: (a) The behaviour of the fundamental and market price in the modified model. (b) Moving averages of group occupancies for $b = 1.25$. At time step 3217, the model transitions into a bubble [74].

the traders differ in the time they need to make a decision.

The model assumes that the times at which traders decide to buy are Poisson distributed, with a mean of $(\kappa\sigma^{\rho})^{-1}$, where ρ is a parameter for the influence of the utility on the time-to-buy. The mean waiting time is dependent on the utility σ, thus making the cumulative distribution time dependent if σ is time dependent. The model also assumes imitative behaviour between agents. This is implemented through a shift in the utility when a nearby agent buys. However, this information is limited to agents in the same group as the buying agent. This informational asymmetry reflects a cost-effective minimal information strategy: traders of other groups might incorporate this information, but it would come with significant cost. That is why the model assumes this limitation of information.

During my bachelor thesis, I was able to show how combining the Lux-Marchesi Model with such a hierarchy can lead to the emergence of speculation with corresponding bubbles and crashes. A graphic from the corresponding publication is depicted in Fig. 1.15.

A follow-up investigation during my PhD looked at the effect of a fluctuating interaction strength. We let the hierarchy strength evolve according to

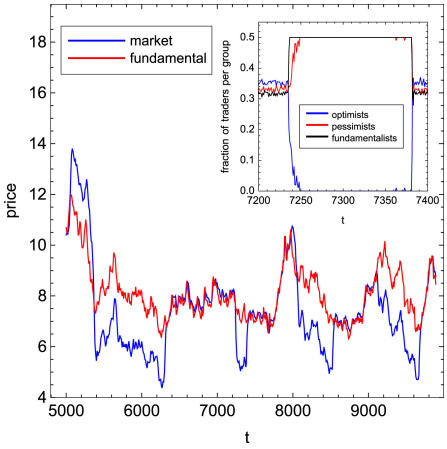


FIGURE 1.16: Price behaviour for a simulation with interaction strength fluctuating around $b = 1.25$. The inset shows the distribution of traders during the bubble around $t = 7300$. There, it is clearly visible that bubbles burst much faster than they emerge. [74]

an Ornstein–Uhlenbeck process. This is similar to dynamics proposed by Westphal and Sornette [79], where the coupling strength follows an Ornstein–Uhlenbeck process. This method is particularly useful for our system since it allows us to balance on the edge of the hierarchical regime, which resembles a market close, but not always at the critical point. In this regime, bubbles emerge and collapse again. Fig. 1.16 depicts the corresponding market movements.

While the market can now transition between the regimes, the most important finding can be taken from the inset. The depicted bubble emerges over a much larger time span than it collapses (which happens almost instantaneously). This is a pattern that could be observed throughout all the simulations in this regime. The fact that bubbles can collapse almost instantaneously strikingly reproduces a feature of the original hierarchical model, in which crashes can arise rapidly. While one might assume that the market diverges when b is above a certain threshold and converges with the fundamental value, otherwise, this is not the case. It was not possible to identify a clear pattern with regards to market price deviations and the corresponding hierarchy strength at a given time. This seems to suggest that, while small fluctuations are necessary, they are not the driver for the collapse of bubbles.

This work is a clear example of how system dynamics can be massively shaped by the **structure** of the interaction. Another prominent example of this is the use of networks in descriptions of financial markets.

1.5.4 *Networks in financial markets*

Networks have been used extensively to describe a variety of phenomena, ranging from ecological systems like food webs [80] to epidemiological networks [81] as well as economic applications like interbank networks [82], supply chain networks [83], and many others [84]. The origins of network science can be traced back to Leonhard Euler, who introduced graphs to solve the Königsberg Bridge Problem in 1736 [85]. Königsberg, a city in what was formerly Prussia (now Kaliningrad, Russia), was divided by the Pregel River into two parts, along with two islands, Kneiphof and Lomse. As illustrated in Fig. 1.17, both parts of the city were connected to the islands by three bridges each, and the islands themselves were linked by a single bridge, resulting in a total of seven bridges.

The question was whether it was possible to devise a path crossing all seven bridges without crossing the same bridge twice. Euler proved that this is not possible.

While this in itself might not sound particularly groundbreaking at first, the method he used for his proof would have far-reaching consequences and lay the groundwork for an entirely new strand of mathematics: graph theory and network science [84]. His revolutionary insight was to abstract the problem into a list of land masses and the connections between them, thereby converting the problem into a network. Networks are an abstract way to encode real-life structures and relationships and are commonly used today to describe real-world phenomena. By realizing that each of the landmasses can be represented as a node (or vertex) in a network while the bridges can be represented as links (also called edges), the problem becomes solvable purely by counting the number of edges each node is connected to.

Figure 1.18 illustrates how the seven bridges problem can be converted into a network. The key realization that led to solving the problem is that if there is a path over all of the bridges, all but two of the landmasses should have an even number of bridges connected to them. That is because each

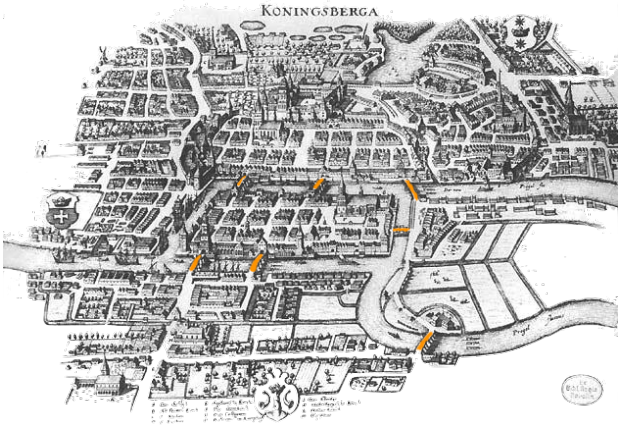


FIGURE 1.17: The 7 bridges of Koenigsberg, map by Merian-Erben (1652). The bridges have been marked in yellow for visibility.

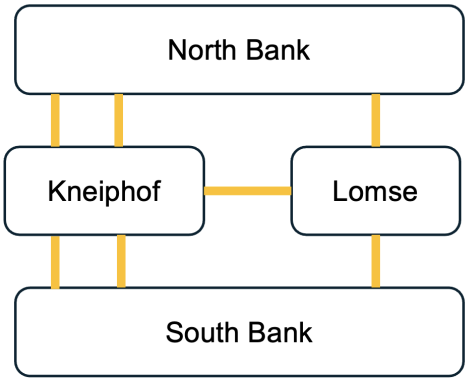


FIGURE 1.18: The seven bridges of Königsberg can be described as a network problem.

trip over a landmass requires a pair of bridges: One to enter the landmass and one to leave it. Since crossing the same bridge twice is forbidden, only the start- and endpoint can have an odd number of bridges.

Euler's initial insight sparked a plethora of further investigations into the properties of networks and graphs. In his 1923 book *Who Shall Survive?: A New Approach to the Problem of Human Interrelations*, Jacob Moreno introduced the idea of sociograms, a graphical way to describe interpersonal relations within a group. One of his key points was that the behaviour of a group could not simply be described by the sum of individuals but heavily relied on the relationship of the individuals within this group [86]. This once again confirms the importance of emergence when describing social interactions.

In a series of seminal papers published between 1959 and 1960, Paul Erdős and Alfréd Rényi laid the foundation for modern probabilistic graph theory. They introduced the concept of a random graph $G(n, p)$, which is a graph of n nodes, where the probability that any given nodes (n_1, n_2) are connected is given by the probability p [87, 88]. Because of their well-defined statistical properties, random graphs are still used today as null models, which can serve as a benchmark to evaluate empirical networks, helping to identify non-random structures [89]. They have also been used in situations where interactions are random, even though most models today assume that interactions follow more complex underlying rules.

The relevance of network science to describe human connections is best illustrated by a 1967 study by Stanley Milgram [90]. Milgram sent envelopes to 160 random people in Omaha, Nebraska, asking them to forward the envelope to the person they knew most likely to know – or be close to knowing a specified individual in Boston. Milgram tracked the packages and found that, on average, packages were sent along 5 "middlemen" until they reached his friend in Boston, prompting him to formulate the theory of "six degrees of separation", meaning that any two individuals are, on average, just 6 connections away from each other. Even though the study was criticized since not all packages actually arrived, subsequent studies, for example, by Microsoft [91] have shown similar results. Milgram's experiment excellently demonstrates why networks are such a powerful tool to study social systems: By encoding social structures into a network, we can come to fascinating insights, which would not appear obvious otherwise.

The networks Milgrim studied are referred to as the *small-world networks*. These networks are characterized by two key features: short average path lengths and high clustering coefficients. The average path length l_G of a network of n nodes is defined as

$$l_G = \frac{1}{n \cdot (n-1)} \cdot \sum_{i \neq j} d(i,j), \tag{1.19}$$

so as the average of the sum of shortest paths d of any nodes i, j in G. The shortest path between the nodes is the connection between two nodes requiring the least amount of edges.

The clustering coefficient N_i was introduced in 1998 by Duncan J. Watts and Steven H. Strogatz [92]. It is a measure describing how strongly nodes in a graph cluster together. Real-world networks, especially social networks, often exhibit tightly knit groups with a higher density of connections than would occur randomly. To determine the clustering coefficient of a node A, we look at all of its neighbours, so the nodes are directly connected to it. We then count the number of edges between the neighbours of A and divide it by the number of *possible edges*, so the number of edges that could possibly exist if all of the neighbours of A were connected to each other. An illustration of this can be found in Fig. 1.19.

$$N_i = \frac{\text{Number of edges between neighbours of node A}}{\text{Possible edges between neighbours of node A}}. \tag{1.20}$$

Watts and Strogatz also introduced the small-world network model, now known as the Watts-Strogatz model, in 1998. Their model demonstrated that networks can exhibit both high clustering, similar to regular lattices, and short average path lengths, like those found in random graphs, thus managing to describe the small-world properties observed in many real networks. An example is depicted in Fig. 1.20. The model is constructed as follows:

1. Start with a ring lattice of N nodes, where each node is connected to k nearest neighbors ($k/2$ on each side).
2. For each edge in the lattice, rewire it with a probability p. Rewiring involves connecting one end of the edge to a randomly chosen node, avoiding duplicate edges and self-loops.

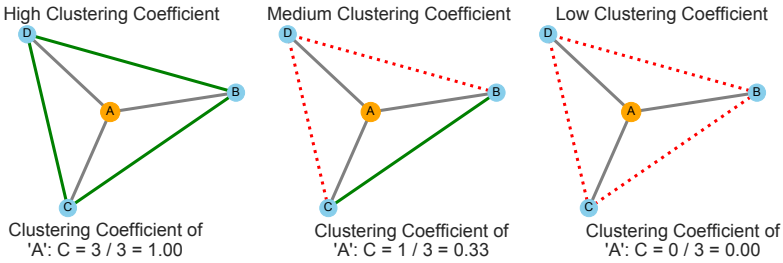


FIGURE 1.19: Clustering coefficient of node 'A' with varying neighbour connections. Green lines show existing edges, red dotted lines show missing edges. The coefficient decreases as connections are removed.

The parameters N, k, and p control the network's structure:

- N: Total number of nodes in the graph.
- k: Number of neighbours each node is initially connected to.
- p: Probability of rewiring an edge.

By adjusting p, the model transitions from a regular lattice ($p = 0$) to a random graph ($p = 1$), with intermediate values producing small-world properties. This makes the Watts-Strogatz model useful for studying real-world networks where nodes are locally clustered but globally connected.

Building on the work of Watts and Strogatz, Albert-László Barabási and Réka Albert introduced the concept of scale-free networks to describe the World Wide Web in 1999 [93]. They observed that many real-world networks exhibit a power-law degree distribution, meaning that a small number of nodes (hubs) have a large number of connections, while the majority of nodes have only a few connections (see Fig. 1.20). To explain the emergence of these scale-free networks, they proposed the Barabási-Albert (BA) model, which relies on the principles of growth and preferential attachment. The model is governed by a single parameter m, which determines the structure of the network. We start with a network $n \geq m$ nodes, then subsequently continue to add nodes. Each added node is connected to m randomly sampled already existing nodes, which results in the probability

$$p_i = \frac{k_i}{\sum_j k_j} \tag{1.21}$$

for the new node to be connected to the already existing node i.

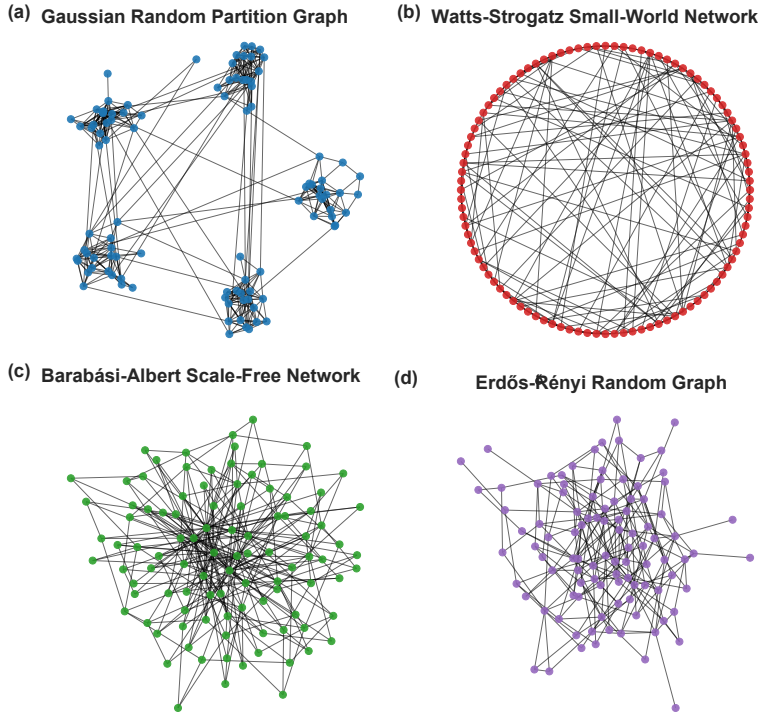


FIGURE 1.20: Visualization of four distinct network models: (a) Gaussian Random Partition Graph, (b) Watts-Strogatz Small-World Network, (c) Barabási-Albert Scale-Free Network, and (d) Erdős-Rényi Random Graph.

All of the described networks are depicted in Fig. 1.20, with tile (a) additionally depicting a Gaussian random partition graph. While not as widely used as the other networks, these networks can be used to describe community structures in social networks and are an excellent tool to visualize community behaviour in certain models. We will rely on them in Chap. 2 to illustrate model behaviour.

1.5.5 *Network Applications in Economics and Finance*

Network analysis has grown to be an important factor in financial and economic analysis. Due to their availability to encode complex relationships into a mathematical framework that allows comprehensive analysis, their importance continues to grow with the increasing amount of data and processing power at researchers disposal. One of the most sought after use of networks is in the mitigation of financial risk in the banking sector.

Especially since the 2007-2008 financial crises, both researchers and policymakers have understood the need for better approaches to model financial risk for financial institutions. Banks are highly interconnected and often highly leveraged, leaving them vulnerable to sudden liquidity shocks which can spread quickly. By describing the liabilities between banks as edges in a network, interbank networks can be used to model what is called "financial contagion", meaning the quick spread of financial distress amongst banks and other institutions [94–96]. Boss *et al.* studied the Austrian Interbank network and found that interbank networks exhibit properties common to many real-world systems, such as power-law degree distributions and small-world characteristics [97]. Acemoglu *et al.* studied the theoretical properties of such financial networks and identified a phase transition in network resilience: while dense networks are stable under small shocks, they become fragile under large shocks [98].

Systemic risk in banking networks can emerge through multiple channels: Interconnected liabilities can propagate shocks when one institution defaults, affecting others through cascading failures [94, 96], Liquidity shortages or solvency issues in one part of the system can escalate due to interdependencies [95, 99] and shared exposures to common assets can amplify risks [100]. By understanding these risks, policymakers can develop better regulations to limit the risk exposure of financial institutions.

In addition to looking at the balance sheet data of institutions, networks can also be created from observed correlations of stock prices. This has the advantage that data is often widely available. However, due to the amount of noise in stock markets, sophisticated methods are needed to draw empirically founded conclusions. Onnela *et al.* employ a Minimum Spanning Tree framework to study stock market correlations dynamically. By constructing "asset trees", they aim to gain insights from vast correlation

matrices. Using this method, they manage to find evidence for hierarchal structures in markets [101].

Tumminello *et al.* investigate clusters of investors in the Nokia stock market by utilizing Statistically Validated Networks. They find behavioural synchronization among investors using trading data, which confirms some of the stylized facts of financial markets, like herding behaviour and heterogenous responses to traders [102].

Particularly interesting insight can be drawn when combining both approaches: By looking at fundamental connections between different firms, we can try to understand market behaviour. One of the best sources of data for this is Input-Output tables, which allow us to create *production networks*. These networks describe the interplay in production among different entities by linking firms in multiple states of production processes. Similar to financial risks, shocks to certain production capabilities or demands can have widespread application across the whole sector [103]. Consider, for instance, the car manufacturing industry: A decline in car sales will affect not only car manufacturers but also countless other smaller firms manufacturing specialised parts necessary for vehicle production. Acemoglu *et al.* describe how such shocks can lead to increased tail risks (so larger than-expected price changes). They find that large downturns are disproportionately shaped by network heterogeneity and sectoral dominance [83].

Herskovic finds that network measures can predict the financial risk for stocks. He introduces two network-based factors, network concentration and network sparsity, as determinants of systematic risk. Both factors can significantly impact asset returns [104]. Marti *et al.* provide a comprehensive review of the state-of-the-art methods used to create and analyse financial networks [105].

While the mentioned studies only cover a small part of academic research into financial networks, they still show the clear link that can be drawn between the network structure of both traders and firms and market behaviours, like the stylized facts we described previously. After looking at the structure of the system, it is time to examine the other crucial part of agent-based modelling dynamics: The behaviour of agents.

1.6 BEHAVIOURAL SCIENCE AND ECONOMICS

Arguably, even more important than the structure of interaction might be the nature of the agents behaviour. How do they choose between different options? How do they incorporate the actions of others into our behaviour? How do they deal with risk?

Economic models have depended on assumptions about human behaviour from the very beginning of the discipline. The core concept of economics, with which every first-year student is confronted, is the "homo economicus", the all-rational being that always picks the option with the most utility. However, beginning in the 20th century, a new subclass of economics emerged to challenge the old assumptions of the field: behavioural economics.

Behavioral economics incorporates insights from psychology, neuroscience and other sciences into economic analysis. It relaxes assumptions of perfect rationality and recognizes that human decision-making is subject to various cognitive and emotional biases. These studies often rely on experiments, surveys, field studies and tools from psychology and neuroscience like brain scans. By utilizing these approaches, the hope is to provide a more psychologically realistic foundation for economic analysis [106].

The themes explored in behavioural economics are wide-ranging: From looking at the motivations of behaviour, which often deviates from financial incentives to self-image or emotions, to social influences. This a theme also prominent in agent-based modelling, since social influence may shape behaviour through social norms, pressures or comparisons. Many insights can be gained when taking a closer look at the behavioural aspect of decision making. Humans make systematic mistakes and display biases, often due to limits in time, attention, memory or processing capability. They often misjudge probabilities, focus too much on the short term, and make decisions based on personality and emotion.

One of the key achievements of behavioural economics, prospect theory, was already discussed in Sec. 1.3.2. It explains the risk-averse nature of most human actions with two key insights: We evaluate options relative to reference points and not in absolute terms, and we feel the pain of losses much more strongly than the pleasure of gains. That is why we choose a smaller guaranteed reward over a riskier option with the higher expected

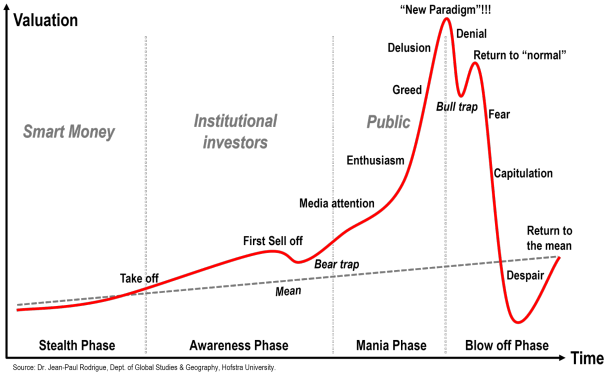


FIGURE 1.21: The four stages of a bubble according to Rodrigue: Stealth, Awareness, Mania, and Blow off. This description identifies herding as the main driver behind bubbles and crashes [108].

value. We are, in some sense, "over-risk-averse". In a similar sense, we overvalue present goods, even when faced with much higher payoffs in the future. This has been described as a struggle between our "patient self" and "impatient self" [107]. This has profound implications for modelling and policy since most models assume that agents will happily forgo consumption in the current period for greater consumption in the following period.

A great feature of behavioural economics is that it allows us to find explanations for some of the observed stylized facts in financial markets. It does so by taking a close look at the drivers of our behaviour and how they deviate from rational choices. A good example of this are herding phenomena. While many models accept the existence of herding in financial markets, many simply assume these to be due to some natural alignment between interacting agents (just remember the Ising-like opinion models discussed in Sec. 1.5.1). Behavioural economics allows a more intricate view of this phenomenon.

There are two main drivers of herding behaviour: Normative and informational influences [106]. The informational influence is quite in line with what we would expect from rational agents: We follow the choices of others because we assume that they have information that might not be available to us, so we hope to benefit from the "wisdom of the crowds". This can, of course, go horribly wrong, however, there are good arguments to

be made why such a strategy could be profit maximizing in the absence of better alternatives. The normative influence however refers to an entirely different driver of our behaviour. Instead of referring to some sort of profit-maximizing or information-gathering motivation, it speaks to the human desire to conform with the crowd. Reasons for that can be social acceptance or reassurance in picking the same options. This behaviour is irrational but has been observed in experiments [109]. There is evidence that we are hardwired on a neural level to do this, as studies on mirror neurons suggest [110].

There are many other insights that behavioural economics was able to gather through social experiments, but we will focus on a particular aspect of the discipline, the attempt to identify the neuronal foundations for economic decision-making, also known as Neuroeconomics [111]. A typical example of deviation from profit-maximising behaviour is the tendency of humans to decline offers that are perceived as unfair. If I have to split 100 Dollars with you and offer you just one Dollar, chances are that most people will decline the offer because they perceive the split as unfair. This adversity towards unfair deals has very real implications for economic systems. Using fMRI imaging, researchers have been able to investigate brain responses to fair and unfair offers, thereby giving tangible explanations as to how humans react to unfair offers on a neuronal level. As it turns out, our brain reacts to unfair offers similar to observing someone breaking social norms [112].

The analogy used earlier between the patient and impatient self when valuing future payoffs can be put onto a solid neuronal footing when looking at experimental data [107]. McClure *et al.* found evidence that we use separate neural systems that drive impulsive versus patient behaviours. They relate it to a "multiple selves" model – the impatient, present-focused self activates one set of neural circuits, while our more rational, future-oriented self activates another. Whichever system is more active drives our final decision. The final nails in the coffin for complete rationality are studies performed by Knutson *et al,* [113], which show that we respond emotionally to safe versus risky investments.

While many researchers have since accepted concepts like bounded rationality and herding and most agent-based models operate under these assumptions, there still remains a gap that we aim to bridge with this thesis.

Instead of focusing on some of the findings of behavioural economics and trying to replicate these through a model, the work presented here starts at the very fundamentals of every economic interaction, the neuronal dynamics governing the decision-making process.

1.7 A NEW PARADIGM: NEUROLOGICAL BASIS FOR DECISION MAKING

Economists have tried to find the underlying principles driving decisions ever since von Neumann and Morgenstern introduced the classic expected utility theory in 1944 [114]. According to this theory, people choose the option maximising their expected utility. This became a centrepiece of the neoclassical economic revolution for modelling human behaviour. However, empirical evidence started to show deviations from the predictions of the expected utility theory.

The Allais and Ellsberg paradox experiments in the 1950s-60s demonstrated choice patterns inconsistent with the axioms of expected utility theory [115, 116]. When faced with a choice between different urns, participants irrationaly favored some urns over others with equal payoff. This led to the development of alternative theories like the previously mentioned prospect theory. However, these approaches are still limited in their ability to describe the processes because they view the system from the outside perspective, essentially as a black box. This is why, when advances in neuroscience allowed more detailed insights into neuronal processes, researchers began to relate neural activity and choice behaviour [117].

Today, the consensus among researchers is that there are two underlying principles of human decision processing: It is a gradual process, meaning that we accumulate information over time and it is governed by stochastic processes because the information absorbed at each time is subject to small fluctuations [118].

In their 1993 paper introducing Decision Field Theory [119], Busemeyer and Townsend identify multiple stages that try to tackle the neuronal basis of this problem with increasing detail and realism. Starting with Deterministic Subjective Expected Utility (SEU), they define the first stage as a simple binary decision model in which the agent simply chooses the option

with the higher expected payoff. This initial stage again reproduces the behaviour of a rational economic agent. This fundamental approach does not account for the variability in preferences, which is why the second stage, Random SEU Theory, allows subjective probabilities to vary across trials. Further extensions include introducing time-varying preferences within a trial, a description of preferences as a random walk and the introduction of growth and decay rates as well as changing gradients. Their final model generalized all of these extensions into a single approach.

Building on this work, Usher and McClelland introduced the leaky computing accumulation model in 2001 [118]. They emphasise two important features of decision models: The leaky nature of the information accumulation process and the fact that different options compete with each other for attention, which they interpret as an inhibitory phenomenon. They describe the accumulation process with a simple two-layer network in which a layer of inputs is connected to a layer of "accumulator units". Accumulator units are described as being populations of neurons. Their state is defined by two key entities: their activation u and their output t. (The symbols chosen in this thesis to describe these entities might seem arbitrary, but will make sense by the end of this chapter). The output is governed by an activation function f, which is a threshold function

$$f(u) = \begin{cases} 0 & \text{for } u < \vartheta, \\ u & \text{for } u \geq \vartheta \end{cases} \tag{1.22}$$

with threshold ϑ. When the threshold is reached, the unit fires. Depending on the context, either the unit with the highest firing rate or the first to reach the threshold triggers the response. They then draw on the cascade equation, first introduced by McClelland in 1979 [120], to find a mathematical description of the corresponding behaviour. The cascade equation describes the evolution of the activation and is given as

$$du_i = [I_i - \lambda u_i]\frac{dt}{\tau} + \xi_i\sqrt{\frac{dt}{\tau}}. \tag{1.23}$$

Here, I_i describes the input from other units, λ is the decay rate, τ is the timescale and ξ_i is a Gaussian noise term with $\langle \xi_i \rangle = 0$ and $\langle \xi_i \xi_j \rangle = \sigma^2 \delta_{ij}$. Equation (1.23) describes a leaky, stochastic accumulator. While the 1979 paper then investigated cascade effects by describing a system comprised of multiple of these layers, Usher and McClelland then adapted this equation to the problem of competing choices. They do this by introducing an

inhibition term, which corresponds to the fact that high activation for one option reduces the activation for another option.

The final equation they arrive at is

$$\begin{aligned} du_i &= [I_i - \lambda u_i - \beta \sum_{j \neq i} u_j] \frac{dt}{\tau} + \xi_i \sqrt{\frac{dt}{\tau}} \\ u_i &\to \max(u_i, 0), \end{aligned} \tag{1.24}$$

where β governs the inhibition term. With this model, they show in a series of experiments that they can reproduce experimental data, making the model a valuable step towards a description of decision processes. The model can describe multiple options competing for attention and has a threshold mechanism for selection. It is a powerful framework to describe decision-making.

Following a genuine bottom-up approach, the next step is to find a way to incorporate the neuronal basics for information processing and decision-making into a model that is capable of not just handling a single decision-making process but can be used to describe the decisions of many individuals faced with the repeated necessity to make choices. The goal is to find an agent-based model that can utilize not just advanced decision mechanics but is also capable of representing the importance of structures like networks and hierarchies in financial markets. To achieve this, we take our inspiration from the neural basis for decision-making and investigate a model that exhibits many of the same dynamics but that is capable of describing the interactions of many neurons and has been used to model brain activity ever since Lapique conceived it in 1907 [121].

1.8 INTEGRATE-AND FIRE MODELS

Given that every human decision results from interacting neurons, it seems intuitive to use a model for brain activity to model decisions. However, that is easier said than done. The computational power to model even a single human brain is (at least at the time of writing this) nowhere to be found, so modelling the many millions or even billions of people that interact within the stock market seems beyond infeasible. However, as discussed in Sec. 1.6, decision-making processes can be abstracted so far that we just model certain populations of neurons reacting to incoming information. Still, there are a lot of adjustments and abstractions necessary to make the model work.

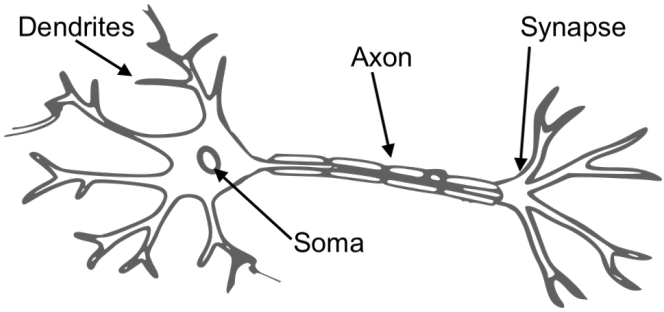


FIGURE 1.22: Sketch of a neuron. Adapted from [124].

The function of a neuron can be best understood by looking at the three parts responsible for its actions: a dendrite, a soma, and an axon [122]. The dendrites are the input device, the soma process inputs, and axons are responsible for transmitting outputs to other neurons. The connections between neurons are called synapses. The actual dynamics of a neuron are very complicated, so there are many models which try to simplify the behaviour. One of these is the integrate-and-fire model [123]. Since this is one of the simplest models and, as such, can be used with great computational efficiency, it is a suitable starting point for applying neuronal dynamics to traders. The integrate-and-fire model represents the following aspects of neurons: They integrate incoming currents until a certain critical voltage is reached, at which time they fire. The behaviour of a neuron is defined by its action potential u, the incoming current I, its resting potential u_r, and its threshold ϑ. The current I at time t can be expressed as the sum of the resistance current I_R and a current I_C:

$$I(t) = I_R + I_C. \tag{1.25}$$

The current I charges the neuron, which is described as a capacitor with capacity C and resistance R. Ohms law for I_R and $I_C = dq/dt = C\,du/dt$ yields

$$I(t) = \frac{u(t) - u_\mathrm{r}}{R} + C\,\frac{du}{dt}. \tag{1.26}$$

With the decay constant $\tau = R\,C$, we see that the action potential follows the differential equation

$$\tau \frac{du}{dt} = -\big[u(t) - u_\mathrm{r}\big] + R\,I(t). \tag{1.27}$$

This linear differential equation describes the behaviour of a neuron like a leaky integrator. The only missing part is the the reset condition. This introduces nonlinearity in the system, which is one of the reasons why it is attractive for financial modelling. To achieve this, we introduce the firing time t_f. At the firing time, a neuron transfers its integrated current back into the network, which is called the "fire event". The firing time is defined as the time at which the action potential reaches the threshold

$$u(t_f) = \vartheta. \tag{1.28}$$

After firing, the neuron returns to a resting state, which can be expressed as

$$\lim_{\delta \to 0; \delta > 0} u(t_f + \delta) = u_r \,. \tag{1.29}$$

This behaviour is clearly visible in Fig. 1.23.

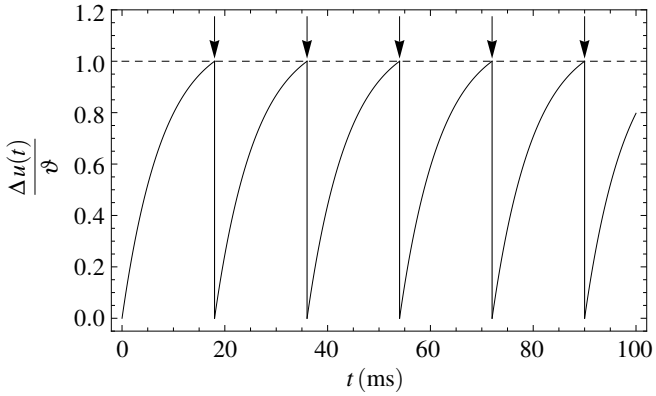


FIGURE 1.23: Change in action potential for a neuron subjected to a constant current. Δu is normalised by the threshold ϑ. The action potential rises according to (1.27) until it reaches the threshold (indicated by the broken line), where it is reset (indicated by arrows). Adapted from [123].

The integrate and fire model is one of the simplest descriptions of neuronal behaviour available. It draws its inspiration directly from a common description of neurons as capacitors. It is also clearly a model of information accumulation, just the level of observation is different since it describes individual neurons, not populations.

1.9 OUTLINE OF THE THESIS

This thesis is structured into seven chapters, each building on the previous to progressively develop, extend, and apply the integrate-and-fire framework.

- In this first Chap. 1, we discussed the joint history of economics and physics as well as some of the common critiques of classical economic theory. We introduced approaches from behavioural economics and the integrate-and-fire model for neurons. This lays the groundwork and explains the intution for applying the integrate-and fire model to financial markets.
- In Chap. 2, we will adapt the integrate-and fire model so that it can be used as part of agent-based models, recovering many of the features of the leaky competing accumulator model described in Sec. 1.7. In the new model, agents integrate information that they receive from observing price changes and from other agents and "fire" when they have enough information to make a decision. Afterwards, the effect of parameters and network structures is investigated. We show how the model can replicate wide ranges of behaviour, including volatility clustering, spontaneous opinion formation, and critical transitions.
- In Chap. 3, we introduce inhibitory effects into the model. These represent psychological biases and irrational behaviour. We discover that inhibition can lead to speculative bubbles and their eventual collapse. The chapter further explores entropy and phase transitions, using tools from statistical mechanics to interpret market dynamics. By introducing inhibition into the model, we are able to establish a direct link between irrationality in the neuronal level of decision-making and markets deviating from efficient pricing.
- Chapter 4 introduces hierarchical structures into the agent network. These allow us to simulate how information propagates through layered social systems. The results reveal how information siloes can lead to consistent mispricing and rapid market transitions.
- In Chap. 5, we generalise the integrate-and-fire model to a framework for social interaction. This new framework introduced constraints into the model, a crucial feature for real-world applicability. We use this framework to extend the trader model to a multi-asset order-book model. In this model, market makers, volatility clustering, realistic

return and wealth distributions which align with real-world data naturally emerge.

- In Chap. 6, we demonstrate the wide applicability of the integrate-and-fire framework to by applying it to a different discipline: epidemiology. We model the spread of disease using an integrate-and-fire approach and show that the model can effectively capture many patterns found in infection dynamics. This includes effects like secondary peaks and modelling interventions. We also develop a method to calibrate the model to data and manage to do so with high accuracy.
- Chap. 7 concludes the thesis with a summary of findings, theoretical implications, and directions for future work, including calibration strategies, additional applications like banking risk models, as well as potential improvements to behavioural and structural realism.

2

FIRST MODEL

The mission is clear: we want to find a way to incorporate the neuronal level of decision-making that drives economic behaviour in an agent-based model. However, on the path to achieving this, a couple of crucial challenges need to be overcome. The first – and the most pressing – is the question of how to reconcile classical drivers of decisions in economic models with this new neuronal foundation. Even though the model is supposed to provide a more accurate description of the decision-making process, it must still be anchored in economic incentives. We also must not fall into the trap of making the model overly complicated. The introduction of a neural decision-making "engine" bears the risk of introducing arbitrarily many parameters, which would render the model useless because it would be prone to overfitting. We need to make sure to have an appropriate level of abstraction so that we really capture the fundamental behaviour of the system, but do not fit to noise. Finally, we need to appreciate the massive leap in hierarchy that this model is trying to assume. The global economy is made up of billions of interacting entities, the behaviour of each of which is the result of a hundred billion of interacting neurons [125]. In complexity science, similar behaviour is often found on many scales, and this is what we are counting on here. The first steps to doing this were done during my master's thesis. However, a much more elaborate version with a more solid footing in literature is presented here.

2.1 ADAPTING MODELS TO THE INTEGRATE-AND-FIRE APPROACH

To begin the first attempt at constructing a model that can make use of a neuron-like mechanic, we take an already established agent-based model and adapt it to this new dynamic. The reason for this is apparent: It tackles the most significant issues that we might face. By choosing a relatively simple but powerful model, we can avoid making our approach overtly complex while still profiting from the economic intuition that underpinned the initial model. We rely on the Lux-Marchesi model, already established in Sec. 1.5.2 as a starting model. For completeness, we will concisely sum-

marize the parts of the model relevant to the new approach again.

The Lux-Marchesi model operates under the assumption that a trader can be in one of three states (or opinions): a fundamentalist, an optimist or a pessimist. Fundamentalists are traders who believe that the market price will return to an underlying fundamental value p_f in the long run. The expected future earnings of the asset determine this value. Their strategy is to buy when the market price is under the fundamental value and to sell when it is above the fundamental value. On the other hand, optimists and pessimists are chartists who do not consider the fundamental value. They form their opinion based on movements of the market price. Optimists expect an increase in market price, while pessimists expect a decrease. To summarise, we are faced with a system of traders that can follow three different opinions that are influenced by the market and fundamental price and by social interactions.

The first important realization is that while these two influences are very different (one is based on the individual assessment of market movements, while the other is a social phenomenon), both can be reduced to simple sources of information that help drive the agent's decisions. During the decision making process, we collect all of the information that is relevant to the issue and then, when our information threshold is met, we decide. This exact dynamic, which also underpins the information aggregation models introduced in Sec. 1.7, is why we have selected the integrate and fire model as a base model. It describes the behaviour of neurons in the exact same way: They integrate currents, which represent information, up until a threshold is met and they fire.

The leaky stochastic accumulator [120] describes the decision-making process by a stochastic differential equation, given as

$$\tau \frac{du_i}{dt_i} = [I_i - \lambda u_i] + \xi_i \sqrt{\frac{\tau}{dt_i}}. \tag{2.1}$$

This equation is strikingly similar to that of the Integrate-and-Fire model

$$\tau \frac{du}{dt} = -\big[u(t) - u_{\mathrm{r}}\big] + R\, I(t). \tag{2.2}$$

Please refer to Secs. 1.7 and 1.8 for definitions of all the model parameters. The most important difference betweem the models is, of course, that the

neuron only works in one dimension since it works with electrical current I. The other important difference is the reset condition that drives neuron behaviour, resetting the action potential u after the firing event.

$$\lim_{\delta \to 0^+} u(t_f + \delta) = u_r \,. \tag{2.3}$$

This is critical for a dynamic model of decision-making since market actors are faced with decisions repeatedly. This is why the "integrate-and-fire" dynamic is so powerful: It resembles the way in which we accumulate information and decide and how, after our decision, we essentially "reset" until new information drives us to a new decision. Still, in its one-dimensional form, the integrate-and-fire model is hardly helpful in describing the decisions that guide our actions in financial markets. It has been previously used to model social network dynamics [126], however, never for people facing choices between different alternatives. The Lux-Marchesi model offers three alternatives of opinion: optimism, pessimism or "fundamentalism".

That is why we need to generalize the electrical currents driving the integrate-and-fire model in a multidimensional current carrying information flows. To account for the three possible opinions, a three-dimensional current is introduced, where each component of the current corresponds to one possible opinion, fundamentalist (I_f), optimist (I_+), and pessimist (I_-):

$$\mathbf{I} = \begin{pmatrix} I_f \\ I_+ \\ I_- \end{pmatrix} . \tag{2.4}$$

This current allows the traders to communicate information about each opinion over a separate channel. This change requires, in turn, a modification of all the other quantities in the model. The action potential is now the three-dimensional vector $\mathbf{u}$,

$$\mathbf{u} = \begin{pmatrix} u_f \\ u_+ \\ u_- \end{pmatrix} , \tag{2.5}$$

and the resting potential is now given as

$$\mathbf{u}_r = \begin{pmatrix} u_{r,f} \\ u_{r,+} \\ u_{r,-} \end{pmatrix} . \tag{2.6}$$

Again, the indices $f, +, -$ indicate which opinion a dimension corresponds to. In a similar fashion, the decay constant τ and the resistance R are now expressed as vectors. This allows the reformulation of (2.2) as a component-wise differential equation:

$$\tau_i \frac{du_i}{dt} = -\left[u_i(t) - u_{i,r}\right] + R_i \, I_i(t), \tag{2.7}$$

where $i = 1, 2, 3$.

We have now arrived at a set of differential equations that describe how information regarding different opinions is aggregated, or as we will say from now on integrated, over time. This allows us to describe how traders would collect information regarding all of their possible opinions. In order to have a working model, however, the aggregation of information needs to result in actions.

2.1.1 *The firing event*

The actions of a neuron are easily described: It acts when the integrated current surpasses a threshold and fires. Similar to the neuron model, the traders act at the firing time $t_{i,f}$ defined by the condition

$$u(t_{i,f}) = \vartheta_i \,, \tag{2.8}$$

which corresponds to the fact that a neuron fires when the current potential surpasses the threshold ϑ_i in any of the dimensions. When a trader fires, two things happen. The first is obvious: The trader releases a current (so communicates information) pointing in the direction of the information it received after that amount of information reaches a threshold. The second one is important because it tackles the notion of state in the Lux-Marchesi model.

In the Lux-Marchesi model, traders are always in one of the three possible states, and their state plays an instrumental role in determining the market movements. While the Lux-Marchesi model has mean-field dynamic governing state changes, in this newly updated model, agents change state when they fire. The rule is rather simple: When a current potential u_i surpasses the threshold ϑ_i, the agent fires and switches state towards opinion i. In other words, once an agent has aggregated enough information to make a decision for any of its available opinions, it will switch to that opinion and

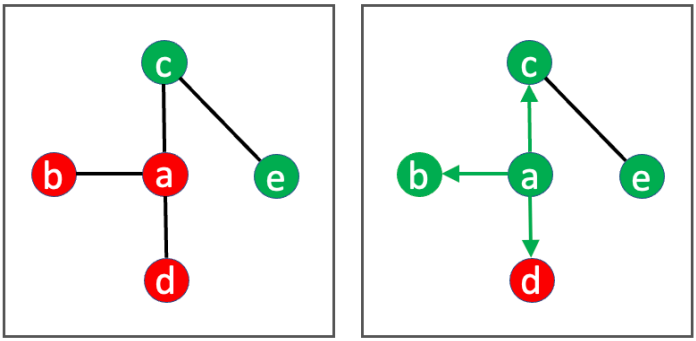


FIGURE 2.1: Simplified network of 2 kinds of traders before (left) and after (right) a firing event. After trader a fires, trader b switches opinion, while traders c and d do not change state. See text for details.

communicate that information to other connected agents (by firing current). After firing, the action potential regarding the particular opinion is reset according to

$$\lim_{\delta \to 0^+} u(t_{i,f} + \delta) = u_{i,r} \,. \tag{2.9}$$

An illustration of the firing event is shown in Fig. 2.1. If trader a receives an input which lets the action potential reach the threshold, that trader will fire and communicate that information to all connected traders. Each trader has three possible scenarios, represented by traders b, c and d. Trader b is close to the threshold, so b changes state (from red to green), while trader c already is in a green state, so that cannot change states. Trader c will, however, continue to aggregate the information until the threshold is reached, at which point c will fire once again. The last trader, d, was not close to the threshold, so it did not change states. However, the action potential of d has increased.

2.1.2 *Market dynamics*

We now have a network of traders but still lack market dynamics. For this model, we introduce market dynamics inspired by the Lux–Marchesi model (Sec. 1.5.2). Similar to this model, the neuron trader model assumes an underlying fundamental value to each asset, also referred to as the fun-

damental price p_f. The market price p can deviate from the fundamental value and is subject to the decisions of traders.

The traders inform their decisions by examining the movement of the market price, its deviation from the fundamental price, and the actions of other traders. Instead of the driving forces in the Lux-Marchesi model, the new model is driven by the current $\mathbf{I}$ received by each trader, which has three parts:

$$\mathbf{I} = \mathbf{I}_{\text{price}} + \mathbf{I}_{\text{net}} + \mathbf{I}_{\text{rand}}. \tag{2.10}$$

(i) $\mathbf{I}_{\text{price}}$ represents the information traders receive from the movement of prices and is the same for every trader:

$$\mathbf{I}_{\text{price}} = \begin{pmatrix} I_f \\ I_+ \\ I_- \end{pmatrix} \propto \begin{pmatrix} \text{abs}(p - pf) \\ \max(0, \frac{dp}{dt}) \\ \max(0, -\frac{dp}{dt}) \end{pmatrix}. \tag{2.11}$$

Every trader receives currents that can point in any direction. Large movements of the market price might lead the trader to estimate increasing ($\frac{dp}{dt} > 0$) or decreasing ($\frac{dp}{dt} < 0$) prices, while large deviations of market and fundamental price might lead to a bet on the long run return of the market to its fundamental value. Thus, the chartist opinions are driven by the derivative of the market price, while the fundamentalists are driven by deviations of market and fundamental prices.

(ii) $\mathbf{I}_{\text{net}}$ is the sum of incoming currents for each trader m:

$$\mathbf{I}_{\text{net,j}} = \sum_n \mathbf{T}_i k_{nm} \tag{2.12}$$

We introduce two new quantities: the trigger $\mathbf{T}_n$ and the connection matrix K, whose entries are k_{nm}. The trigger is the output of neuron n into the network at each time step. Given the nature of the model, the output is zero for most of the time: $T_n = 0$ for all t not in F_n, where F_n is the set of all firing times $t_{f,n}$ of the neuron n. The connection matrix K defines the traders who communicate with each other. For simplicity, we limit connections to the binary case, where $k_{nm} = 1$

(resp. 0) if trader m does (resp. does not) receive information from trader n.

(iii) $\mathbf{I}_{\text{rand}}$, the last term in (2.10), is the random current each trader receives. The random current represents random information that could push each trader into any of the three opinions. This current follows a normal distribution and is different for every trader and every opinion at every time step.

Having introduced how the traders react to the market, we must determine how the market reacts to the decisions of traders. While traders make their decisions based on market and fundamental price, the market price is influenced by the decisions of traders, creating a feedback loop. The Lux-Marchesi model provides excellent market dynamics for this case, which can be directly translated. These rules are summarized here for completeness.

The price is driven by excess demand or supply

$$\Delta E = D_c + D_f, \tag{2.13}$$

where

$$D_c = (n_+ - n_-)\, t_c\,, \tag{2.14}$$

$$D_f = \gamma\, n_f \frac{p_f - p}{p}. \tag{2.15}$$

The number of fundamentalists defines n_f, while n_+ and n_- are the number of optimists and pessimists, respectively, t_c is the mean trading volume per unit of time and γ scales the responsiveness to mispricing. The probabilities $\pi_{\uparrow p}$ and $\pi_{\downarrow p}$ for price changes are

$$\pi_{\uparrow p} = \max[0, \beta\,(\Delta E + \mu)]\,, \tag{2.16}$$

$$\pi_{\downarrow p} = -\min[\beta\,(\Delta E + \mu)\,, 0]\,, \tag{2.17}$$

where μ defines the long-term drift of the price and β the speed of price adjustment.

Now that all the currents and the price behaviour are defined, evaluating model behaviour is straightforward. For each simulation step:

1. Generate currents based on the triggers from the last step, market movement and random noise
2. Update the action potential of the traders depending on incoming currents and evaluate which traders fire and switch state
3. Aggregate the output (triggers) of all connected firing traders for each neuron, which generate the $\mathbf{I}_{\text{net}}$ for each trader
4. Update the market price according to the market dynamics given the new trader distribution
5. Apply decays to the action potentials according to τ

2.1.3 *Model analysis*

We have now finally finished the definition of the model, which combines and applies neuronal decision-making with an agent-based model of financial markets. The model produces a plethora of outcomes, depending on both network structure and parameters, which is why careful attention is needed when analysing model behaviour. A typical simulation produces the following outputs:

Market price and group occupancies

As depicted in Fig. 2.2, the simulation produces a market price, which tends to follow the fundamental price in most scenarios, as well as group occupancies, which determine the movement of the market price.

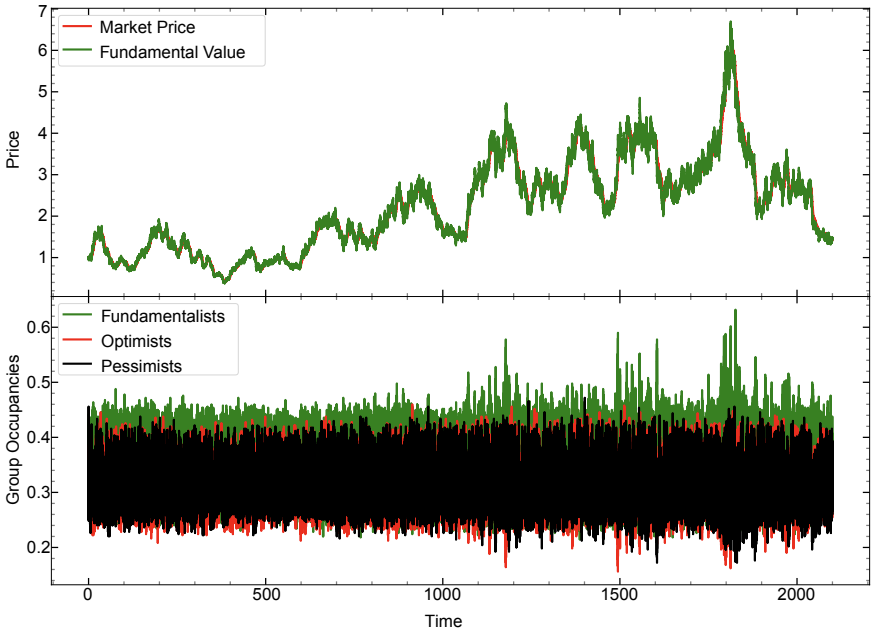


FIGURE 2.2: Simulation output for Barabasi-Albert type graph. Top: Market and Fundamental price output; Bottom: Group Occupancies. The market price closely follows the fundamental value.

Potentials and triggers

The potentials and the triggers allow a glimpse into the models' inner workings. The action potential of each trader determines how close the trader is to making a decision (remember, if the potential hits a threshold, the trader "fires"), while the triggers allow monitoring of the decision-making of agents by measuring their output. An example of this can be seen in Fig. 2.3. Some of the analysis in Chap. 3 will rely heavily on using the action potential to understand agent behaviour.

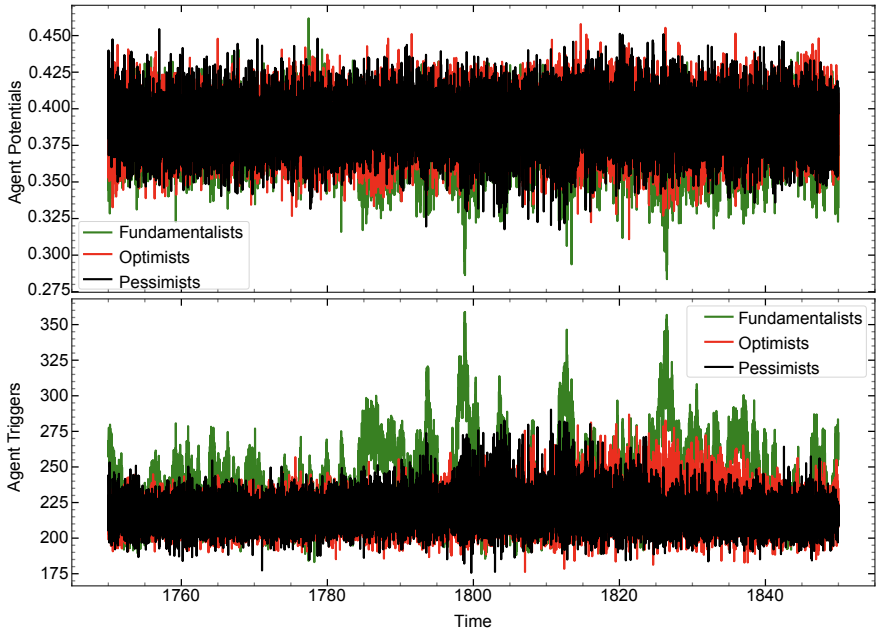


FIGURE 2.3: Simulation output for Barabasi-Albert type graph. Top: Agent potentials; Bottom: Agent triggers.

Excess demand and trader states

Figure 2.4 shows a network of traders at a given timestep, with the colour of the traders indicating their state. We will use the states of agents for some of the more advanced analyses, especially for entropy analysis in Chap. 3. In addition to that, we can also track the excess demand and supply, as defined in Sec. 2.1.2.

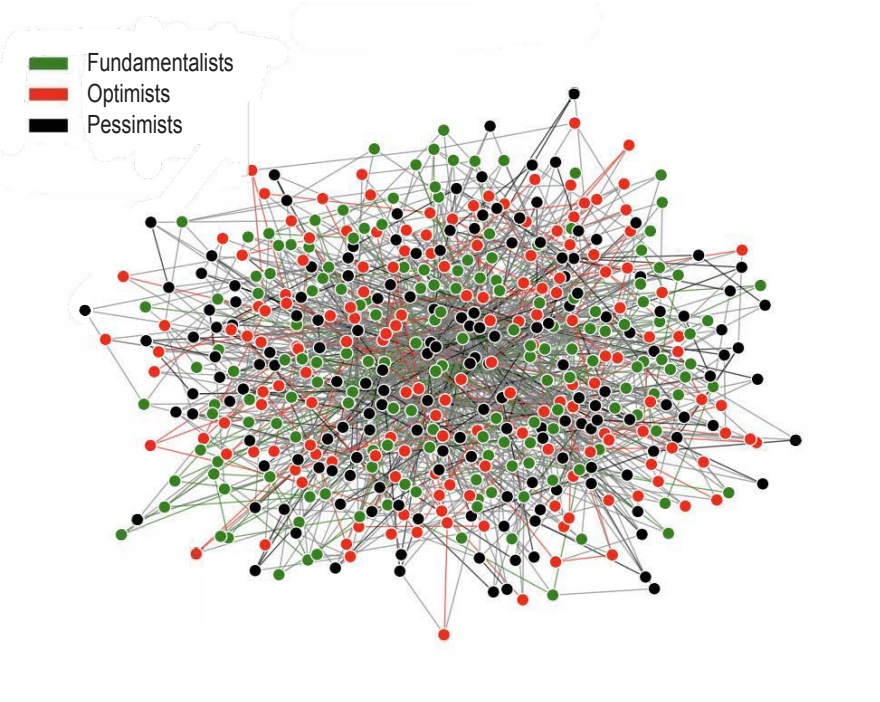


FIGURE 2.4: Network of agents, coloured by states.

The following chapters will show that, despite the interesting observations that can be made by analyzing the above quantities, the individual trader states, together with the network structure defined by the connection matrix K, give even more interesting insights. The rest of this section is two-fold: We will first apply the model to some of the network structures defined in Sec. 1.5 and try to establish relationships between parameters, network structures and model behaviour. Following that, we will apply the model to a variety of more exotic networks, selected to demonstrate the wide variety of behaviour that can be produced by the model. The following chapters will then use the model to investigate concrete scenarios. After investigating the effects of irrationality and information asymmetries (hierarchies) in Chap. 3 and Chap. 4, we will formalize the model Chap. 5 and apply it to a more elaborate model of stock traders. Finally, we will show in Chap. 6 how the integrate-and-fire model can be applied in the real world to accurately model contagious processes.

2.2 THE ROLE OF PARAMETERS

We are now faced with a daunting problem - How do we analyse a model that has the potential for almost endless complexity? Since the agents are organized in a network structure, there are, in theory, 2^N possible parameters for the connections between agents alone. We decide on a methodological approach: We first run the model on different configurations of the most important network used to describe social interactions with varying parameters. This allows us to not only investigate the effect of model parameters, but also will help us to demonstrate how network structure can drastically change the model response to parameter choice.

As we will see in the following analysis, the effect of changing a parameter can sometimes vary when applied to different network structures. The key parameters we will investigate are the decay constant τ and the parameters ω_1 and ω_2, reflecting the relative influence of market information and trader influence respectively.

We will investigate the effects of these parameters for three of the most relevant networks for social applications, the Erdos-Reny, the Watts-Strogatz and the Barabasi-Albert type network, which have all been introduced in Sec. 1.5. This first step in model analysis is necessary to understand how the fundamental elements of the model interact; however, as we will see later in this chapter, it allows only a glimpse into the fascinating model behaviour. Thankfully, with custom-made tools and measures we will be able to uncover some unexpected features of the model and will continue to expand on the analysis in future chapters.

2.2.1 *The effect of market price information ω_1*

We introduce the parameter ω_1, which defines to which extent the traders are subjected to market price information. As explained in Sec. 2.1.2, the currents a trader receives are given as

$$\mathbf{I}_{\text{price}} = \begin{pmatrix} I_f \\ I_+ \\ I_- \end{pmatrix} \propto \begin{pmatrix} \text{abs}(p - pf) \\ \max(0, \frac{dp}{dt}) \\ \max(0, -\frac{dp}{dt}) \end{pmatrix}. \tag{2.11}$$

The traders receive this information, and the parameter ω_1 determines the fraction of this information that is absorbed and thus added to the action potential. In the context of the integrate-and-fire model, it can be thought of as a kind of resistance. Fig. 2.5 shows the relationship between ω_1 and the average firing rate of chartists for different kinds of network configurations.

There are two important takeaways from this figure: First, increasing market information leads to an increased firing rate. The firing rate describes the rate of state changes, so the frequency of traders' decisions. This increase is to be expected - when there is more information available, it is not a surprise that agents make more decisions. On the other hand, it exemplifies the amplifying effect the network can have on traders' reactions to information. The more strongly connected networks ($m = 8$ for Barabasi-Albert, $p = 0.03$ for Erdos-Reenyi and $p = 0.1, k = 10$ for Watts-Strogatz) react much more strongly to an increase in information than the other networks.

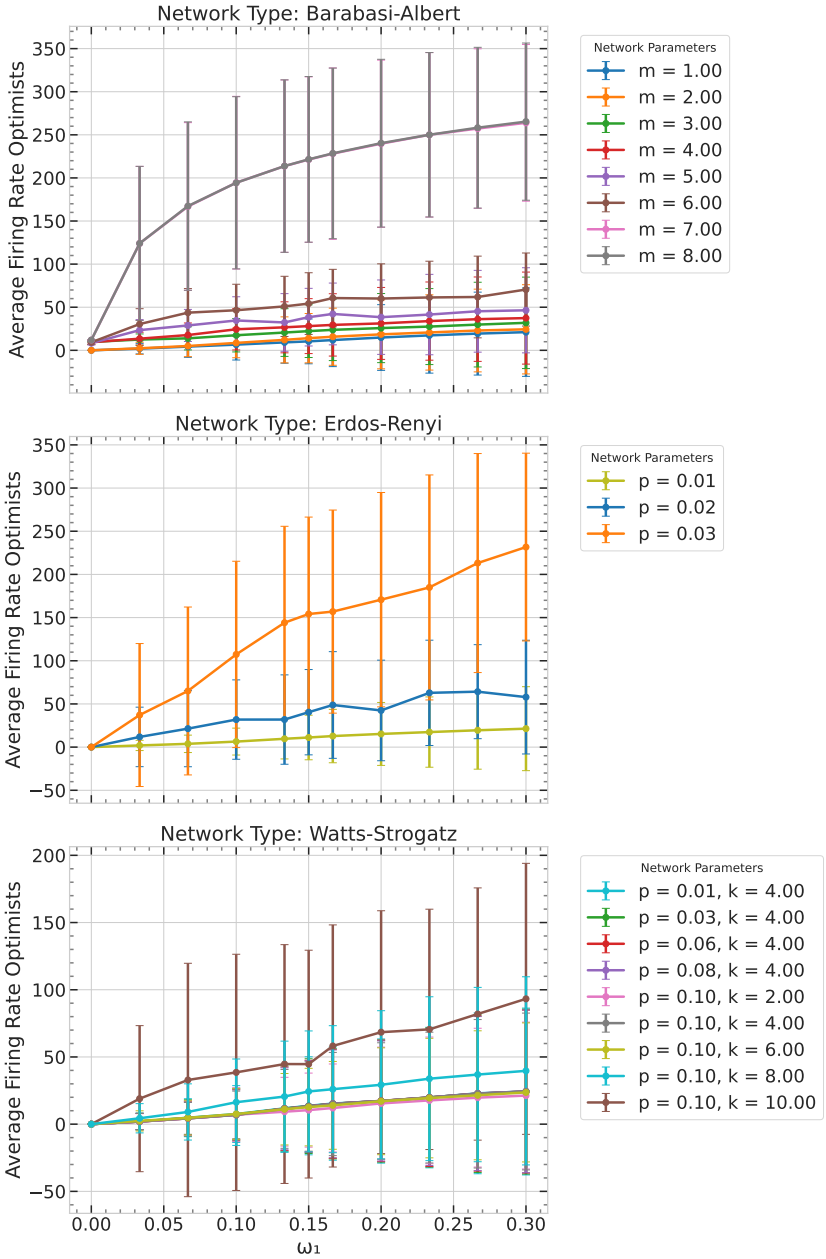


FIGURE 2.5: Average firing rate as a function of the amount of market information in the integrate-and-fire trader model for different network configurations. Error bars for this and and the following plots are standard deviations of measured values.

A similar picture can be seen in Fig. 2.6, which shows volatility with increasing market information. We have introduced volatility in Sec. 1.3 as a crucial measure of market behaviour. Volatility also tends to increase, though in the Barabasi-Albert network, there are clearly two regimes: the highly connected networks always display more volatility, while the low-connected networks need significantly more information.

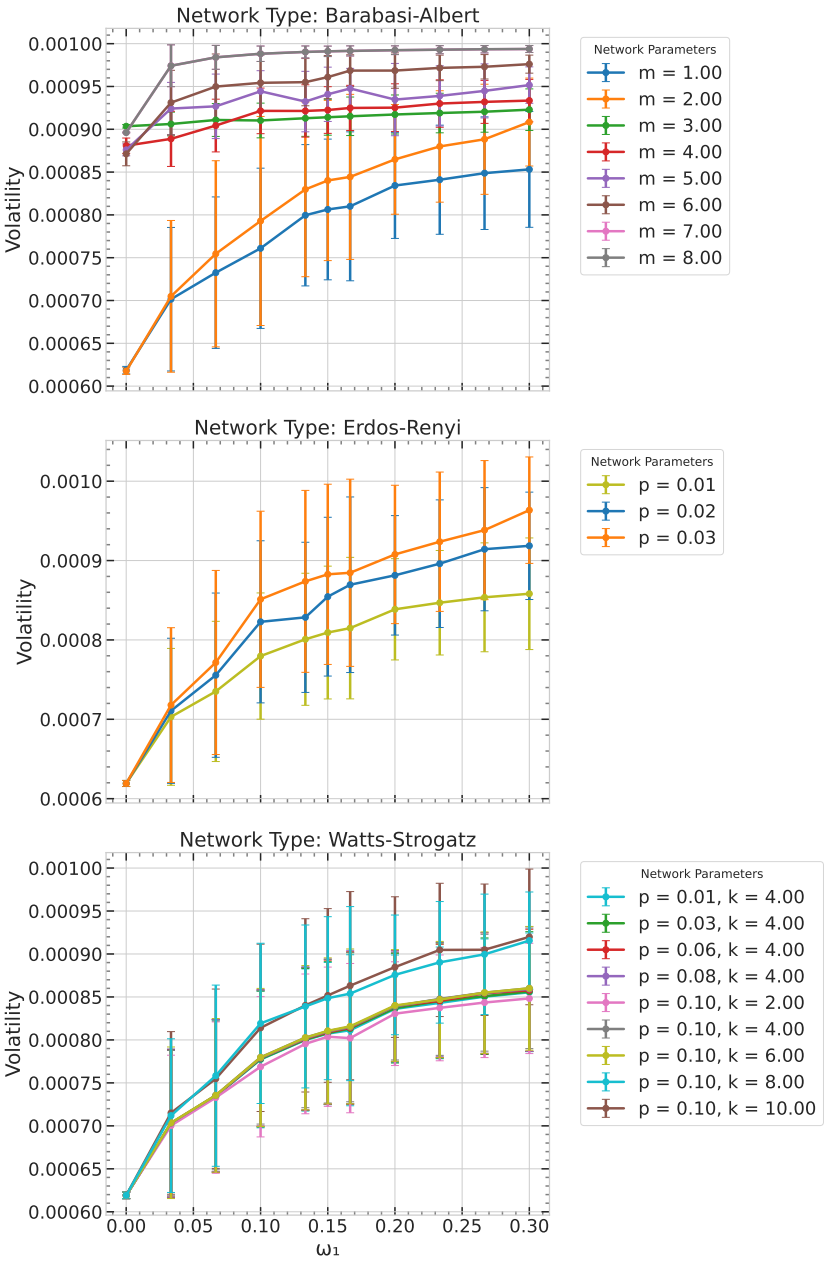


FIGURE 2.6: Volatility of market price as a function of the amount of market information in the integrate-and-fire trader model for different network configurations.

Figure 2.7 underscores the important effects that different network structures can have on agent behaviour: It displays the agent potential for different parameter choices. Agent potentials are a measure of how close agents are to make a decision - In a regime with high potentials, small influxes of currents can lead to avalanches of trader behaviour, so they are one of the indicators of criticality in the model. We will focus on the unique insights that can be gained from the agent potentials in Chap. 3.

For the Barabasi-Albert network, the average potential is mostly defined by the network structure, with additional information only having a minor influence on the strongly, connected networks. On the other hand, for both the Erdos-Renyi and the Watts-Strogatz networks, we see an increase in activity for weakly connected networks and a decrease for strongly connected networks. This is mostly due to the fact that in very strongly connected networks, slightly more information can lead to a significantly higher firing rate (Fig. 2.6), which means that most received information is immediately shared with others.

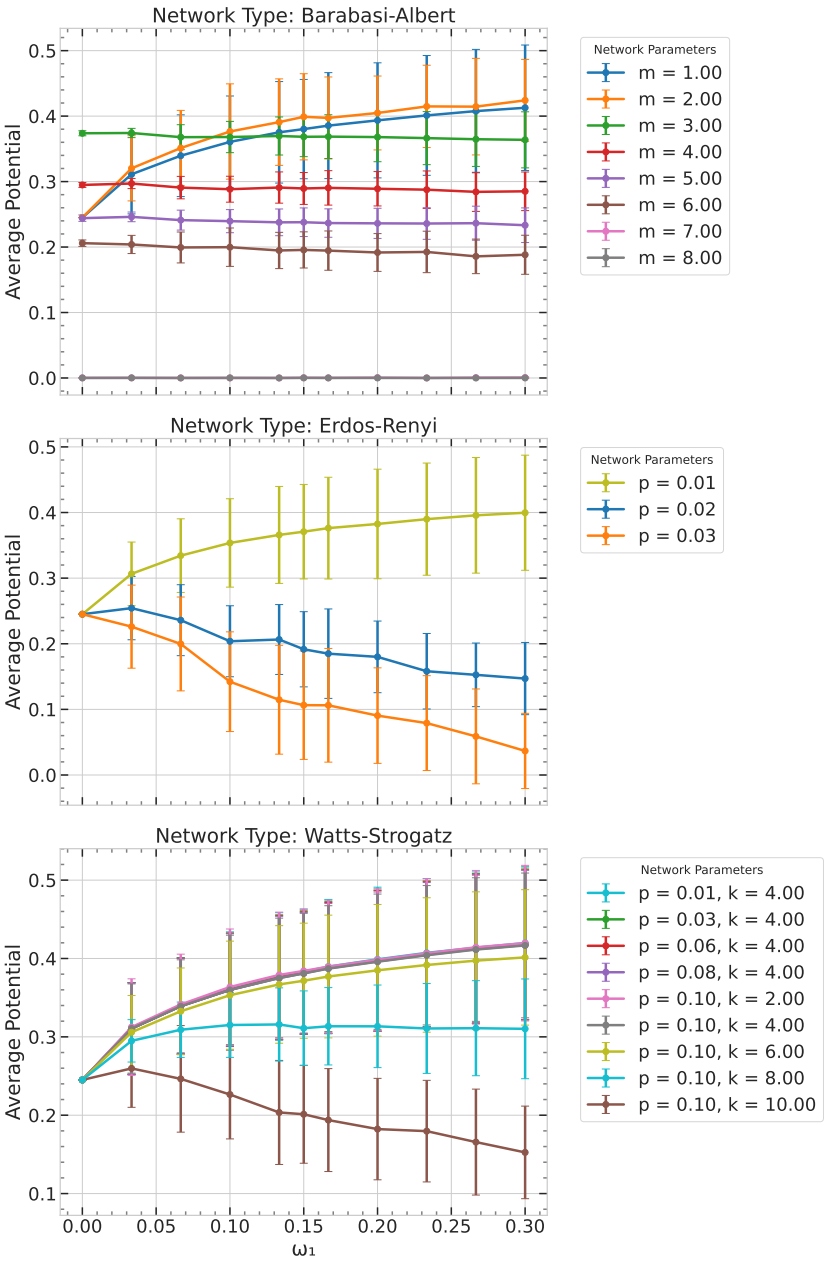


FIGURE 2.7: Average potential as a function of the amount of market information in the integrate-and-fire trader model for different network configurations. The threshold $\vartheta = 1$ for all simulations.

2.2.2 *The effect of network information*

While ω_1 controls the effect of market information, ω_2 controls the information directly received from other traders. Again, we see a similar picture for both the firing rates and the volatility. As Fig. 2.8 reveals, while more network information generally increases the firing rate, the effect is much more pronounced in strongly connected networks, which can undergo a radical shift in activity with a small increase in available information.

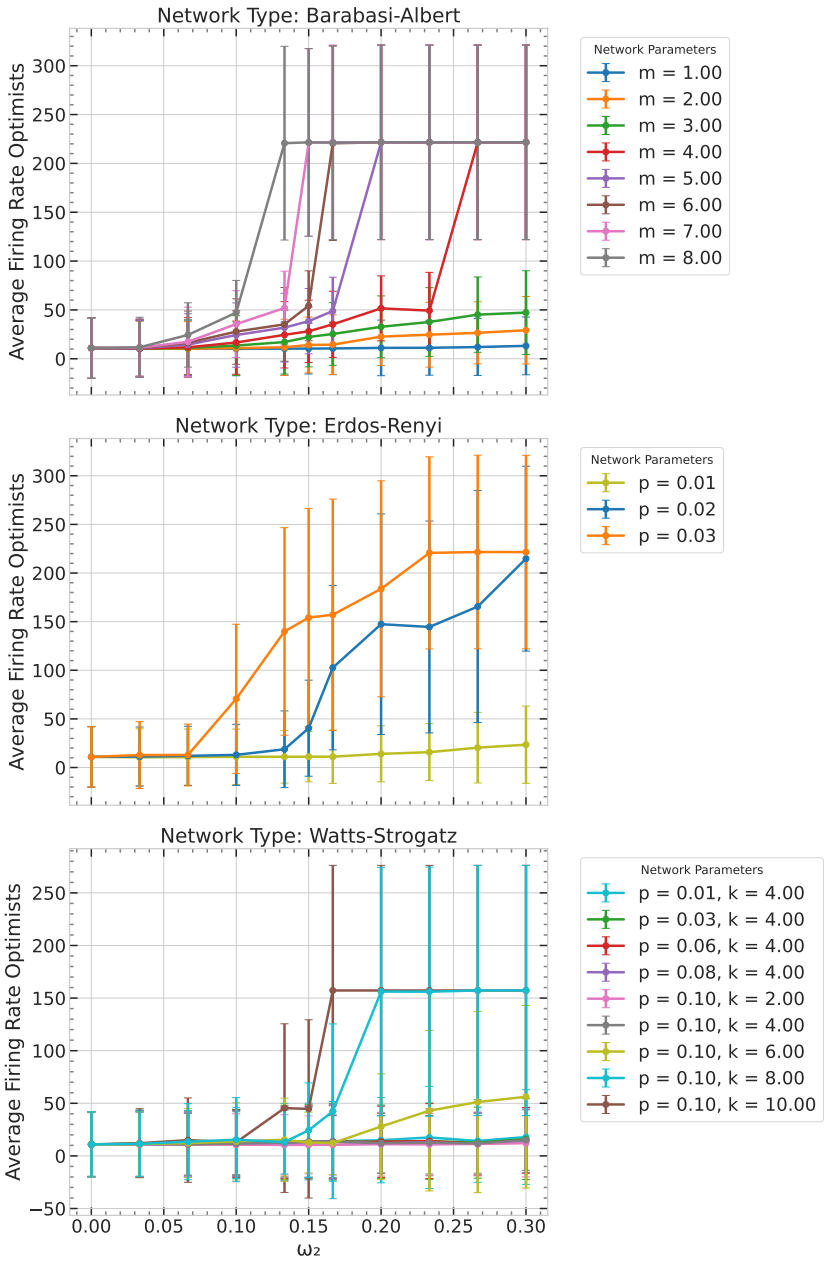


FIGURE 2.8: Average firing rate as a function of the amount of network information in the integrate-and-fire trader model for different network configurations.

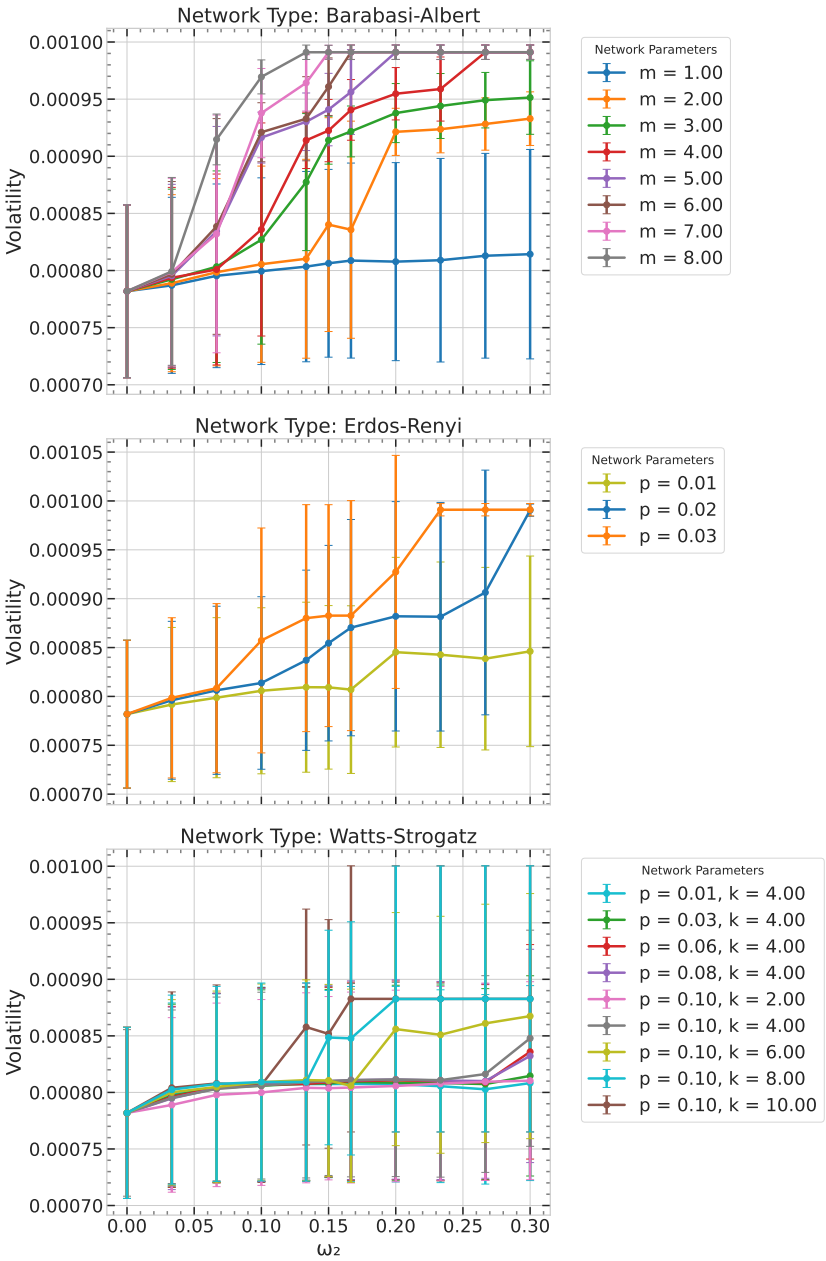


FIGURE 2.9: Volatility as a function of the amount of network information in the integrate-and-fire trader model for different network configurations.

Similarly, Fig. 2.9 shows that volatility increases most in strongly connected networks, once again confirming that the network structure has an amplifying effect on the reaction of traders to the available information.

2.2.3 *The effect of the decay variable τ*

The decay variable τ determines at which rate accumulated information is lost in the agent potential. With fast decay, only recently acquired information counts, so market movements only happen when there is a lot of new information at a given timestep. However with slow decay, agents will make their decisions given the information of many previous timesteps. Figure 2.10 reveals that the Barabasi-Albert type network reacts very little to an increase in decay, with the average firing rate mostly determined by the network structure. On the other hand, increasing the decay leads to a significant reduction in activity for the Erdos-Renyi type network.

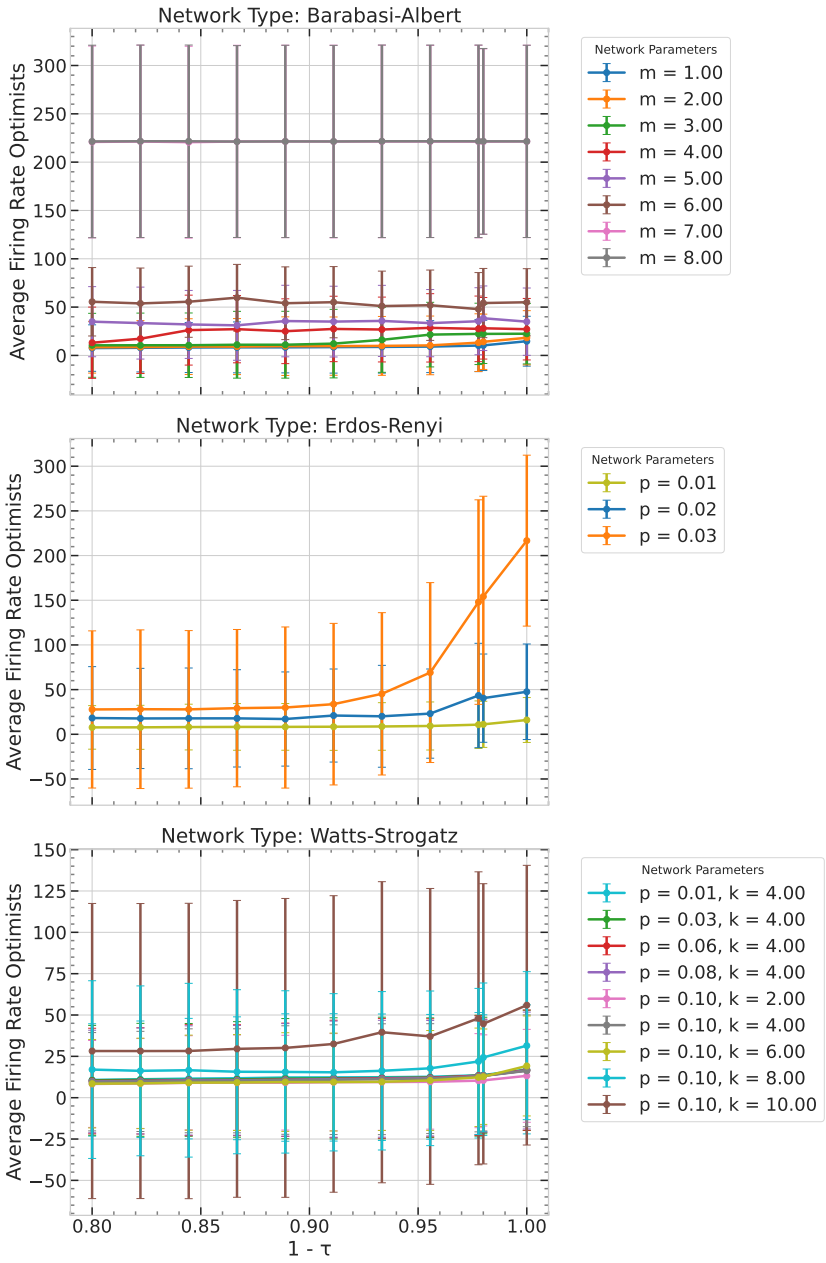


FIGURE 2.10: Average firing rate as a function of the decay of trader potentials in the integrate-and-fire trader model for different network configurations.

Curiously, this trend is not confirmed when investigating the average potential for the same sets of parameters. Here, an increase in potential can be observed with decreasing decay for the more weakly connected Barabsi-Albert type networks, even though Fig. 2.10 confirms that this increase does not convert into increase activity. On the other hand, in the Watz-Strogatz model, the potential generally increases with decreasing decay, while for the Erdos-Renyi network, it can have positive or negative effects, depending on connectivity.

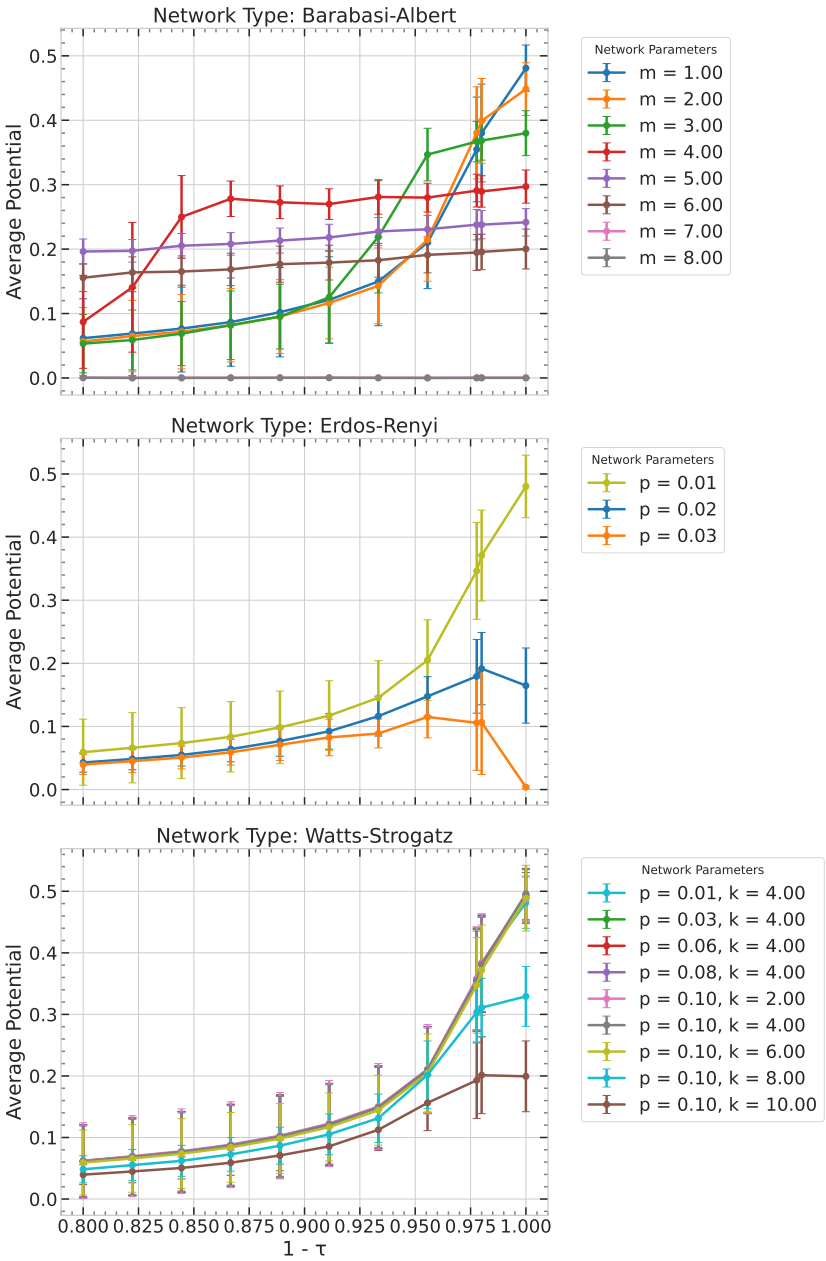


FIGURE 2.11: Average potential as a function of the decay of potentials in the integrate-and-fire trader model for different network configurations. The threshold $\vartheta = 1$ for all simulations.

Volatility shows similar behaviour as the firing rate: For the Barabsi-Albert network, volatility is mostly defined by the network structure, while other network types see an increase in activity with decreasing decay.

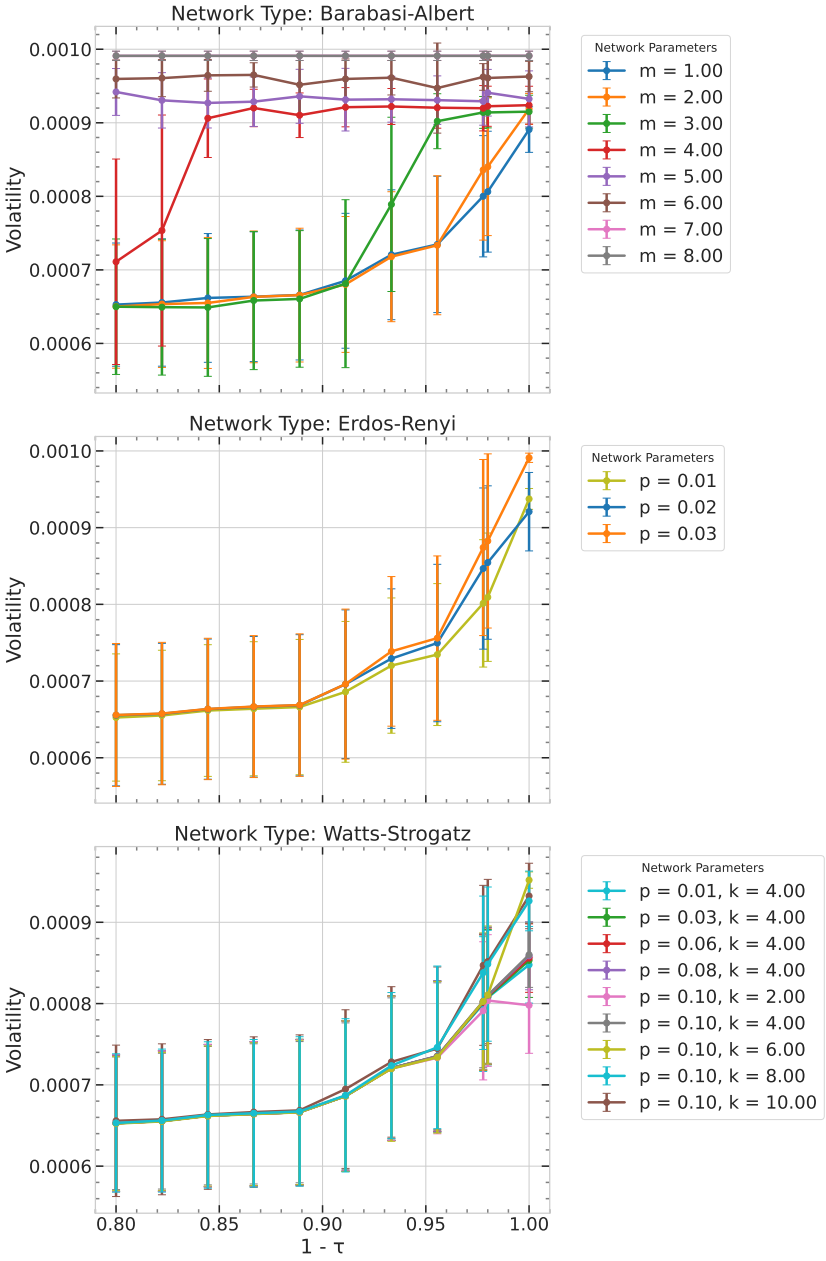


FIGURE 2.12: Volatility as a function of decay of information in the integrate-and-fire trader model for different network configurations.

Finally, we fit the return distributions for multiple sets of parameters to investigate the tail structure. We have established in Sec. 1.3 that we expect returns to be fat-tailed for an accurate description of stylized facts. We fit an α-stable distribution, where the α parameter determines the shape of the distributions, with $\alpha = 2$ giving a normal distribution and $\alpha = 1$ giving a Lorentz distribution. Again, the shape of the return distribution is largely determined by the network structure for the Barabasi-Albert type network, while the Erdos-Renyi and the Watts-Strogatz model depend heavily on the decay variable, with non-normal distribution only emerging when there is little decay. An example of fitting such a non-normal return distribution is depicted in Fig. 2.13.

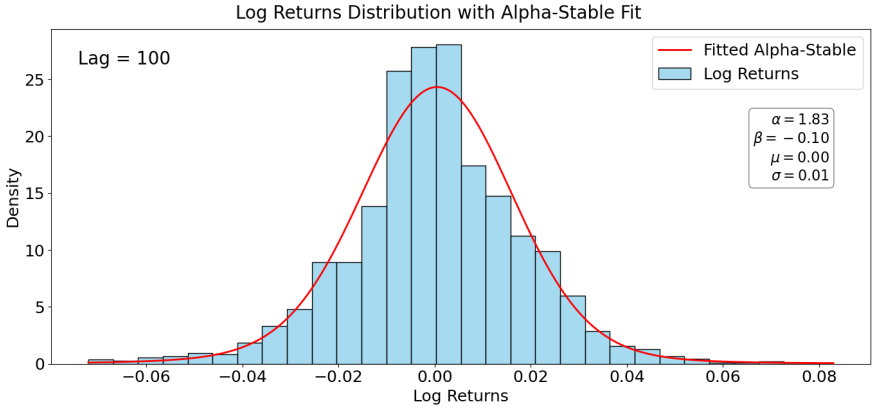


FIGURE 2.13: Return distribution for simulation run on a Barabasi-Albert type network with $m = 2$, $n = 500$ and parameters $\tau = 0.98$ and $\omega_1 = \omega_2 = 0.15$

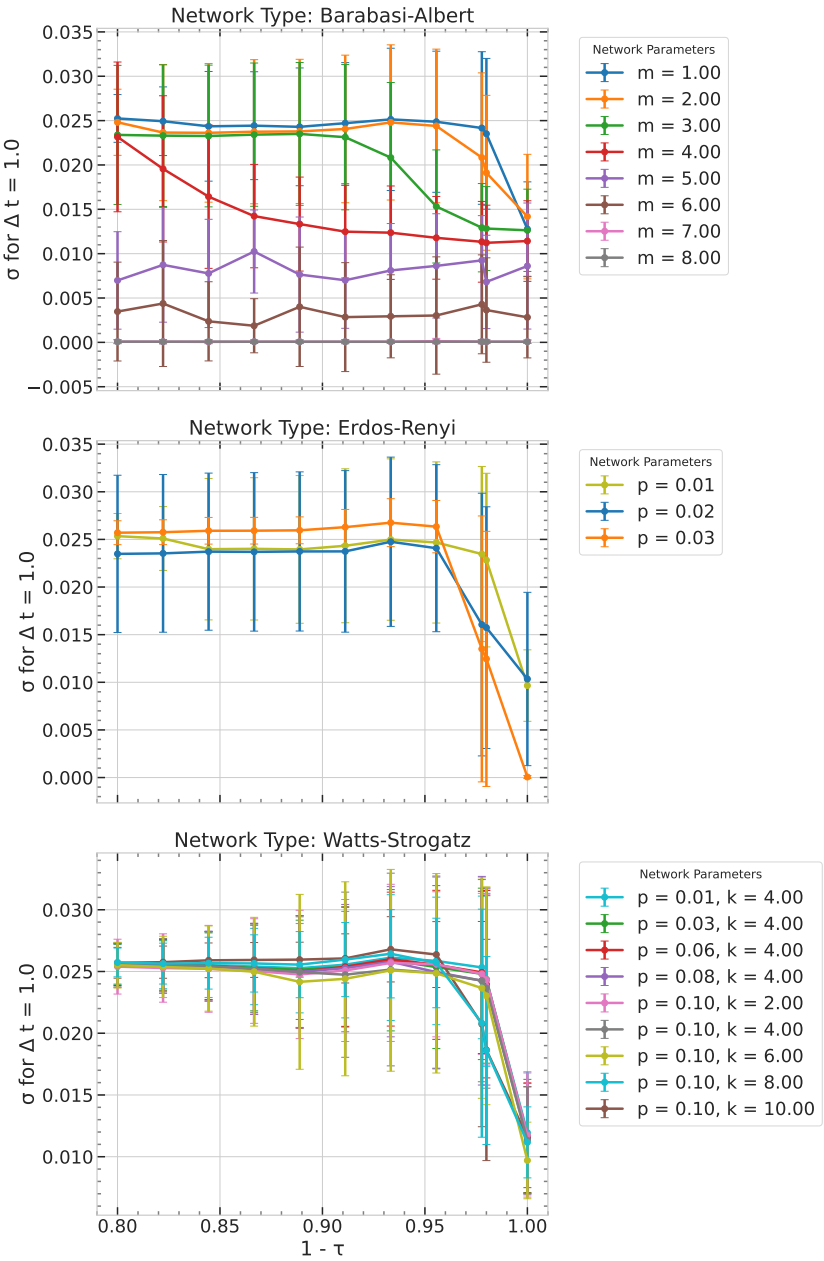


FIGURE 2.14: Stable distribution α as a function of the decay of information in the integrate-and-fire trader model for different network configurations

Now that we have a basic understanding of the effects of parameters, we will delve into a much more detailed analysis of trader behaviour – It turns out that this is where we find the most interesting behaviour.

2.3 NETWORK ANALYSIS

Despite other parameters, which, of course, affect the traders, we have seen that the network structure, defined by the connection matrix K, has a large effect on trader formation. Figure 2.15 shows a network of agents called a Gaussian random partition graph [127]. As the name suggests, such graphs have multiple partitions, which are the clearly identifiable subgroups in Fig. 2.15. The structure of such a graph is defined by the probabilities p_{in} and p_{out} for nodes to be connected to other nodes that are inside or outside their own cluster, respectively, the number of agents N, the mean cluster size n and a shape parameter v. To generate such a graph, the nodes are first split into k multiple partitions, each size of which follows a Gaussian distribution $N(n, v)$.

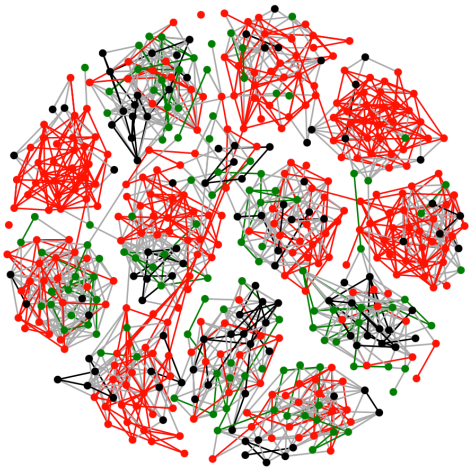


FIGURE 2.15: A Gaussian random partition graph network with 500 nodes, $p_{in} = 0.1$, $p_{out} = 10^{-5}$, mean cluster size 40 and shape parameter 50. Optimists are red, pessimists black and fundamentalists are green. Links are coloured accordingly if both nodes belong to the same group.

Gaussian random partition graphs are particularly useful for investigating the effects of clusters in networks and also allow for easy visualisation of some of the mechanisms of the model. That is why the simulations in the following section have been obtained with such network structures, while the results in Sec. 2.5 are based on a variety of network structures.

2.3.1 *Network structure and trader behavior*

In this section, we introduce several measures to characterize network structure and agent behaviour. We show that the connectivity drives critical behaviour, so the average number of connections per agent that can be observed with energy and entropy measures.

2.3.2 *Centrality measures*

There are three general measures for centrality, termed "degree", "closeness", and "betweenness" [128]. The simplest of these, the degree of a node n_i, is the number of adjacent nodes. Closeness centrality measures how "close" nodes are to each other. Distance usually means the geodesic distance, which is the minimum number of edges between two nodes or, in other words, over how many edges one node is connected to another [128].

The third measure, betweenness, is the number of times a node is on the geodesic between two other nodes. This measures how central a node is since any node that is on a path between two other nodes can control the information passed on. There is no guarantee that information will take the shortest path to travel between nodes, especially in complex network structures. However, these measures are an important factor in identifying the flow of information and can be related to patterns in our model, as described in the next section.

A lesser-known measure, the current flow closeness centrality, has a significant impact on individual trading behaviour. This is described in the next section.

2.3.3 *Current flow closeness centrality and cluster size distribution*

Current flow closeness centrality [128] is particularly suitable for our model, since it has the same underlying assumption: information flows through

the network – like electrical current. It takes all paths into account, unlike other measures, which look only at shortest paths. The fact that current flow closeness centrality is also referred to as information centrality underscores its applicability.

To calculate this centrality, we first define the resistance distance R_{nm} as the effective resistance between the nodes n and m when each edge has a given resistance. The current flow closeness centrality is then given as

$$C(n) = \frac{N}{\sum_{m=1}^{N-1} R_{nm}}, \tag{2.18}$$

where N is the number of nodes in the network. Since our networks are generated before the simulation, analyzing these has little value in itself. Since the exchange of information is central to the model, it is necessary to obtain measure of the travel of information. One such measure is the cluster size distribution.

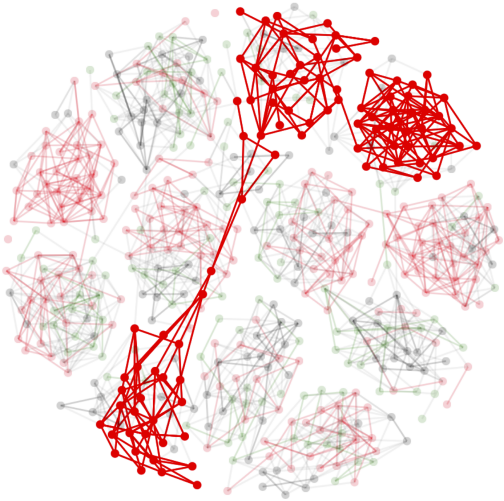


FIGURE 2.16: Illustration of a cluster in a network. We define clusters as groups of connected agents with the same opinion, analogous to percolation theory [129].

A cluster is defined as a group of connected traders with the same opinion (Fig. 2.16). The span of such a cluster gives a measure of whether there is global trader belief or whether small groups of traders form their opinions independently. The cluster size distribution is also a measure of complexity. By comparing the distributions with each other and with randomly occupied states, we can get a glimpse at the formation of communities. Fig. 2.17 shows that the system indeed follows complex dynamics.

For a system of randomly occupied states, the cluster size distribution would not differ between individual trader groups. In the neuron trader model, however, the fundamentalists are typically in much smaller clusters than the chartists, whereas the optimists and pessimists have the same average cluster size. This might result from the fundamentalists being less community-driven. Additionally, the fraction of chartists is higher than the fraction of fundamentalists. The fact that the overall frequency of the optimist clusters is higher can be attributed to the algorithm, which identifies all clusters in a simulation step and sorts them by trader group. Even though the mean size is similar, there are typically more optimist clusters in this simulation. This is interesting since the information travel distance seems to be similar even if the popularity of an opinion differs slightly. The next step is to contrast the cluster size distribution with a network measure, the current flow closeness centrality.

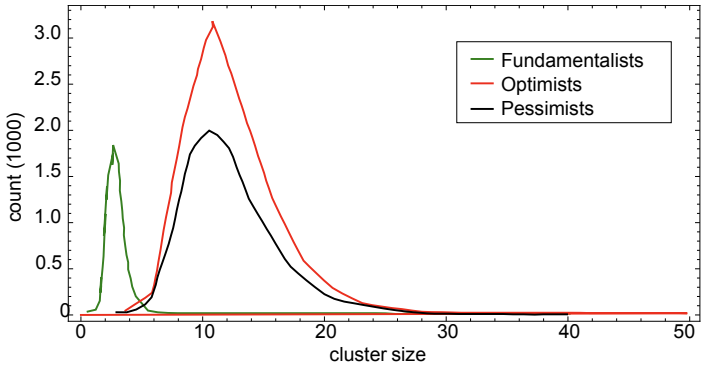


FIGURE 2.17: Cluster size distribution over a simulation for each of the three different trader opinions.

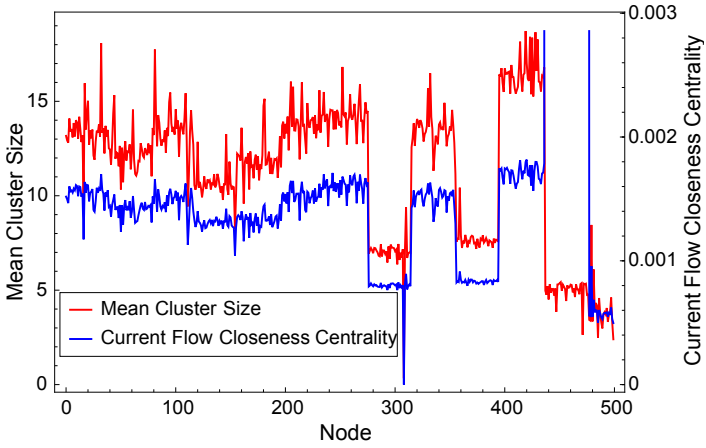


FIGURE 2.18: Mean cluster size and current flow closeness centrality as function of node number. At each simulation step, cluster sizes are computed for each node, then we average over the simulation run. The sharp drop at 308 can be attributed to a single node with no connections.

Figure 2.18 shows the relationship between mean cluster size and current flow closeness centrality. For the majority of the nodes, there is a clear linear relationship. However, a pattern change is seen for the partition which incorporates nodes between 440 and 480. There is still a linear relationship, but with a change in the scaling parameter. Since the network is randomly generated, a partition with a much higher than average current flow closeness centrality can occur. The effects of such a partition are not easily described by this centrality measure, so other measures will be needed to explain this shift.

2.3.4 *Energy*

While entropy calculation requires approximations and computational optimization to be feasible, energy calculations are straightforward for this system. Since we have a multi-state system of connected nodes, the energy can be calculated using a Potts-like Hamiltonian:

$$H = J \sum_{\langle n,m \rangle} \left(\delta(s_n, s_m) - 1 \right), \tag{2.19}$$

with

$$\delta(s_n, s_m) = \begin{cases} 1 & \text{if } s_n = s_m \\ 0 & \text{if } s_n \neq s_m. \end{cases} \tag{2.20}$$

This means that there are nonzero contributions for every pair of connected traders that are not in the same state.

While the calculation may be straightforward, the interpretation of the value is still interesting. Since this measure of energy corresponds to the alignment of connected traders, it can be seen as a measure of how much connected traders tend to agree or disagree. As can be seen in Fig. 2.19, the model dynamics lead the system to a state of higher energy than random states. Since energy is a measure of how much adjacent traders tend to agree, it can be interpreted as a measure of community formation. The higher the energy, the more traders tend to agree with their neighbours. Despite having a higher mean energy, the spread in energy is much wider than for a random system. This corresponds to the fact that the market is not in equilibrium. The mean energy is not constant in time but varies significantly, just as the entropy does. We will delve into how to calculate entropy for the model in Chap. 3, where we will show it to be an essential tool for understanding market behaviour.

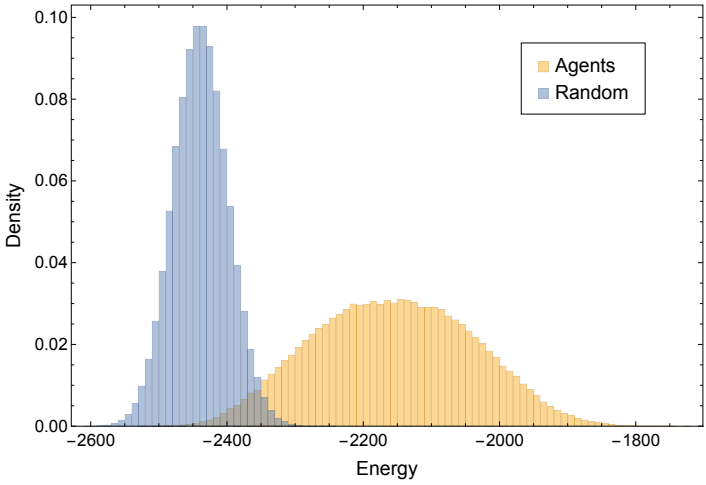


FIGURE 2.19: Histogram of values of energy for a simulated system and randomly occupied nodes.

2.3.5 *Connectivity and criticality*

Having established energy as a measure of community formation, we now investigate how this measure reacts to changes in the network. The most straightforward way to do this is to increase the connectivity of the system. The connectivity of the system is the probability of one node being connected to a randomly selected different node. A sparse network has a low connectivity, while a dense network has a high one.

When measuring the energy for systems with increasing connectivity, an interesting pattern is observed (Fig. 2.20). Here, we contrast the energy of the system with the shannon entropy. See Chap. 3 for a detailed explaination how we calculate entropy for the integrate-and-fire model. While initially decreasing with connectivity, which corresponds to traders agreeing more when more densely connected, the system enters states of higher energy when the connectivity is increased even further, until a critical point is reached. After the critical point, which represents the maximum of energy, so maximum disagreement and the minimum of community formation, the energy drops sharply, indicating that the traders come to a state of widespread agreement. Contrary to the energy, entropy rises continuously until it reaches the critical point. However, the critical point denotes the maximum of entropy and is followed by a sharp drop.

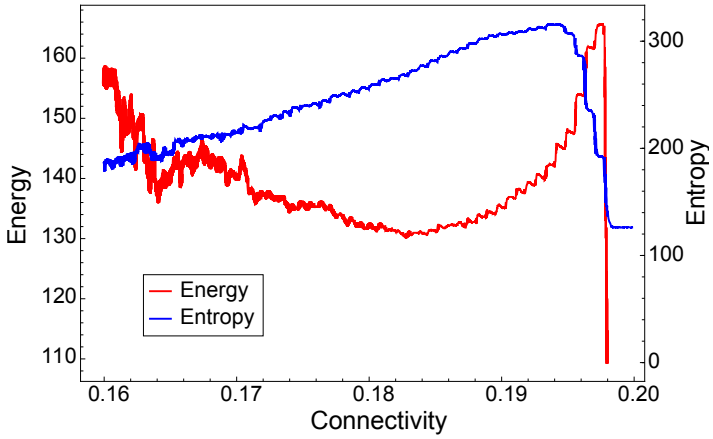


FIGURE 2.20: Absolute energy and entropy as a function of the connectivity. Note that the drop in entropy precededes the peak in energy.

Figure 2.20 leads to an interesting conclusion: While increasing links can lead to the formation of larger communities and more agreement, a more disordered system results until a critical point is reached. At this critical connectivity, the market enters in a different phase which is characterized by less disorder and greater agreement.

2.3.6 *Price change scaling*

In addition to the interesting behaviour with regard to connectivity, the model shows surprising effects when the probability for price changes is increased. As discussed in Sec. 2.1.2, the probability for a price change at each time step is defined by the excess demand or supply created by the opinions of traders. If this excess demand is rescaled, the probability for price changes can be adjusted. We expect a higher probability for price changes, so a more volatile market, to attract more chartists, as is indeed the case (Fig. 2.21). While the decrease in the number of fundamentalists with higher market movements is not surprising, the effects on the actual market movement are indeed fascinating. Perhaps contrary to expectations, the increased number of chartists has a stabilizing effect. To explain this, we first calculate the Hurst exponent.

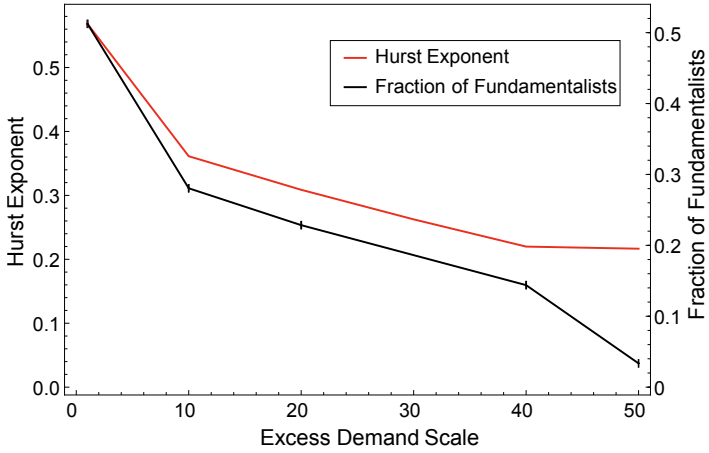


FIGURE 2.21: Fraction of fundamentalists in the system for different scaling of excess demand

2.3.7 *Hurst exponent*

The Hurst exponent [130] H describes the long-term dependence of a time series. There are three cases: (i) $H < 0.5$ (mean-reverting), (ii) $H = 0.5$ (uncorrelated), and (iii) $H > 0.5$ (trend-reinforcing). A common way to estimate the Hurst exponent is with a rescaled range calculation [131]. For a partial time series X_i, the rescaled range is calculated as the ratio of the range $R(n)$ and the standard deviation of the mean adjusted time series $S(n)$. The range $R(n)$ is

$$R(n) = \max(Z_1..Z_n) - \min(Z_1...Z_n)\,, \tag{2.21}$$

where Z_t, the cumulative sum up to time step n,

$$Z_t = \sum_{i=1}^{t} Y_i, \tag{2.22}$$

Y_i is the mean adjusted time series

$$Y_i = X_i - \bar{X} \tag{2.23}$$

with $S(n)$ the standard deviation of Y_i, with $i = 1, \ldots, n$. The Hurst exponent measures how the rescaled range $R(n)/S(n)$ scales with the interval length. If the rescaled range does not scale with increasing intervals, it is mean reverting and will have a Hurst exponent close to zero. Strong scaling will lead to a Hurst exponent close to one.

Given the dominance of chartists in systems with increased probabilities for price changes, the Hurst exponent might be expected to increase since a larger number of chartists would lead to more speculation. However, Fig. 2.21 shows the opposite is the case. Instead of leading to a higher rescaled range, chartists actually have a stabilizing effect and prevent large deviations of market and fundamental price. This is due to every surge of optimists, which drives the price up, being immediately followed by a surge of pessimists, which drives it back down and *vice versa*.

Even though we have presented several measures in this section, the model exhibits behaviour which cannot be quantified as easily. In the following section, the challenges of model visualisation and calibration are discussed.

2.4 VISUALISATION AND CALIBRATION

2.4.1 *Visualisation*

The measures introduced in Sec. 2.3 might lead to the impression that there are easily implementable measures suited to capture the entirety of the model dynamics. Unfortunately, this is not the case. Vastly different model structures can lead to similar results for these measures, even though the underlying dynamics are dissimilar.

One of the biggest problems when analyzing a system is that it is not possible to implement a measure to investigate a pattern without observing it first. This is why the development of a visualisation tool was necessary for a meaningful analysis. It is a particularly difficult challenge to display the dynamically evolving system of agents with static measures. Behaviour is often distinct not just between different points in time but also between different traders.

For this purpose, an online tool was developed which can display the network of agents, all meaningful timelines and real-time states of the agents interactively. By scrolling through the timelines, we can discover patterns which would have been difficult to discover otherwise. The results and some of the plots in Sec. 2.5 were obtained with this visualisation tool.

2.4.2 *Calibration*

One of the biggest problems of agent-based models is proper calibration [59]. During simulation runs, the model produced a wide range of behaviour, including fat-tailed return curves, volatility clustering, and mean-reverting and trending behaviour, depending on parameters. The large variety of behaviour the model can produce suggests that it should be well-calibrated.

The goal is to find the network structure with the set of parameters that can accurately reproduce stylized facts of financial markets. There is, however, an obvious problem: the model is reliant on a network of agents, which makes its parameter inputs high-dimensional, since for N agents, there are 2^N possible connections. This fact, combined with the issue that any proper analysis of simulations can only be done in its entirety by manual inspection, makes this a nearly impossible challenge.

Having the ability to analyze a simulation in its entirety was an important step in analyzing model behaviour, but it is not of much use when trying to evaluate millions of potential simulations. To calibrate the model, better tools are needed. In this work, a first step in this direction is used to evaluate simulation results.

The use of machine learning has had an increasing impact in science over the last several years and is increasingly used in applications such as the search for new galaxies [132] and particles [133]. In econophysics, machine learning has been shown to have the ability to identify stylized facts in financial time series [134]. Here, this ability has been used to evaluate simulation results. A neural network was trained to distinguish "real" from generated financial time series. After training, over a million simulations with random parameters were generated and the results were evaluated using the neural network as a discriminator. This made it much easier to evaluate simulation results because an overwhelming majority of simulations can be discarded without further consideration.

Patterns from the resulting simulations were analyzed with the help of snapshots of the tool described in the preceding section and the most interesting observations are described in Sec. 2.5. It must be stated, however, that the calibration using these methods is still insufficient since the space of possible parameter combinations is too large to evaluate all possible simulations. In order to find the best parameters, other methods are needed.

2.5 SIMULATION RESULTS

This section relies heavily on images to describe patterns. This is primarily due to the complexity of the model and the lack of tools to describe it fully. Finding better ways to characterize this behaviour will be an important future undertaking (even though we have already made significant progress with the tools used in Sec. 2.3).

The complexity of the patterns is one of the main reasons why the use of machine learning has immense potential for this kind of system, particularly concerning pattern detection and relating the occurrence of patterns to network structure and other simulation parameters. Nevertheless, it is essential to describe some key patterns which were identified to show the wide range of behaviour that the model can produce.

2.5.1 *Network structure and community formation*

In the following, we will illustrate the patterns by showing and explaining examples of simulation runs.

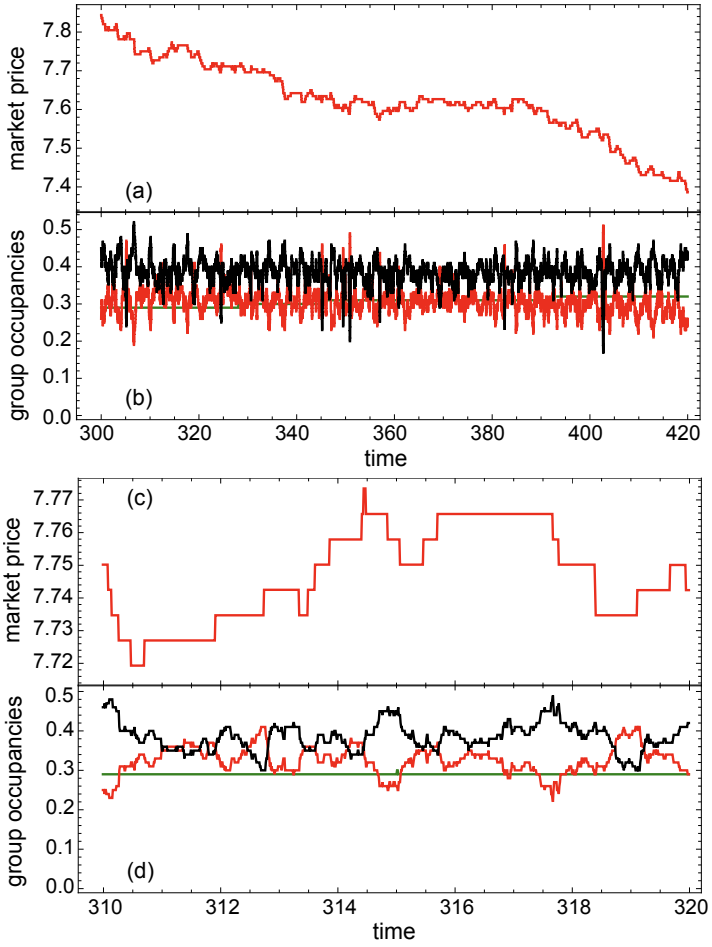


FIGURE 2.22: Trader distribution and market price for a simulation which shows community formation. (a,b) displays an overview of the simulation, and (c,d) a small region of the simulation. For the group occupancies, optimists are depicted in red, pessimists in black and fundamentalists in green.

Section 2.3.4 discussed community formation from the dynamics of the model. The simulation shown in Fig. 2.22 and Fig. 2.23 is a perfect example of such behaviour. Figure 2.22 shows timelines of trader distributions and market prices. Several interesting observations can be made. First, the number of fundamentalists stays relatively constant throughout the simulation run. Secondly, loops of trader movement can be identified when looking at Fig. 2.22c and d, which have a greater time resolution. While fundamentalists keep their opinion, there are frequent transitions between optimists and pessimists.

Figure 2.23 reveals what could already be suspected from the previous analysis: There is a community of fundamentalists who are constant in their opinion. The fascinating observation, however, that can be made from this graph is that the fundamentalists are clearly separated from the other traders and form a central cluster. This core of a densely connected cluster of fundamentalists that is connected to a periphery of chartist nodes is especially remarkable because this behaviour emerges naturally just through the given network structure. This sort of dynamics could be compared to the hierarchies in financial markets, where there are large institutional investors who tend to be more conservative than individual traders, who might be more speculative [74].

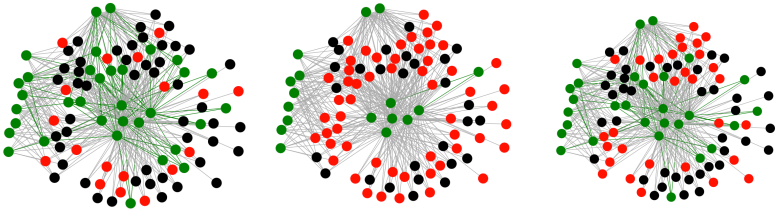


FIGURE 2.23: Opinions of agents during three different times in the simulation. The green fundamentalists do not change their opinions, while the optimists (red) and the pessimists (black) change often.

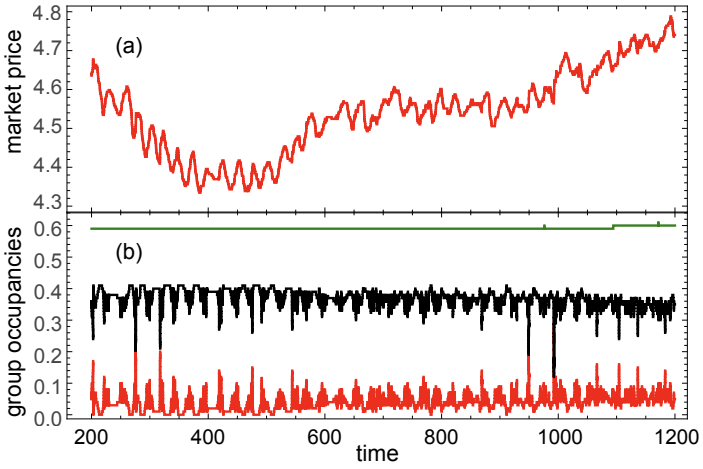


FIGURE 2.24: Trader distribution (b) and market price (a) for a simulation with strong community formation. Communities of fundamentalists and pessimists dominate.

An even more extreme version of community formation can be observed in Fig. 2.24 and Fig. 2.25. Similarly to the previous simulation, there is a core of fundamentalists who are constant in their opinion. However, in addition to that, there is also a clear dominance of pessimists over optimists. In this regard, the simulation differs significantly from the simulation depicted in Fig. 2.22, where optimists and pessimists trade blows. The effect of this is clearly visible in the price behaviour, as the price drops quckly and significantly and only rises back slowly, driven back to a much higher fundamental value by the fundamentalists.

A comparison of the network structures in Fig. 2.23 and Fig. 2.25 reveals a similarity concerning the fact that there is a densely connected group of fundamentalists at the core. However, this core is much less densely connected in Fig. 2.25, which might explain why such a robust pessimistic community could form.

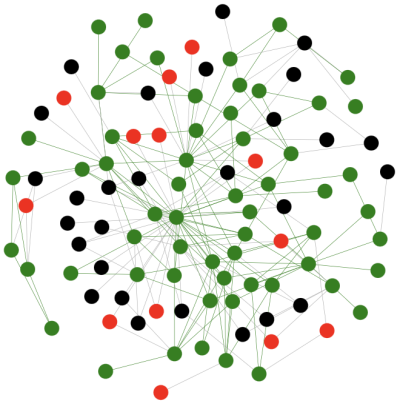


FIGURE 2.25: Network structure for the simulation in Fig. 2.24.

2.5.2 *Spontaneous rise of chartists*

One key dynamic of the Lux-Marchesi model is the spontaneous rise of chartists in a system usually dominated by fundamentalists [70]. Even though the results in the previous section seem to suggest different behaviour for the neuron model, it is indeed possible to observe such dynamics, given the proper network structure.

Snapshots of the model states are shown in Fig. 2.26. This figure shows how the information percolates through the system. In the first step, there are some optimistic traders, which then influence their connected traders and each other. This leads to a growing cluster of optimists in step two and three. The latter two snapshots show how the cluster of optimists shrinks again, and the fraction of fundamentalists and pessimists increases.

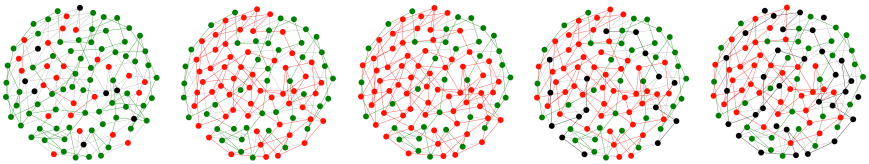


FIGURE 2.26: States of agents selected between time steps 535.5 and 536 in Fig. 2.27. Chronological order is from left to right.

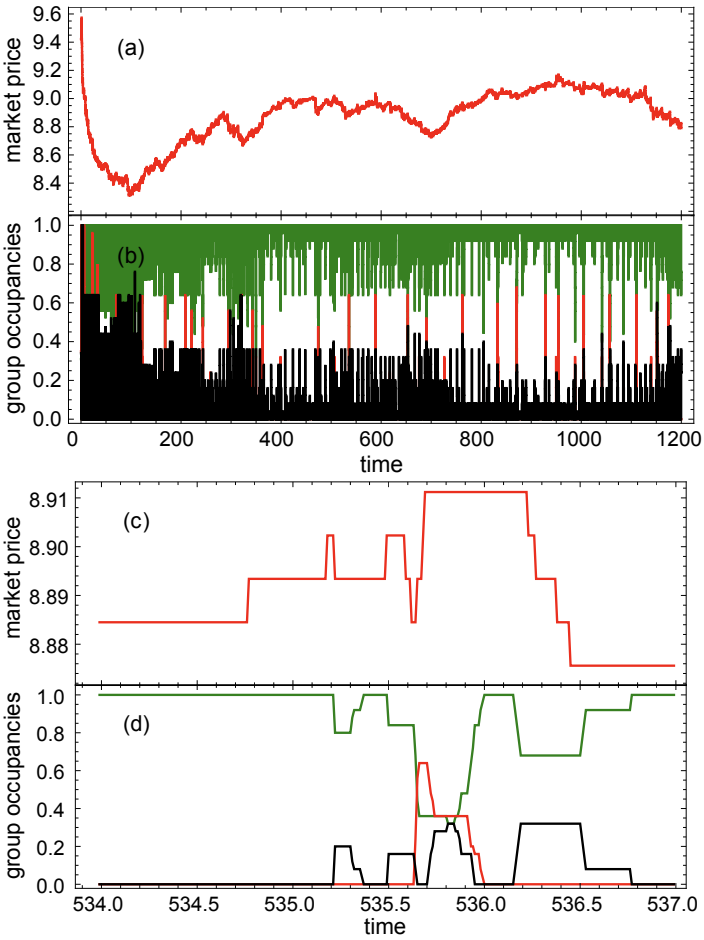


FIGURE 2.27: Simulation which shows the spontaneous rise of chartists. In (a) and (b), an overview of the simulation is given, and in (c) and (d), an event of rising chartists is depicted.

Figure 2.27 shows a simulation which exhibits just this behaviour. While the majority of traders usually have a fundamentalist opinion, there are periods of spontaneously rising chartists. There are regular occurrences of pessimistic opinion and larger sudden spikes of optimist opinions. This is particularly interesting because this system displays an asymmetry between optimists and pessimists, as was also the case in different form in Fig. 2.24. However, much more about this behaviour can be learned when looking at

the evolution of the states of individual traders during the spike shown at the bottom of Fig. 2.27.

A volatile version of the emergence of chartists is seen in Fig. 2.28. The fraction of fundamentalists never drops significantly below 50%, and most of the chartists, if any are present, are pessimists. Comparison with Fig. 2.27 reveals a much more volatile price movement.

Figure 2.29 shows one of the frequent spikes of pessimists. A fascinating aspect of the plateau at time step 579.8 is that there is a single cluster of all pessimists in the system, which is stable over 50-time steps before the model dissolves back into fundamentalist dominance.

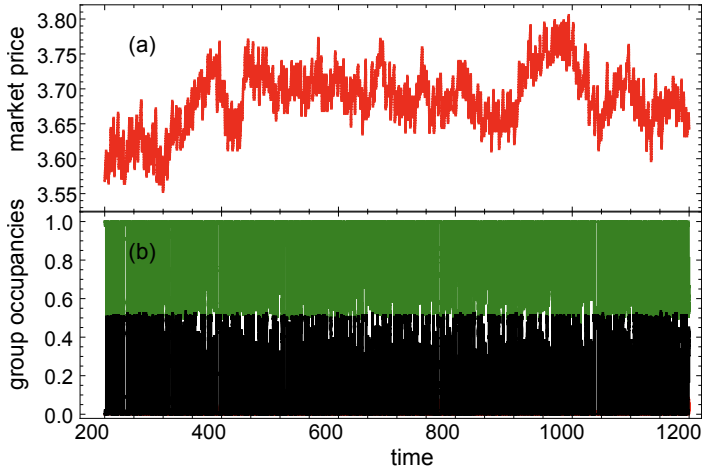


FIGURE 2.28: Trader distribution (b) and market price (a) for a simulation with many spikes of pessimistic chartist.

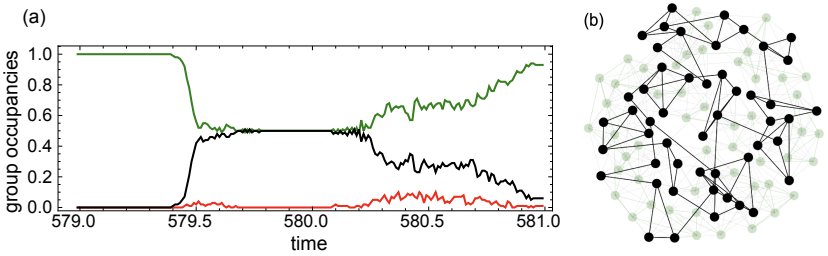


FIGURE 2.29: Spike of pessimists in Fig. 2.28. The model snapshot (b) was taken during the plateau depicted in (a) around time step 579.9, and the cluster of pessimists is highlighted.

2.5.3 *Extreme volatility clustering*

The model frequently produces volatility clustering, which refers to correlation of absolute returns in financial markets because large absolute returns tend to cluster together. Figure 2.30 shows an edge case of this behaviour. This effect leads to intervals with clustered smaller returns and intervals with clustered higher returns.

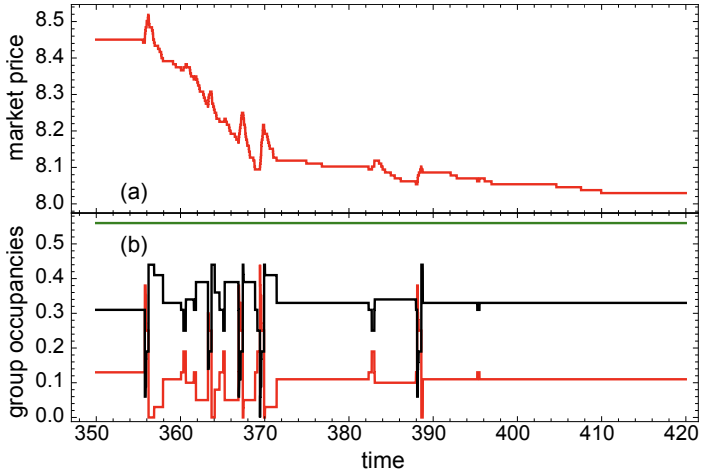


FIGURE 2.30: Spontaneous market movement after a period of constant price and opinion.

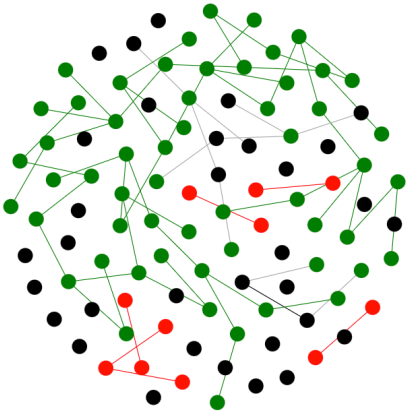


FIGURE 2.31: Network structure for the simulation in Fig. 2.30.

In this simulation, the model is often static over prolonged periods until randomly entering a state of increased activity, after which it returns to its rest state. This can be explained with traders fixed in their opinion without the arrival of new information, which then leads to a continuing exchange of information throughout the network of agents until a new configuration is found. This pattern represents the fact that there is no reason for agents to act without new information. In the absence of new information, traders will not act.

Analysis of the network structure in Fig. 2.31 provides a theory for this behaviour. The network is relatively sparsely connected, which reduces the amount of information that traders exchange. With a low amount of random information $\mathbf{I}_{\text{rand}}$, there is just not much information arriving at each trader in non-volatile time spans. Moreover, a significant portion of individual traders only follow market movement. They are active in areas of price volatility or when they receive enough random information to drive them to an opinion change, but they are not subject to crowd behaviour.

2.5.4 *Infectious Behavior*

Fascinating dynamical behaviour is the fixed cluster growth observed in Sec. 2.5.1. The simulation in Fig. 2.32 corresponds to the star graph in Fig. 2.32, where a single node is connected to all other nodes. This central node assumes a fundamentalist opinion and does not deviate. What is fascinating is that while all peripheral nodes are chartists at the beginning, some of them change to fundamentalists during the simulation run. This in itself is not noteworthy, as traders change their opinions all the time. What is remarkable, however, is the fact that, once fundamentalists, they do not change their opinion but instead form a cluster of fundamentalists. The left side of Fig. 2.32, shows that the fraction of fundamentalists increases over time without a single drop.

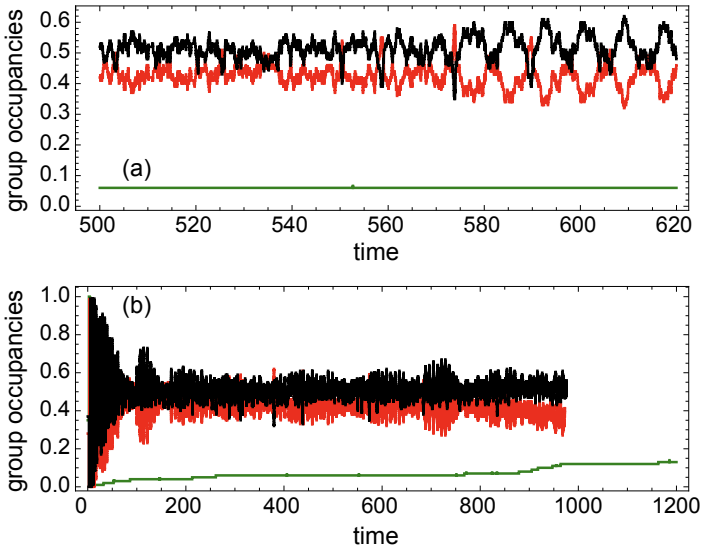


FIGURE 2.32: Trader distributions with high (a) and low (b) time resolution for sticky fundamentalists.

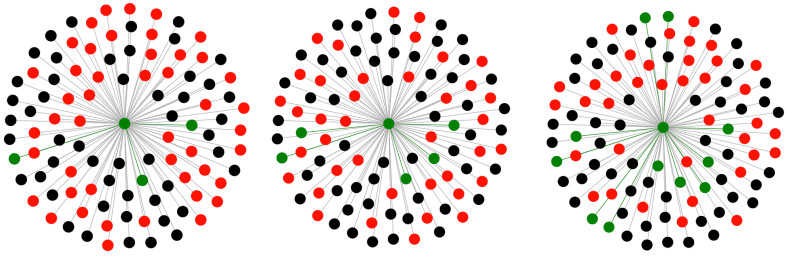


FIGURE 2.33: States of agents during different times of the simulation runs. Once a fundamentalist joins the cluster, it will not switch opinion.

This behaviour could be referred to as "sticky fundamentalists": a trader exposed to a fundamentalist opinion and connected to the network will not deviate from that opinion. While there is frequent switching between optimists and pessimists, the fundamentalists are unaffected. In Fig. 2.33, we can observe that not just the fraction of fundamentalists is constant, but that individual fundamentalists indeed do not change their opinion.

2.5.5 *Unifying opinion*

The last two examples are included not because of their great simulation of financial markets but because they show interesting, complex patterns. In the simulation depicted in Fig. 2.34, all traders have a unified opinion, as clearly visible on the right side. However, this opinion changes frequently and reacts to the movement of prices and the arrival of new information. In the hierarchical model [77], Sornette and Johansen say that a crash is characterized by all traders in a market imitating one another. The market movement here is quite contrary. Even though all traders act at once, the market price is still relatively stable. The system in itself stabilizes the price through frequent changes in opinion.

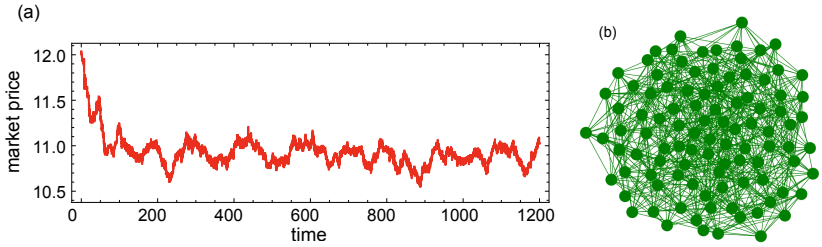


FIGURE 2.34: (a): market price for a simulation where all traders share the same opinion. (b): Illustration of such a network.

This change in opinion is due to the currents $\mathbf{I}_{\text{price}}$, which each trader receives. The network is connected strongly enough that all traders will join one majority opinion. However, this connection is not enough to detach the traders from reality, the price movements. This behaviour could be adjusted by changing the scaling of the market information.

2.5.6 *Sudden shifts in behavior*

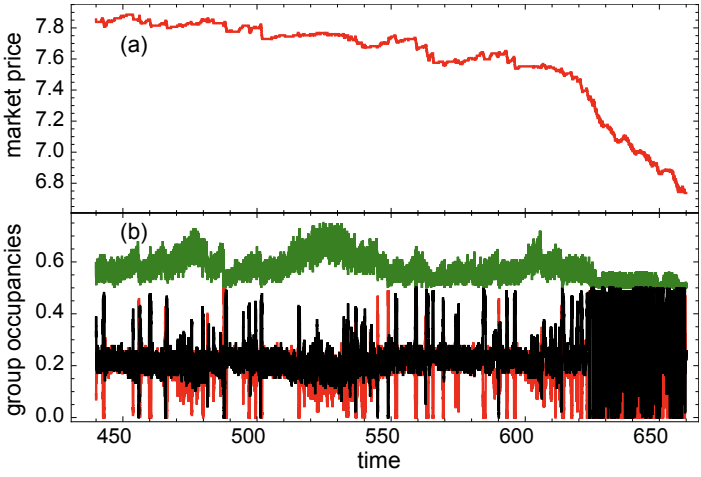


FIGURE 2.35: This simulation run depicts market price (a) and trader distribution (b) during the sudden change in behavior around time step 625.

Most of the simulations follow the same patterns throughout a run. Extreme volatility clustering or sticky fundamentalists may have a disrupting effect on the market dynamics, but they do not change the system behaviour in a way that it is changed in Fig. 2.36. At the time step 625, there is a change not just in price movement but in trader distribution as well. After a period of relative stability, the market price enters a downward spiral, with much-increased volatility. Similarly to that, the volatility of trader distribution changes drastically as well. However, the changes are more profound. While the volatility of chartists increases drastically, the volatility of fundamentalists decreases. The apparent increase in activity by the chartists is accompanied by a cluster of fundamentalists whose opinion is more stabilized.

This critical behaviour is different from the one described in Sec. 2.3.5. While it was possible to present a critical point with regard to entropy and energy there, the criticality appeared when changing the connectivity. This simulation, however, generates such a change without a change in parameters.

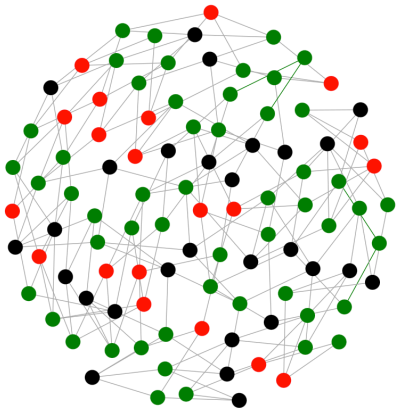


FIGURE 2.36: Network structure for the simulation run in Fig. 2.35.

2.6 CONCLUSION

In this chapter, we introduced and explored a first version of an agent-based model that integrates a neuronal-level decision-making mechanism. Building upon the Lux-Marchesi framework, we replaced mean-field opinion dynamics with a more biologically inspired integrate-and-fire process that controls agents' transitions between fundamentalist, optimist, and pessimist states.

We started by adapting the original one-dimensional integrate-and-fire model of neuronal behaviour to a three-dimensional setup that can capture the diverse opinions of traders. We then showed how these neuron-like agents interact within a network, exchanging information through "currents" that, once integrated to a threshold, trigger a "firing event" and hence a change in opinion. This approach ensures that decision-making occurs intermittently and endogenously, influenced both by the incoming information (from markets, social interactions, or random noise) and by the network structure that underpins agent connectivity.

Next, we analyzed the influence of parameters—such as the weight of market and network information and the decay rate of trader potentials—and noted how they can dramatically change the dynamics. Depending on how strongly agents are influenced by market signals or their peers, we observed that the model can generate a broad range of stylized financial phenomena: from volatility clustering, fat-tailed return distributions, and periods of stability interspersed with sudden bursts of activity, to the spontaneous formation and dissolution of opinion clusters.

We found that network structure is at least as important as agent-level parameters. Varying network topology and connectivity can push the system toward critical points, beyond which collective behaviour changes radically. Dense connections can cause global consensus or orderly clustering around the fundamental value, while sparser or more uneven structures might produce volatility spikes, persistent communities with fixed opinions, or even rapid oscillations of collective sentiment.

To gain deeper insights, we introduced measures such as current flow closeness centrality, cluster size distributions, and energy. These measures allowed us to quantify how traders tend to organize into communities and

how global patterns like stable clusters of fundamentalists or spontaneous waves of chartists emerge from local interactions. We also highlighted the challenge of meaningful model calibration: the multidimensionality of parameters and structural degrees of freedom makes a systematic parameter search and fit to real market data nontrivial. The use of advanced tools, such as machine learning techniques, present a promising path toward efficient parameter estimation and model validation.

In conclusion, the results presented in this chapter underscore the richness and complexity that emerge from integrating neuron-like decision dynamics into an agent-based economic model. Although this first model is still a simplified demonstration, it reveals the potential of the approach. Future chapters will build on these insights, refining the methods, introducing inhibiton and hierarchical structures, and integrating our findings into more sophisticated frameworks that come closer to real-world financial systems.

3

INHIBITORY EFFECTS - IRRATIONALITY DRIVES BUBBLES

In the last chapter, we combined the integrate-and-fire model with the Lux-Marchesi model and arrived at a working model that can replicate many of the stylized facts found in financial markets and beyond. We investigated how the model behaviour changes with different network structures and how parameter choices can affect model behaviour. Now that we have a fundamental understanding of how the integrate-and-fire trader model works, the natural next step is to use it to get insights into some of the most interesting questions in financial markets. The integrate-and-fire trader model has information at the core of its mechanics, and the agent behaviour is exclusively driven by the information that they receive. What differentiates the model from many further approaches is the explicit focus on information, which allows for some unique analysis. In this chapter, we will focus on a key aspect of trader behaviour: Inhibitory behaviour that can manifest itself as irrationality in markets.

3.1 INTRODUCTION

The insight that human behaviour deviates from perfect rationality has informed both agent-based models and broader economic modelling approaches for decades. As discussed in Chap. 1, behavioural and neuroeconomic perspectives have shown that decision-making is shaped by psychological factors, cognitive biases, and underlying neural mechanisms [111, 135]. In particular, it has become clear that processes in the brain influence how we aggregate and evaluate information, and these insights have begun to shape economic theories and models of choice under uncertainty.

In Chap. 1, we gave an introduction to some of the key motivators for behavioural- and neuro-economics. In their 2001 paper, McClelland and Usher [118] focus heavily on the importance of inhibitory effects when making decisions between multiple options. They contrast two different kinds of inhibition: feed-forward and lateral inhibition.

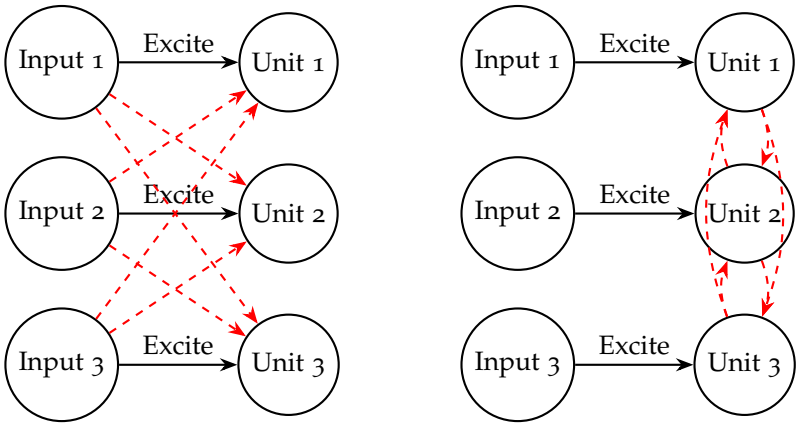


FIGURE 3.1: Comparison of feed-forward (left) and lateral inhibition (right) mechanisms. Feed-forward: Units receive excitory (black) and inhibitory (red) currents from the input level. Lateral inhibiton: while units receive excitory units from the input level, inhbition just depends on the activation of other units.

Figure 3.1 illustrates the difference between the two: Feed-forward inhibition describes a process where each input into the neuronal network directly excites some neurons while inhibiting others, the inhibition at any given moment only depends on the inputs at the same time. Lateral inhibition, on the other hand, occurs between competing units themselves, such that the activation potential of one neuron will directly lower the action potential of other neurons. This means that the inhibition of inputs at time t depends on the action potentials and, thus, on the previous inputs. McClelland and Usher strongly argue for lateral inhibition: This process induces natural competition between different options: As the potential for one option builds up, it simultaneously reduces the activation for all of the other options, thus theoretically ensuring that the option that is most supported by the gathered information comes out on top. This induces a dynamic form of competition for "attention" or "dominance" among options, as the accumulated evidence for one choice option partially determines the processing of alternatives [118].

The preference for lateral inhibition in models of decision-making is supported by extensive evidence in neuroscience and psychology. For instance, lateral inhibition has been implicated in early sensory processing in the visual cortex [136], as well as higher-level cognitive tasks where multiple alternatives compete for selection [137, 138]. In such models, lateral inhibition encourages the system to converge on the choice option best supported by the available evidence, reducing the likelihood of indecision and improving the overall signal-to-noise ratio during choice tasks. In the leaky competing accumulator model [118], leakage and lateral inhibition together ensure that random fluctuations do not accumulate indefinitely, which would otherwise degrade decision quality. It essentially improves the signal-to-noise ratio when making decisions. It is thus a great model to make well-informed decisions.

While lateral inhibition is seen as a benefit in the leaky accumulator model, it might not necessarily be the most efficient way to process information in financial markets. Financial decision-making often occurs under conditions of rapidly changing information, uncertain forecasts, and strategic interactions among heterogeneous agents [139]. In such environments, the strong filtering effect provided by lateral inhibition can have unintended consequences, potentially introducing systematic biases into the decision-making process.

Two well-known effects highlight how biases can factor into such decision-making mechanisms: recency bias and primacy bias. Recency bias refers to the tendency of agents to overweight more recent information. Within the leaky competing accumulator framework, this bias arises naturally from the leakage (decay) parameter τ, which causes older information to become less influential over time. As a result, sudden price movements or fresh pieces of news can disproportionately shape expectations, even if older, more fundamental signals would otherwise counsel caution. Studies in psychology on belief-updating have repeatedly shown such recency effects [140].

Primacy bias, on the other hand, emerges because of inhibition. When an agent first receives a strong piece of information favoring a particular choice, subsequent competing information finds it harder to influence the decision process, since the initially dominant option continuously suppresses its rivals. This means that "early movers" in an agent's information set gain a persistent advantage. In financial contexts, this phenomenon

resembles well-known anchoring effects, where initial estimates or benchmarks disproportionately influence subsequent judgments and trades [141]. Markets characterized by strong early signals - perhaps driven by charismatic opinion leaders, sudden news flashes, or initial jumps in asset prices - can thus become "locked in" to certain states that do not necessarily reflect the most updated fundamentals.

These biases are particularly important in finance, because traders are constantly subjected to new information. Yesterdays news might be irrelevant today, and relying on outdated signals can prevent optimal responses to market changes. These factors might contribute to well-known market anomalies like speculative bubbles and fat-tailed return distributions [142].

Considering the significant effect that irrationality can have in financial markets, we aim to extend the previously introduced model by incorporating inhibitory effects. By doing so, we can investigate how neuronal-level decision processes affect agent behaviour in complex financial environments. More specifically, we will explore how lateral inhibition and the interplay between leakage and inhibition can lead to emergent patterns of biased decision-making at the market level, with consequences for price dynamics, volatility, and other stylized facts.

In the following sections, we will introduce inhibitory effects into the integrate-and-fire trader model and examine their effects. This approach allows us to better understand the role of neuronal-level processes in shaping economic outcomes, bridging the gap between micro-level cognitive architectures and macro-level market phenomena.

3.2 UPDATED MODEL

Thankfully, adding inhibition to the integrate-and-fire trader model requires only minor adjustments. For recollection, in Sec. 1.7 we explained the leaky competing accumulator, where the action potential u_i is governed by the following equation:

$$\begin{aligned} du_i &= [I_i - \lambda u_i - \beta \sum_{j \neq i} u_j] \frac{dt}{\tau} + \xi_i \sqrt{\frac{dt}{\tau}} \\ u_i &\to \max(u_i, 0). \end{aligned} \tag{1.24}$$

Parameter	Value	Parameter	Value
Number of Agents	500	τ	0.98
ω_1	0.15	θ	1
ω_2	0.15	t_c	1
ω_3	0.01	β	10
y	1	μ	0
		Network	Barabási-Albert with $m = 2$, $n = 500$

TABLE 3.1: Simulation parameters

in which β determines the strength of the inhibitory effect, while $\frac{\lambda}{\tau}$ determines the decay of the action potential over time. In the previous section, we defined the action potential of the integrate and fire model without inhibitory effects:

$$\tau_i \frac{du_i}{dt} = -\big[u_i(t) - u_{i,r}\big] + R_i\, I_i(t), \tag{2.7}$$

where $i = 1, 2, 3$. Now, to add inhibitory effects, we simply add the inhibitory term from Eq. (1.24) to this equation and arrive at

$$\tau_i \frac{du_i}{dt} = -\big[u_i(t) - u_{i,r} - \beta \sum_{j \neq i} u_j\big] + R_i\, I_i(t), \tag{3.1}$$

where we keep β as the variable indicating interaction strength. The rest of the model setup is kept unchanged. For a complete description of the previous model refer to Chap. 2.

3.3 SIMULATION SETUP AND RESULTS

To set up the simulation, we decide on a set of parameters used in the last chapter that produced a stable market with a realistic return distribution (see Fig. 2.13). We decide to use a Barabási-Albert-type preferential attachment network because it best resembles the information structure in financial markets. As described in Sec. 1.5.4, many social networks, including those underlying financial trading, exhibit scale-free characteristics – meaning that a small number of well-connected "hub" nodes have disproportionately many connections, while the majority of nodes remain relatively

sparsely connected [93]. As we have seen in Chap. 2, a Barabási-Albert type network with parameter $m = 2$ sets the model up in a state where it is susceptible to parameter changes and not entirely dominated by the network structure. A parameter of $m = 2$ is also consistent with research on financial networks, which usually puts m between 2 and 3 [20, 97]. An example of the networks used is depicted in Fig. 3.2.

We now let the inhibition parameter range from $\beta = 0$ to $\beta = 0.2$ and investigate changes in behaviour. First, we will look at the time series of group occupancies and price behaviour to identify observable changes. After that, we compute a variety of measures that have proven relevant in Chap. 2 to gain deeper insights. We introduce a new entropy measure, which gives even more interesting insights. We will also look at the role of the action potential, which turns out to be a key factor.

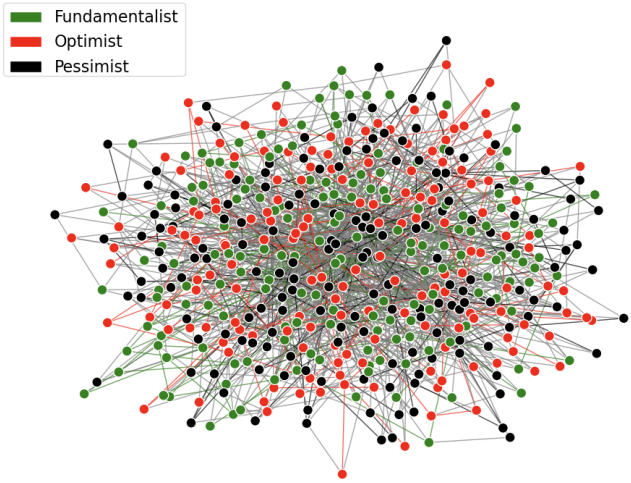


FIGURE 3.2: Example network structure used for simulations. The network is generated using a preferential attachment process with $N = 500$ and $m = 2$.

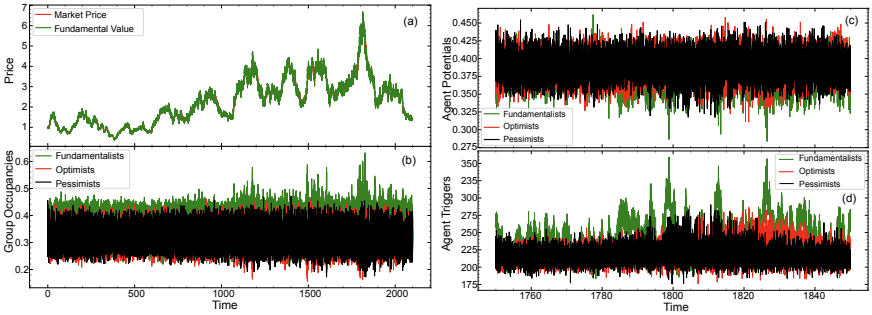


FIGURE 3.3: Price behaviour (a), group occupancies (b), agent potentials (c) and agent triggers (d) for $\beta = 0$.

The behaviour of the model without any inhibitory effects is depicted in Fig. 3.3. In the absence of inhibitory effects, i.e. $\beta = 0$, the model defaults to a state where optimists, pessimists and fundamentalists have an average occupation of one-third each, given the chosen parameters. The market price tends to follow the fundamental value.

We now incrementally apply inhibitory effects to the model until we observe a shift in behaviour. At $\beta = 0.036$, we can observe the first shift in behaviour, as depicted in Fig. 3.4.

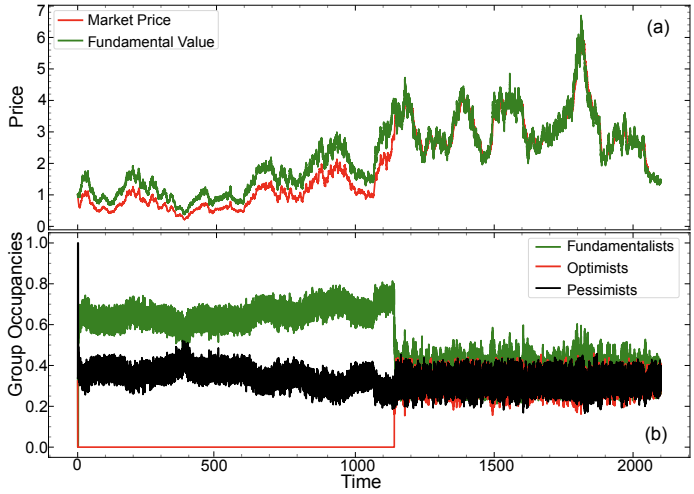


FIGURE 3.4: Group occupancies (b) and price behaviour (a) for $\beta = 0.036$.

The shift is rather dramatic and resembles a phase transition in statistical mechanics. Upon the start of the simulation, the market immediately enters a state in which all optimism is suppressed, and only fundamentalists and pessimists control the market behaviour. This is reflected in the market price, which displays a constant undervaluing when compared to the fundamental value. Indeed, the relative fractions of the market and fundamental value suggest that it is a constant struggle between fundamentalists and pessimists, with the fundamentalists preventing the market price from deviating too much. Around timestep 1150 however, the behaviour shifts dramatically into a regime similar to the no-inhibition model. This change is spontaneous and the group occupancies do not show any precursors of the transition.

The effects of introducing inhibition into the model are in line with what we expected earlier - A very early bias that the traders might have is "locked in" due to the inhibition, leading to a constantly warped perception of the true value of the asset. It takes the agents a long time to overcome this initial bias; the shift happens spontaneously and results in a drastic shift in market price. This becomes even more evident when we continue to increase β to 0.092.

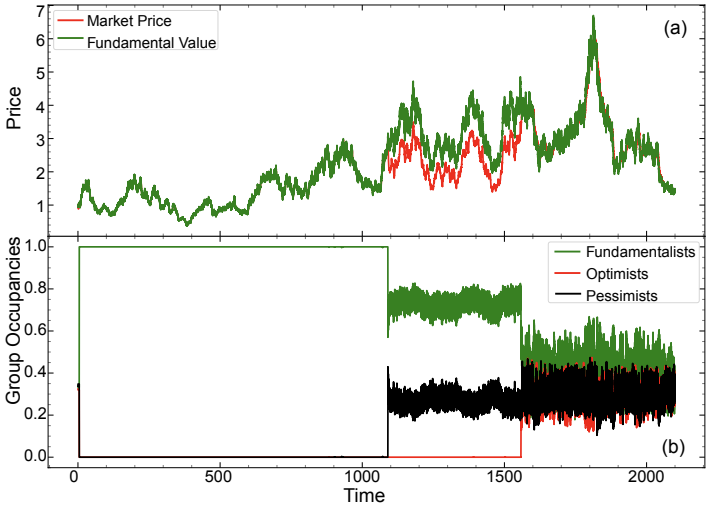


FIGURE 3.5: Group occupancies (b) and price behaviour (a) for $\beta = 0.092$.

Figure 3.5 reveals a total collapse of trader opinion towards the fundamentalist state. This means that the system self-organises into a state without any speculation, with all traders just following the asset value. This situation is stable until it breaks around timestep 1100. Now we are in the regime we started with for $\beta = 0.036$. After around 500 timesteps, the system undergoes another transition into the pre-inhibition phase. So for $\beta = 0.092$, there are three cases: Full collapse, partial collapse and approximate equal distribution. It should be noted that this is also reflected in the price behaviour: The first phase does not exhibit significant market-price deviations, while the second phase does. The third phase again shows the market price fluctuating around the fundamental value.

So far, the simulations have shown a tendency to relax the influence of the inhibitory effects over time, going from complete collapse to one state towards the non-inhibitory state of approximately equal state population. However, this pattern breaks down when increasing the inhibitory effects to $\beta = 0.144$. For this scenario, as can be seen in Fig. 3.6, the simulation does not collapse into the non-inhibiton state, but instead goes in and out of periods of speculation, first dominated by pessimists and then by optimists, resulting in undervaluation and overvaluation respectively. Also noteworthy is the timespan of market turmoil from timesteps 1800 to 2050, in which the group occupations fluctuate randomly.

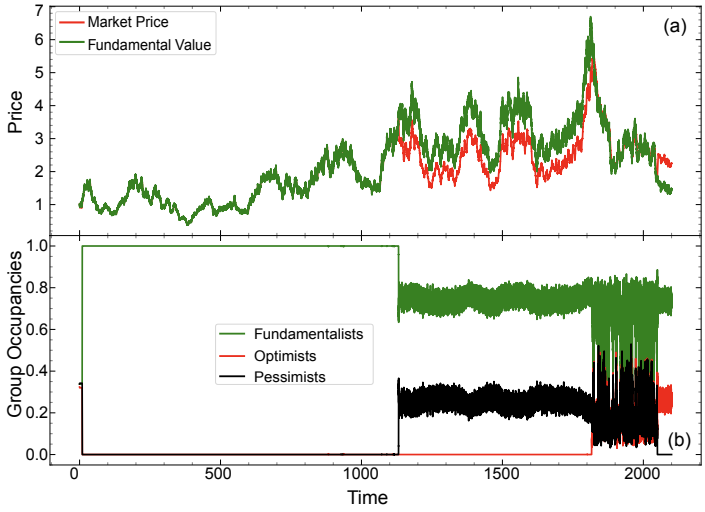


FIGURE 3.6: Group occupancies (b) and price behaviour (a) for $\beta = 0.144$.

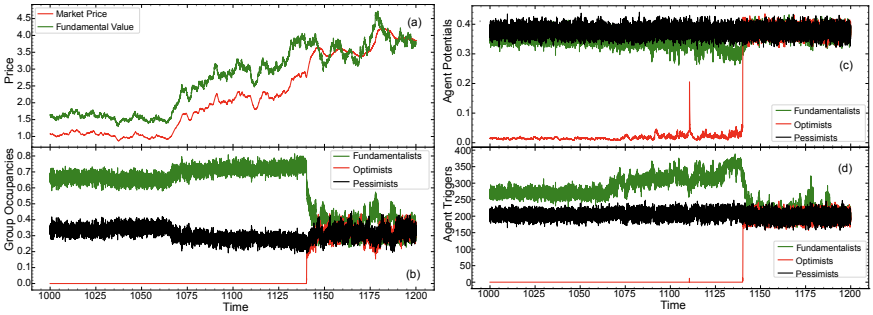


FIGURE 3.7: Left: zoomed in group occupancies (b) and price behaviour (a) for $\beta = 0.036$. Right: agent potentials (c) and triggers (d) over the same period.

3.4 AGENT POTENTIALS: PRECURSORS FOR MARKET PHASE TRANSITIONS

We will now take a closer look at the different phases induced by inhibition and especially at the transitions between the different phases. We start by examining the transition found for $\beta = 0.036$. Figure 3.7 depicts the transition between a period of undervaluation dominated by pessimists and fundamentalists towards the non-inhibitory regime at $t = 1141$. In this regime, the market price fluctuates around the fundamental value. By just looking at the price behaviour and the group occupancies in Fig. 3.7a and 3.7b, these transitions appears to be both instantaneous and entirely spontaneous. This is also confirmed by looking at the triggers in Fig. 3.7d, which measure the information that agents share with each other.

However, a look at the action potential in Fig. 3.7c tells a different story. In the build-up to the phase transition, there is obvious movement in the optimist potential. While it is almost zero (with negligible fluctuations) until timestep 1075, fluctuations begin to increase noticeably after that, resulting in a sudden build-up in trader potential around the transition at timestep 1140. This increase in fluctuation before a transition is a common feature we find before transitions. It highlights one of the key strengths of the model: by focusing on information transfer and processing as the key driver of decision-making, we can gain an "x-ray" vision of the decision processes and the current mood of traders. All of the analysis of the features that can be seen by outside observers does not allow for any indication that a

transition might be closeby. However, when we look at the internal action potentials of traders, these signs become obvious.

This story continues when we analyse Fig. 3.8, which depicts the scenario for $\beta = 0.092$. Here, we can observe two transitions: First, the emergence of pessimism around timestep 1095 and then the emergence of optimism around timestep 1560, which brings the system back to a non-inhibitory state.

Again, the observable features of the model do not indicate that a transition is approaching. The emergence of pessimists around $t = 1095$ is accompanied by a sharp drop in price, which resembles a crash caused by a sudden mood shift in a market. Similarly, the realignment of market and fundamental values happens seemingly spontaneously. However, once again, a look at the action potentials allows us to observe clear precursors in the form of increasing fluctuations leading up to the transitions.

We now look at a last scenario, $\beta = 0.144$. In this simulation, two main features require close attention: The first is the period of market turmoil between timesteps 1800 and 2050, and the second one is the sudden emergence of a bubble at $t = 2050$. The build-up to the first transition looks very similar to previous transitions: we observe the spontaneous emergence of optimists in a simulation which was previously dominated by pessimists and fundamentalists. However, the post-transition phase we observe is strikingly different from the cases with lower inhibition.

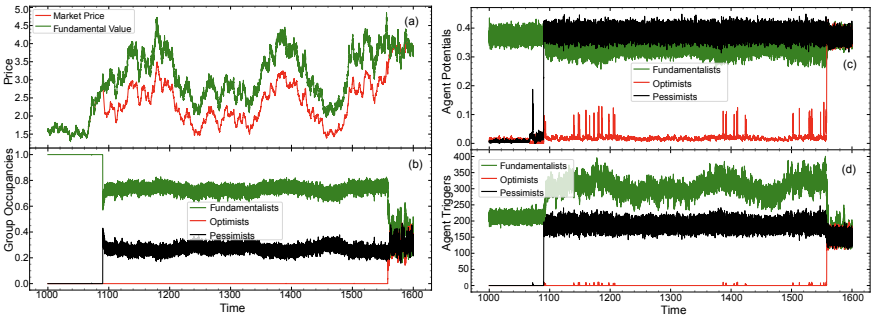


FIGURE 3.8: Left: zoomed in group occupancies (b) and price behaviour (a) for $\beta = 0.092$. Right: agent potentials (c) and triggers (d) over the same period.

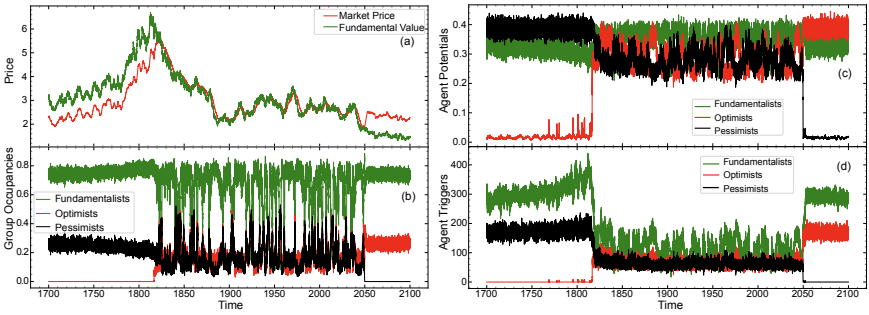


FIGURE 3.9: Left: zoomed in group occupancies and price behaviour for $\beta = 0.144$. Right: agent potentials and triggers over the same period.

Instead of a transition to the no-inhibition state (as observed for $\beta < 0.144$), the fundamentalist dominance is sticky. While we see frequent spikes of speculation throughout the period of "market turmoil", the system always reverts back to fundamentalist dominance. It is not surprising that this would be the case since inhibitory effects lead to primacy bias. When the inhibitory effect is strong enough, a certain fraction of agents will not change their opinion.

There are a few important things to notice when observing Fig. 3.9: First, the agent potentials are all in the same range of about 0.3 to 0.4, while the range of possible values is [0,1) and optimist and pessimist potentials oscillate thus the majority chartist opinion switches frequently. However, at the same time, the triggers, i.e. the information passed between agents, are much smaller than they were in any earlier scenario, meaning that agents fire less. This can be explained by the strong inhibitory effects, leaving a large fraction of fundamentalists "locked in" and refusing to switch their opinion.

Figure 3.10 shows the emergence of a bubble following a period of turmoil. As described previously, the agent potentials of optimists and pessimists oscillate during the turmoil phase. At the point of the transition, the optimistic potential is slightly higher than the pessimist potential, which might explain why the system collapses into a bubble state instead of an undervaluation state.

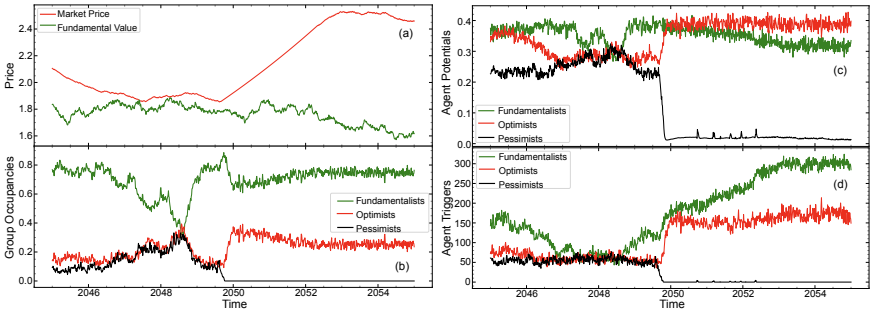


FIGURE 3.10: Left: zoomed in group occupancies and price behaviour for $\beta = 0.144$. Right: agent potentials and triggers over the same period.

We have seen quite dramatic changes in system behaviour, including sudden shifts in trader occupation, price behaviour and action potentials. However, we have only illustrated them by directly describing simulation behaviour. In the next section, we will introduce a measure that will allow us to better quantify some of the observed changes: entropy.

3.5 ENTROPY

The concept of entropy is a measure of the order (or disorder) of a system that is used in many different contexts. In our case, it is useful as a measure to quantify the complexity arising from the interactions between agents. Since we are looking at the entropy in the agent states, which is an entropy of information, one appropriate definition is the Shannon entropy [143]:

$$S_i = -\sum_i p(s_i) \ln(p(s_i)), \tag{3.2}$$

where $p(s_i)$ is the probability of state s_i. The large number of possible states s_i the system can have (for N agents there are 3^N possible states) means it is computationally infeasible to directly estimate the entropy. Therefor we estimate the entropy using the Bethe approximation [144].

3.5.1 *Bethe approximation*

The Bethe entropy originates from the Bethe-Peierls approximation, introduced by Hans Bethe in 1935 to study the statistical mechanics of spin systems on lattices [145]. This approximation was developed to analyze

models like the Ising model, where particles (spins) interact only with their nearest neighbors. The key idea was to simplify the complex interactions in a lattice by considering local clusters of nodes and neglecting long-range correlations. Because of this, the Bethe approximation is particularly useful for systems with local interactions and is thus a tractable way to estimate entropy for a network of agents.

While originally developed for networks without self-loops, for which the approximation is exact, it still provides useful results for networks with loops. In the Bethe approximation, the system is divided into subgroups. Entropy is calculated for each subgroup, then adjusted for over-counting. In our case, we look at pairs of nodes, in which case the entropy is given as

$$S_{\text{Bethe}} = \sum_{(n,m)\in E} S_{nm} - \sum_{n\in\nu}(1-d_n)S_n, \tag{3.3}$$

where S_n is the entropy of node n, S_{nm} is the joint entropy of the connected pair of nodes n and m and ν and E are the set of nodes and edges in the network, respectively. The degree of each node is given by d_i. The joint entropy of nodes n and m is:

$$S_{nm} = -\sum_{s_i,s_j} P_{ij}(s_i,s_j)\log P_{ij}(s_i,s_j), \tag{3.4}$$

where $P_{ij}(s_i,s_j)$ is the joint probability that node n is in state s_i and node m is in state s_j. We first convert the large data of individual states to pair states. As for each pair of nodes there are only $3^2 = 9$ possible states, this calculation can be done directly. After having a list of states for each node and edge, we approximate the probability with the frequency of states, which then allows for entropy estimation using the Bethe approximation. We do this over windows of timesteps to calculate a moving average of the entropy over time. This allows us to contrast the entropy of the system with other measures, like the deviation of market price and fundamental value. Figure 3.11, which shows the rolling average of entropy of a system without inhibition in relation to the market price deviations, reveals a remarkable trend that the entropy of the systems decreases when the market deviates from the fundamental value. This can be attributed to the fact that for speculation to occur, organisation and alignment of opinion is required among agents. Thus, a natural random selection of opinions, which maximises entropy, keeps the market price close to the fundamental value.

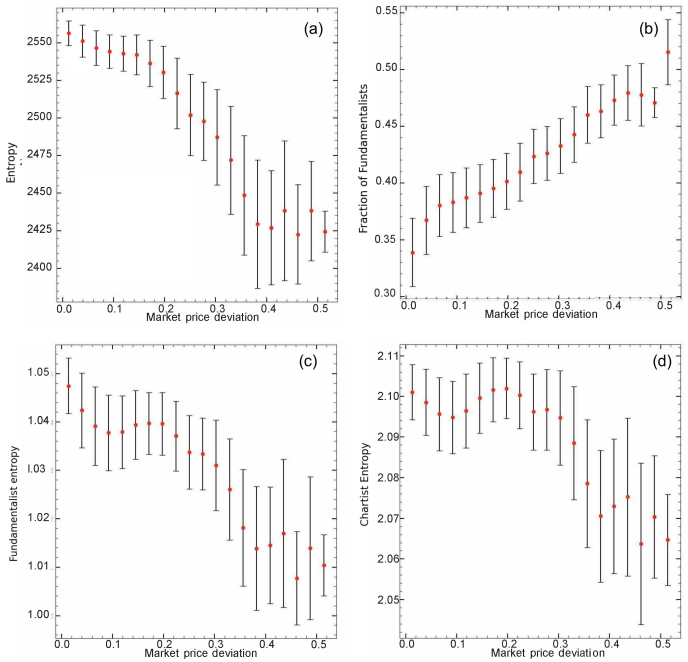


FIGURE 3.11: The entropy of a system is determined by the deviations of fundamental and market price. (a) The mean entropy per trader, (b) fraction of fundamentalists, (c) the fundamentalist entropy, and (d) the chartist entropy.

This behaviour may also explain the mean reverting features often observed for price movements. What is even more interesting is that a similar pattern is observed for the fraction of fundamentalists in the system (Fig. 3.11b). Similar to entropy, the fraction of fundamentalists increases in areas of higher speculation. This can be understood from the fact that deviations between fundamental and market value drive fundamentalists. However, more work is needed with regard to causality: Do the fundamentalists drive the entropy, or do they react to it? By calculating the entropy per trader group, we can answer the question of whether fundamentalists or chartists drive entropy. However, traders switch their opinion quite often, while a certain amount of time is necessary to arrive at good entropy estimations. Care is needed in attributing entropy to each trader group.

We calculate a rolling window of states for each trader so that the window used for entropy calculations exactly overlaps with the window used to assign the majority state for each trader. The results from this calculation are shown in Fig. 3.11. Here, it becomes apparent that the chartists are the main drivers of entropy in the system. While the number of fundamentalists increases with price deviations, their contribution to the entropy is rather small, while the entropy generated by chartists is almost twice that of fundamentalists. This means that while there are fewer chartists present when there is substantial speculation (when the market price deviates most from the fundamental value), these remaining chartists are still responsible for most of the disorder in the system.

3.5.2 *Entropy and inhibition*

So far, we have limited the analysis of entropy in the system to the non-inhibitory case. However, when we repeat the same analysis for the cases of increased inhibition, truly fascinating features emerge.

Figure 3.12 displays the relationship of entropy and market price deviations for different values of the inhibition parameter β. There are a few important observations to make here: Firstly, introducing inhibitory effects to the system lowers the overall entropy, with the cases of no inhibition ($\beta = 0$) and high inhibition $\beta = 0.144$ encapsulating the entropy curves. The second, and even more important discovery is that the shape of the curves changes with increasing inhibition.

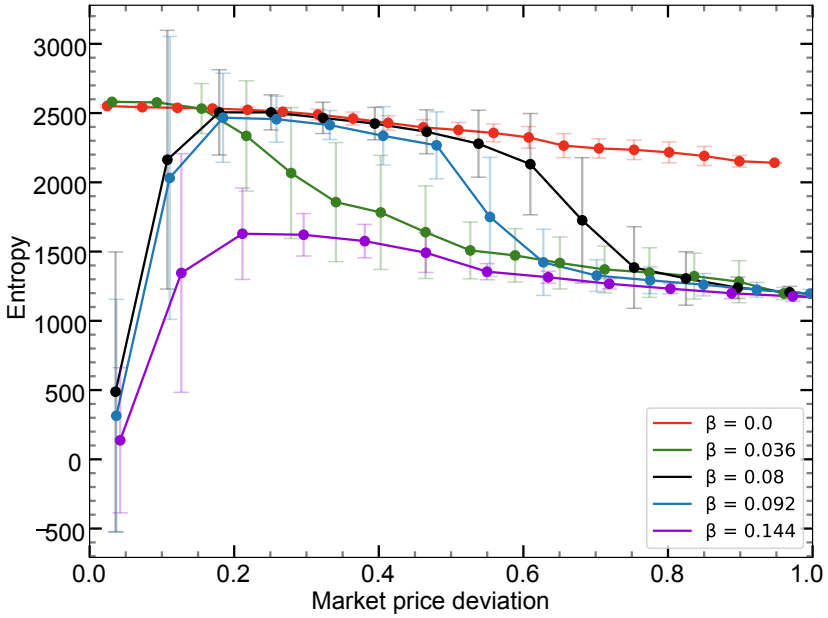


FIGURE 3.12: Entropy as a function of market price deviations for different values of inhibition β.

When we try to describe the system through a thermodynamic lens, we would expect it to assemble in a state that maximises entropy, given enough time. When we try to reconcile this with the notion of economic equilibrium, we expect to find the economic equilibrium at maximum entropy. For both $\beta = 0$ and $\beta = 0.036$, maximum entropy is observed when there is no market price deviation. This is no surprise and corresponds to classical economic literature, in which market and fundamental value should be identical in economic equilibrium [40].

However, upon further increasing the inhibitory effect, we can observe a clear peak of entropy at a market price deviation of around 0.2. This peak resembles a fundamental shift of system behaviour: We now observe a shift in economic equilibrium. This means that certain deviations of market price from the fundamental value are the natural state of the market. This finding is a clear indicator that irrationality, here implemented as inhibition, drives bubbles in financial markets. It is the inherent reluctance of humans

to discard old information, that leads us to cultivate perceptions of the markets that are divorced from reality. Since inhibitory effects are a key factor in our decision-making on a neuronal level, the takeaway is clear: We clearly just can't help ourselves.

3.6 RETURN DISTRIBUTIONS AND OTHER MEASURES

Now that we have established that inhibition leads to fundamental changes in market behaviour, we will analyse some of the other measures to investigate how the changes manifest in quantifiable market behaviour. We start with the most used measure, the return distribution. We observe log returns of the market prices p_t defined as

$$r_t = log(p_t) - log(p_{t-1}). \tag{3.5}$$

Log returns are an established way of analysing price behaviour and are a key factor when determining market risk. If our previous analysis is correct and inhibition does indeed drive the formation of bubbles, we would expect the return distribution to change when we introduce inhibition. Figure 3.13 confirms that this is the case. The figure depicts return distributions for β ranging from 0 to 0.16. It reveals a striking pattern: Without any inhibition present, the returns are normally distributed. When we start increasing β, the tails of the distribution become fatter, once again confirming that extreme market events increase with increased inhibition and thus irrationality in the model. This trend continues until we observe the distribution reverting to a normal distribution at $\beta = 0.16$. This is due to the fact that with strong enough inhibition, the model will collapse to a scenario without state changes, in which the market price behaviour eventually follows the fundamental value.

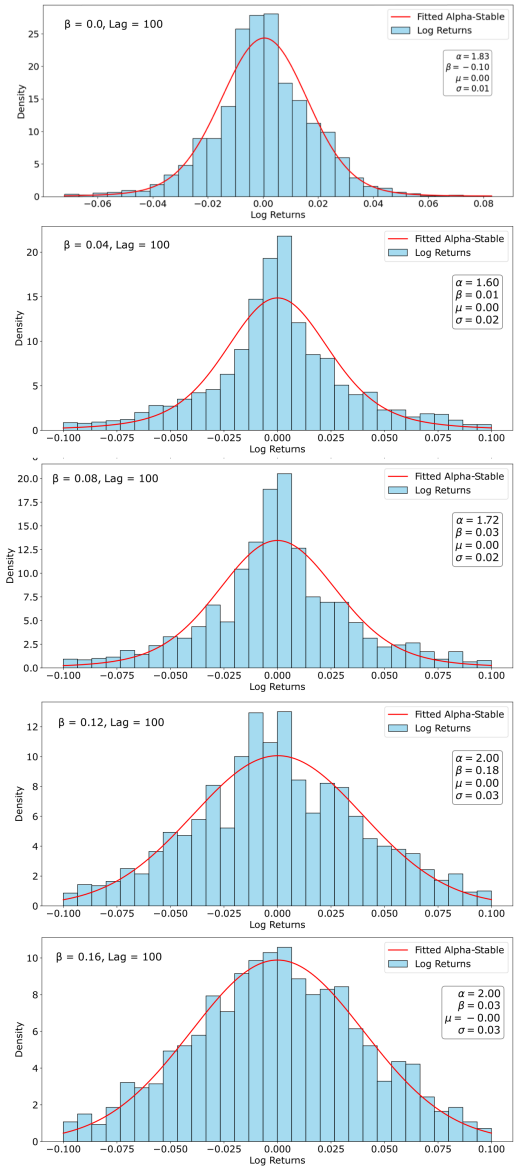


FIGURE 3.13: Return distributions over one time step (100 simulation steps) for different values of β.

3.6.1 *Average potentials and trader distributions*

We have looked extensively at a qualitative picture of both trader distributions and potentials for individual simulations and discovered that model behaviour can shift drastically with small parameter changes. To confirm this, we ran around 100000 simulations with different sets of parameters. Figure 3.14 indeed confirms the patterns discussed earlier, with the average potential showing rapid changes when the model transitions between regimes.

The same pattern can be observed in Fig. 3.15, where the average distributions of traders are depicted. It becomes clear that in the non-inhibitory regime, traders are approximately equally divided between opinions and that with increasing inhibition, we see rapid transitions towards a bifurcared market opinion. It should be noted that both Fig. 3.14 and Fig. 3.15 were created using the same network structure for all of the simulations to ensure consistency. To get a more adequate picture of the relationship between decay and inhibition, we performed a grid search of both parameters, where we also varied the network structure to ensure the general validity of the results. The corresponding phase maps can be seen in Fig. 3.16.

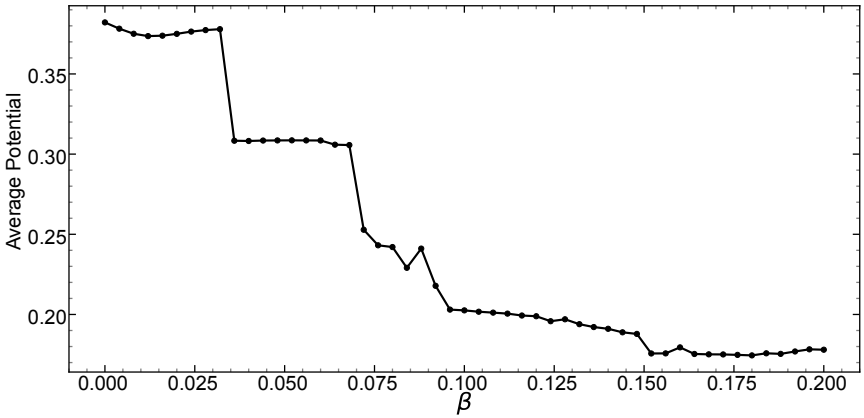


FIGURE 3.14: Average potential as a function of β

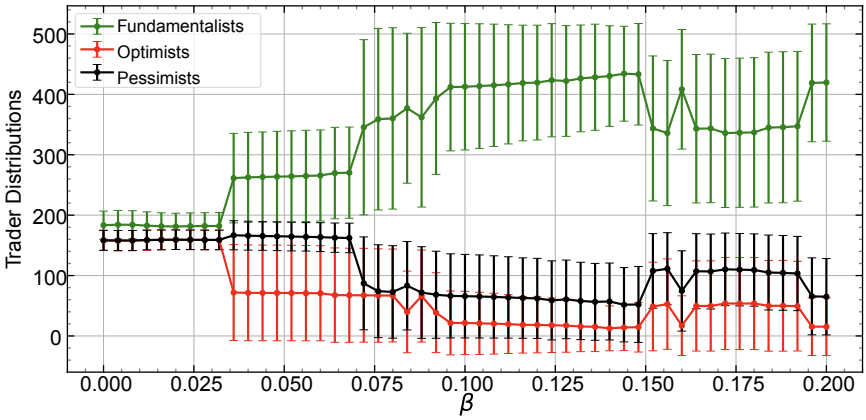


FIGURE 3.15: Oberserved trader distributions as a function of β

3.6.2 *The relationship between inhibition and decay*

As discussed previously, both decay and inhibition are crucial parameters to consider as possible drivers of irrational behaviour. While inhibition can lead to primacy bias, decay can introduce recency bias, because strong inhibition causes aggregated information to be lost quickly. McClelland and Usher describe a perfect balance of both as ideal for rational decision-making [118].

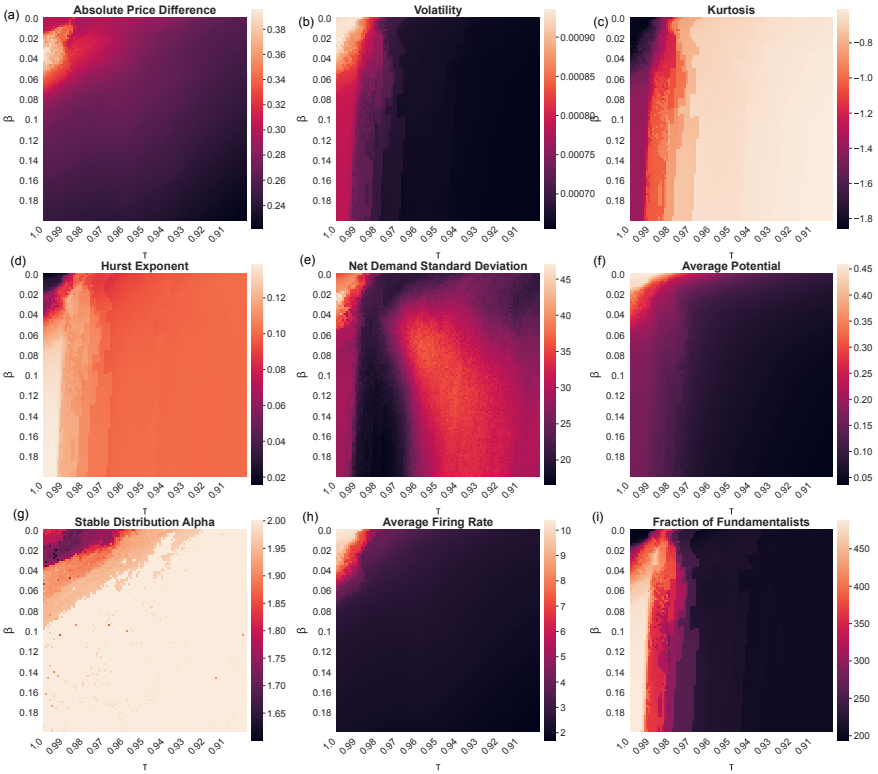


FIGURE 3.16: Phase diagrams for different measures on simulation results for combinations of β and τ.

Figure 3.16 shows heatmaps of some of the measures for different values of inhibition β and decay τ. The plots can be regarded as a sort of phase diagram for the integrate-and-fire model. If a perfect balance between inhibition and decay existed, we would expect to see a cyclical or ellipsoid pattern between both decay and inhibition. We indeed observe almost circular patterns for the average potential, average firing rate and the shape of the return distribution (α). These are to be expected, because both inhibition and decay reduce the potential, so the amount of information aggregated, so it no huge surprise that we see this confirmed here. The stable distribution alpha first increases with both inhibition and decay as irrational behaviour increases and then tends towards a normal distribution with both large α and decay, which concurs with the return distributions

observed in Fig. 3.13. An almost circular pattern can be observed for the volatility, however there is an area of medium volatility in the low decay range, irrespective of inhibiton.

For the absolute price difference (subplot a) and the standard deviation (subplot b) of net demand, we see a pattern emerging in the low-decay area, where there seems to be a peak for both around $\beta = 0.02 - 0.04$, so a certain amount of inhibition is needed to observe strong price differences emerging.

Both kurtosis (c) and the Hurst exponent (d) are largely driven by decay, with high decay pushing the hurst exponent towards a value of 0.08 and kurtosis around -0.8. However, both clearly have a circular pattern in the low-decay and low-inhibition phases. It should be noted, that the kurtosis and the Hurst exponent were calculated at a simulation-step level, which is 0.01 timesteps. The fraction of fundamentalists (i) in the simulation is maximised in the low-decay range and with high inhibition.

3.7 DISCUSSION AND CONCLUSION

In this chapter, we introduced inhibitory effects into the integrate-and-fire trader model to better understand how biases arising from neuronal-level decision mechanisms can shape financial market outcomes. The focus on inhibition, inspired by neuroeconomic decision-making models, allowed us to investigate how primacy and recency biases emerge endogenously and how these biases can drive market dynamics far from what classical rational-agent models would predict.

Our findings are multi-fold. First, we confirmed that even a modest introduction of inhibitory effects can drastically alter market behaviour. Without inhibition, the model exhibits roughly balanced opinion distributions among fundamentalists, optimists, and pessimists, and the market price tends to follow the fundamental value closely. In this scenario, entropy is relatively high, and the return distribution is close to normal.

As soon as inhibitory effects are introduced, the picture changes dramatically. We observed regime shifts similar to phase transitions, which manifest themselves as the market rapidly transitioning between equilibrium and speculatory states or even collapsing into prolonged periods dominated by a single opinion. During these phases, primacy bias is a key driver of agent

behaviour. With strong inhibition, early established opinions and states become "locked in", which causes persistent deviations of market price from the fundamental values. These locked-in states result in irrational valuations, sometimes leading to persistent undervaluation or overvaluation (bubbles) that would be impossible to explain through rational expectations alone.

However, thankfully the model provides not just outside observable market price information, but also internal variables. Especially the agent potentials provide early warning signs of upcoming transitions that would remain hidden if one only looked at observable market outcomes. Before a system-wide shift in opinion and pricing dynamics, the potentials showed increasing fluctuations, acting as a precursor signal. This underscores the strength of our approach: by modelling neuronal-level decision dynamics explicitly, we gain access to internal state variables that can give us clues about impending market shifts.

We also explored entropy as a measure of disorder in the model. Under non-inhibitory conditions, market equilibrium corresponds to a state of high entropy and minimal differences between fundamental value and market price - a result that is in line with the efficient market hypothesis and classical economic theory. With increasing inhibition, however, the entropy landscape changes. At moderate levels of inhibition, we find that states of non-zero price-fundamental deviation maximize entropy, suggesting a fundamental shift in what constitutes the "natural" equilibrium of the system. Rather than converging to fundamental value, the market self-organizes into states of persistent mispricing. This result provides a striking demonstration of how biases at the individual (neuronal) level can lead to systemic biases at the macroeconomic level.

Further analysis of the return distributions showed that the introduction of inhibition and the associated irrationality led to the emergence of fat-tailed distributions. Such heavy-tailed returns are a well-documented empirical feature of real financial markets and have long been a challenge for traditional, rational-based equilibrium models. Our results suggest that processes akin to lateral inhibition in neural circuits can help explain why extreme events - crashes and bubbles - are more common than Gaussian-based models would predict.

Examining phase diagrams in terms of inhibition and decay parameters reinforced that both biases - primacy from inhibition and recency from decay - interact nonlinearly. We observed that particular combinations of decay (how quickly old information loses relevance) and inhibition (how strongly early signals dominate) produce patterns consistent with known stylized facts like excess volatility, speculative bubbles, and bursts of activity. These findings align with our intuition that human cognitive constraints and biases, instead of being irrelevant details, actively shape large-scale market dynamics in fundamental ways.

In conclusion, by incorporating inhibitory effects inspired by neuronal decision-making into an agent-based financial market model, we have shown that primacy and recency biases can drive persistent and systemic deviations from rational benchmarks. These mechanisms lead to the formation of bubbles and crashes, increased volatility, and fat-tailed return distributions, all of which are observed in real-world markets. Our results suggest that understanding cognitive and neural-level processes is crucial for explaining the complexity and frequent irrationalities observed in financial markets.

The next chapter will explore another key aspect of financial markets - hierarchies in information structure, and examine how the interplay between hierarchical network structures and the cognitive biases discussed here can further influence market outcomes.

4

HIERARCHICAL INFLUENCES

In the last chapters, we have introduced the integrate-and-fire trader model and established its ability to reproduce some aspects of trader behaviour. So far, we have observed the effects that inhibition and decay have on market price behaviour. Irrational behaviour causes rapid market movements, which can lead to bubbles and crashes. The integrate-and-fire model is particularly suited to address these crashes because of its focus on information processing.

However, we have only looked at information processing on an individual level. Agents process all of the information that they are subjected to and make corresponding decisions. This is certainly not the only cause of inefficient use of information in markets. On the contrary, the *structure* of the market is a crucial factor to account for when analysing market inefficiencies. As previously explained in Sec. 1.5, a growing body of literature in the field of econophysics and complexity economics emphasizes that the structure of the market itself is critical in determining market behaviour and can contribute significantly to inefficiencies and instability.

Hierarchical market structures, in particular, have received increasing attention as a source of systemic risk and as potential catalysts for extreme events such as bubbles and crashes [70, 146, 147].

Benoit Mandelbrot, when investigating fractals and their applicability to financial data, laid the groundwork to understand how hierarchical and self-similar structures can shape the behaviour of financial markets [47, 146]. Mandelbrot's fractal perspective suggests that some of the observed stylized facts in markets, such as volatility clustering and heavy-tailed distributions, can be explained with hierarchical clustering of information and trading activity.

Following a similar intuition, Mantegna and Stanley investigated the presence of hierarchical organization in cross-correlations of stock returns,

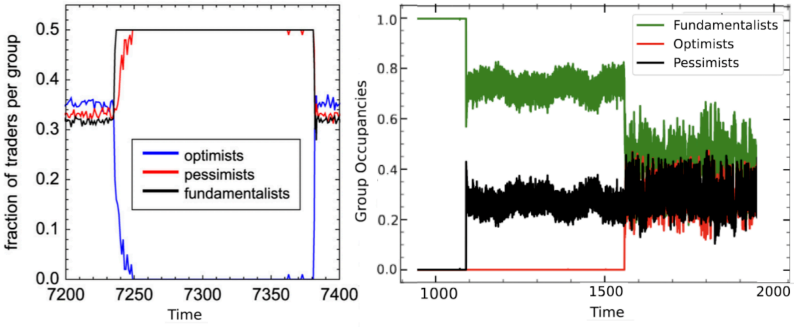


FIGURE 4.1: Comparison of rapid shifts in group occupancies. Left: transitions on the hierarchical adaptation of the Lux-Marchesi model [74]. Right: Integrate-and-fire model.

demonstrating how firms or sectors can form clusters of shared information influences [147]. Hierarchies in these networks reflect asymmetries in information flow and accessibility, implying that certain nodes (e.g., institutions, funds, or particular classes of traders) have more influence and access to information than others.

Hierarchies in markets can be thought of as "silos of information", meaning that certain actors have access to more information than others. Institutional investors or certain highly connected actors may have greater capacity to influence market expectations and prices due to their position in the information network [148]. This inequality in information distribution is a major factor that needs to be considered when investigating information asymmetries in markets. As described in Sec. 1.5, hierarchical structures in markets can amplify small fluctuations into broader market movements, leading to system-wide events like crashes. We were able to show that augmenting the Lux-Marchesi model [70], a well-established agent-based model known for generating realistic bubble-crash dynamics with an explicit hierarchical structure, can lead to the emergence of similar systemic instabilities [74].

As can be seen in Fig. 4.1, the bubbles resulting from hierarchical information imbalances look strikingly similar to the emergence of bubbles due to irrationality in the integrate-and-fire model. This observation suggests that both irrational processing at the individual level and hierarchical

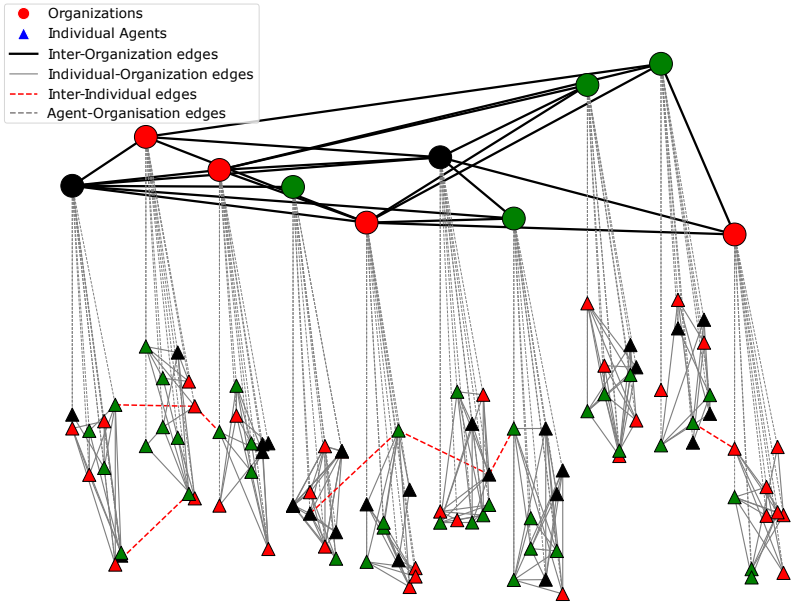


FIGURE 4.2: Two-layer network of agents. Top-level organisations are connected in an inter-organisational network. Bottom-level: Individual agents within an organisation are strongly connected, with sparse connectivity to individuals in other organisations

structuring at the network level share common mechanisms in generating destabilizing feedback loops and large-scale, abrupt market shifts.

Since the integrate-and-fire model focuses on information, it is also a suitable candidate for investigating the effect of hierarchies in markets. This is because the hierarchies in financial markets are *information hierarchies*, meaning they limit the amount of information that the individual has access to. Due to the flexibility in model structure, we only have to make minor changes to the model setup to incorporate hierarchies into this model and to investigate their effects on agent and market behaviour. Figure 4.2 shows how hierarchies can be incorporated into the model.

We decide on a simple two-level setup, with a top-level network that is a network between organizations and a bottom-level network that connects

the individual agents. The agents within each organization are strongly connected and share information, while communication with agents in other organizations is usually sparse.

4.1 HIERARICHAL NETWORK MODEL DEFINITION

4.1.1 *Cluster Formation and Network Structure*

We begin by partitioning the total number of 1000 agents into 50 clusters of 20 agents each. This clustering introduces a two-level network structure:

INTRA-CLUSTER NETWORKS Within each cluster, agents are connected according to a network topology that models the interactions between agents in the same cluster.

INTER-CLUSTER NETWORK Clusters themselves are nodes in a higher-level network, representing the interactions between different clusters. This inter-cluster network models how groups influence each other, with connections representing channels through which information or sentiment can spread between clusters.

INDIVIDUAL INTRA-CLUSTER-CONNECTIONS Agents can themselves share information with individual agents in other clusters.

In this new hierarchical integrate-and-fire model, agents are influenced by both interactions within their own cluster and by influences from other clusters. Agents receive information from their neighbors within the same cluster. These interactions represent stronger ties due to social proximity, shared interests, or frequent communication. Clusters can influence each other by sharing decisions made on an organizational level. Cluster-level decisions can affect connected clusters through the inter-cluster network. This models the spread of sentiment or information between groups, which effectively captures the hierarchical nature of influence in social and economic systems, where individuals are affected by their immediate surroundings as well as by larger societal trends.

4.1.2 *Dynamics of the Hierarchical Model*

To adapt our model to include hierarchical clustering, we must adjust the dynamics to account for the new structure. We retain the core mechanism of agents integrating information and firing when thresholds are met but introduce additional layers of influence due to the hierarchical arrangement.

4.1.3 *Agent Potentials and States*

Each agent maintains an action potential vector $\mathbf{u}$, integrating information towards each of the three possible opinions: fundamentalist (u_f), optimist (u_+), and pessimist (u_-). The potentials evolve over time based on incoming currents, and agents switch their opinion when any component of their potential vector surpasses the corresponding threshold ϑ_i. After firing (i.e., switching opinions), agents reset their potentials and continue to integrate new information. This mechanism is similar to the integrate-and-fire model described in the previous chapter but now includes additional sources of influence due to the hierarchical structure.

4.1.4 *Currents and Information Integration*

Agents receive still currents from three primary sources:

$$\mathbf{I} = \mathbf{I}_{\text{price}} + \mathbf{I}_{\text{net}} + \mathbf{I}_{\text{rand}} \tag{4.1}$$

where:

$\mathbf{I}_{\text{price}}$ represents information derived from market movements, such as price changes and deviations from fundamental values. This current is the same for all agents.

$\mathbf{I}_{\text{network}}$ includes currents from other agents, now separated into intra-cluster and inter-cluster components. Intra-cluster currents come from agents within the same cluster, while inter-cluster currents result from the influence of other clusters.

$\mathbf{I}_{\text{rand}}$ represents random external influences, capturing unpredictable factors that may affect agent decisions.

To model the influence of clusters on agents and other clusters, we introduce the concept of cluster states. The clusters are implemented as "meta-traders" and symbolize the behaviour of an organisation as a whole. We assume scale invariance, so that the clusters follow the same mechanics as the individual agents. The clusters have an action potential, which receives currents from the agents within the same cluster. They do not receive random or market information, since the individual agents already receive this information; instead they represent the organizational level of decison-making. Otherwise, the mechanics are the same: They possess action potentials $u_{c,f}$, $u_{c,o}$ and $u_{c,p}$, which represent the organizational action potentials. They integrate the triggers T_n, $n \in C_n$, from all of the agents within the cluster. As such, they follow the differential equation

$$\tau_{i,nc} \frac{du_{c,i}}{dt} = -\left[u_{c,i}(t) - u_{i,r}\right] + \sum_{i \in nc} T_i(t), \tag{4.2}$$

where $\tau_{i,nc}$ is the decay of the action potential of cluster n_c for opinion i. When an organisation makes a decision, it switches state and then communicates that information through the inter-cluster-network M. As such the new current $\mathbf{I}_{\text{network}}$ received by traders is given as

$$\mathbf{I'}_{\text{network}} = \mathbf{I}_{\text{network}} + \mathbf{I}_{\text{cluster}}, \tag{4.3}$$

where $I_{\text{cluster,j}}$ is given as

$$I_{\text{cluster},j} = \omega_4 \sum_i T_{nc} m_{ij} \tag{4.4}$$

for an agent in cluster j, where T_{nc} is the output of cluster n_c and ω_4 defines the influence of the hierarchal information.

4.1.5 *Updating Potentials and States*

At each simulation step, agents perform the following actions:

1. **Integrate Currents**: Agents update their potentials by integrating the incoming currents, including market currents, intra-cluster currents from connected agents, inter-cluster influence from clusters whose dominant opinion has changed, and random currents.
2. **Check Thresholds**: Agents compare their potentials to the thresholds. If any component surpasses the threshold, the agent fires and switches to the corresponding opinion.

Parameter	Value	Parameter	Value
Number of agents	1000	dt	0.01
τ_{cluster}	0.98	τ	0.98
θ	1	ω_1	[0.05, 0.3, 0.3]
ω_2	0.05	t_c	1
γ	1	μ	0
β	0.05	θ_{cluster}	5

TABLE 4.1: Simulation Parameters

3. **Reset Potentials**: After firing, the agent resets the potential for the opinion that caused the firing event.
4. **Apply Decay**: Potentials decay over time if not reinforced by incoming currents, reflecting the diminishing influence of past information.

4.1.6 *Market Dynamics*

The market dynamics in the hierarchical integrate-and-fire model are the same as those in the previous chapter. The market price is still based on the excess demand ΔE generated by the agents, n_+, n_-, n_f, the number of optimistic, pessimistic and fundamentalist traders, respectively. It is calculated using the chartist demand D_c and the fundamentalis demand D_f:

$$D_c = (n_+ - n_-)t_c, \tag{4.5}$$

$$D_f = n_f \gamma \frac{p_f - p}{p}, \tag{4.6}$$

$$\Delta E = D_c + D_f, \tag{4.7}$$

where t_c is the mean trading volume per unit time, γ is a parameter representing the strength of fundamentalist influence, p_f is the fundamental price and p is the current market price. After defining the fixed model parameters in Tab. 4.1, we now want to investigate the model behaviour.

Since we want to investigate the effect of hierarchical information in comparison to cluster information, we vary the parameters $\omega_2 \in [0, 0.1]$ and

$\omega_4 \in [0, 0.05]$, which describe the effect of the network current and cluster current, respectively.

4.2 SIMULATION AND RESULTS

4.2.1 *Simulations*

We start by investigating changes in price behaviour with increased hierarchical information. The left side of Fig. 4.3 shows how the market price behaviour changes in the model, starting with no hierarchical information (top) to strong hierarchical influences (bottom). The right-hand side shows the excess demands ΔD at each timestep.

There are a few striking observations to make here: Firstly, in the simulation without hierarchical influence, the market price closely follows the fundamental value, while the net demands only show almost no volatility clustering. In the next plot, for $\omega_4 = 0.002$, we observe a shift in the demand behaviour: we now see strong volatility clustering. When we further increase the hierarchical influence, this pattern becomes more pronounced: We observe strong volatility clustering, with the volatility spikes far exceeding the baseline.

At $\omega_4 = 0.003$, we observe a transition both in market price as well as in the demand behaviour. We see a clear deviation emerging between fundamental and market price, which mirrors the behaviour in Chap. 3. With strong hierarchical influence we can observe the emergence of bubbles, similar to how inhibition can lead to bubbles. However, there is another striking transition here: The demand no longer shows any volatility clustering or spikes. When we further increase ω_4 to 0.004, we observe yet another transition: The previously observed deviation of the market price from the fundamental value does not exist anymore, while the volatility shows even less variability than before.

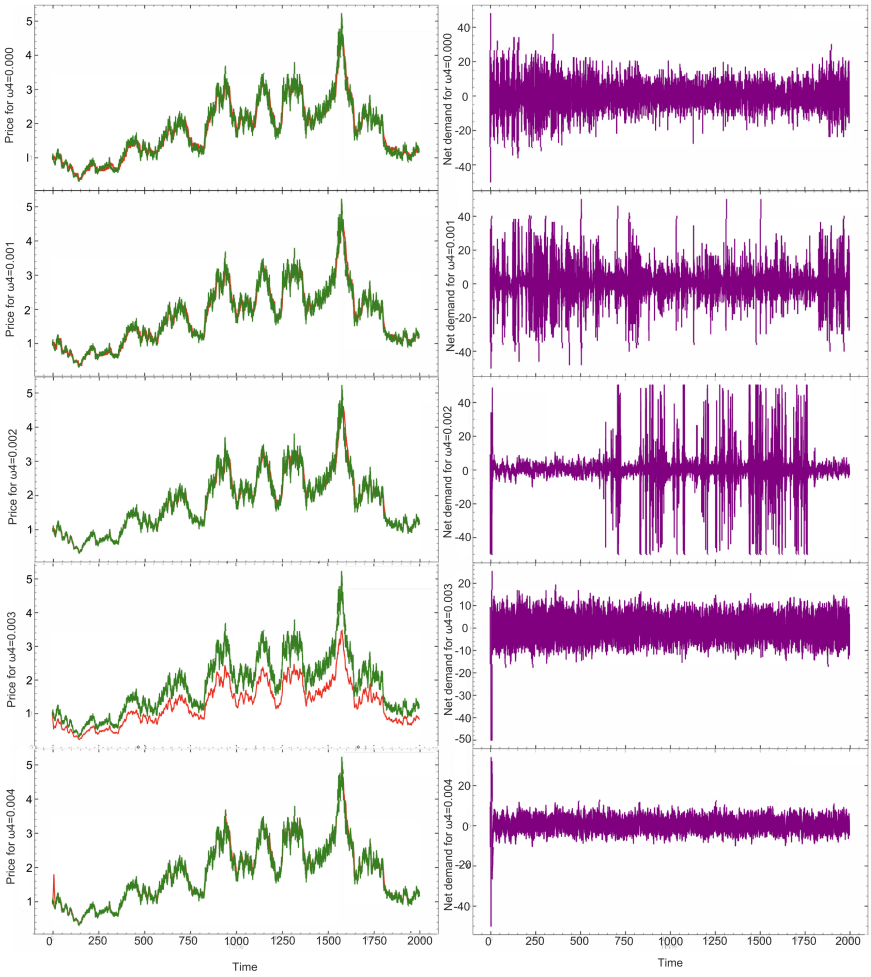


FIGURE 4.3: Fundamental (green) market (red) price (left) and net demand (right) for ω_4=0, 0.001,0.002,0.003 and 0.004 from top to bottom, respectively.

This observed market and price behaviour seems difficult to comprehend from just looking at Fig. 4.3: Both market price behaviour and demand undergo multiple (seemingly unrelated) transitions. To gain a better understanding of this behaviour, we once again have to take a look at the model states to understand this. The states for both the individual agents and the clusters are depicted in Fig. 4.4.

Here, on the left side, the individual trader occupancies are depicted, while on the right side, the cluster occupancies are shown. There are 1000 individual traders organized in 50 clusters. The first three rows of Fig. 4.4 show a seemingly clear picture: Increasing the hierarchal influence leads to a convergence of opinion both among individuals and organisations (clusters). Especially the plot for $\omega_4 = 0.002$ shows the model collapsing to a fundamentalist majority with suddenly emerging periods of speculation, while the lower-hierarchy strength scenarios show the model in constant shift.

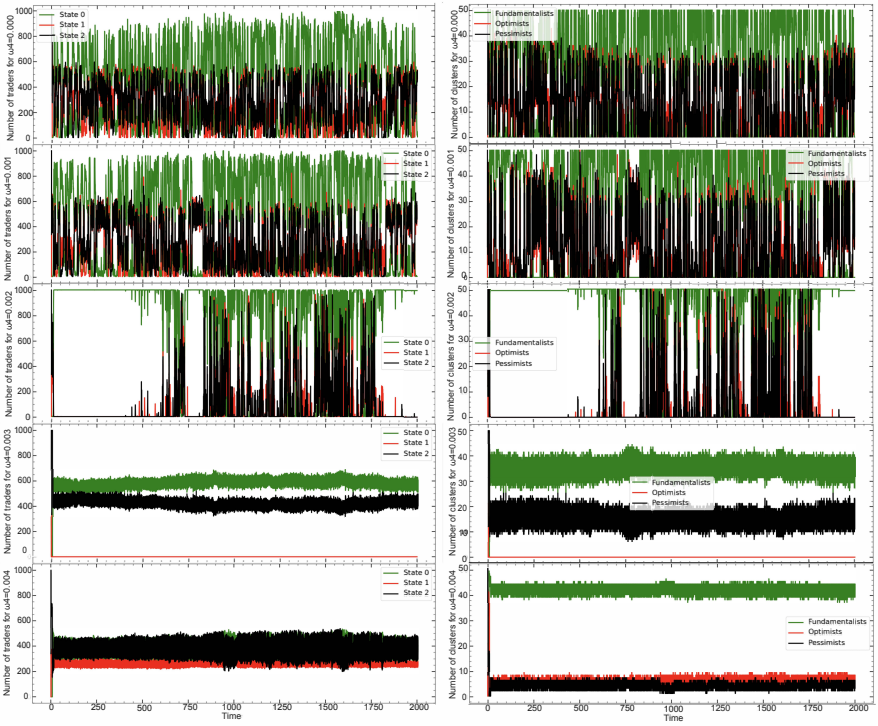


FIGURE 4.4: Occupancy of states for individual agents (left) and clusters (right) for ω_4=0, 0.001,0.002,0.003 and 0.004 from top to bottom, respectively.

This pattern breaks down when we further increase ω_4 to 0.003. Here, we observe a system where one chartist trader opinion is completely suppressed (in Fig. 4.4 the optimist opinion), while the individual trader occupancies for both of the other groups are approximately equal, with the fundamentalists occupying a small majority. Similarly, the cluster states show an absence of optimists, with a stronger dominance of fundamentalists. The hierarchy amplifies small majorities in opinion.

Cluster occupancies usually mirror the individual occupancies, except for the strong hierarchy case $\omega_4 = 0.004$, where the two are uncoupled - while all three trader groups have approximately the same occupation amongst individual traders, the clusters almost all converge on a fundamentalist opinion. So, in the case of a strong hierarchy, organizations default to a conservative fundamentalist trading approach even if the majority of their members do not agree. It should be noted here that a small fraction of organisations remain open to speculation.

In Chap. 3, we determined that agent potentials and triggers are a crucial source of information to understand how the underlying mechanics of the model work, allowing us to observe mechanics in markets that are not usually observable by taking an outside view of trader behaviour. We perform the same analysis here, which is why both potentials and triggers are depicted in Fig. 4.5.

Both an analysis of the triggers and the potentials reveal similar patterns to the ones observed for the trader states: For the first three rows of the plot, we can observe the potentials transitioning from fluctuating around 0.3 to a clearly emerging order of fundamentalist potentials relatively constant around 0.5 (with small fluctuations) and spikes of chartist potentials between 0 and 0.6. Again, potentials show precursors of transitions: For $\omega_4 = 0.002$, the market displays spontaneous emergence of chartists around $t = 500$. While there are no precursors in the visible observable market price, trader occupancies and triggers, the potentials show a clear ramp-up of chartist potential before the transition.

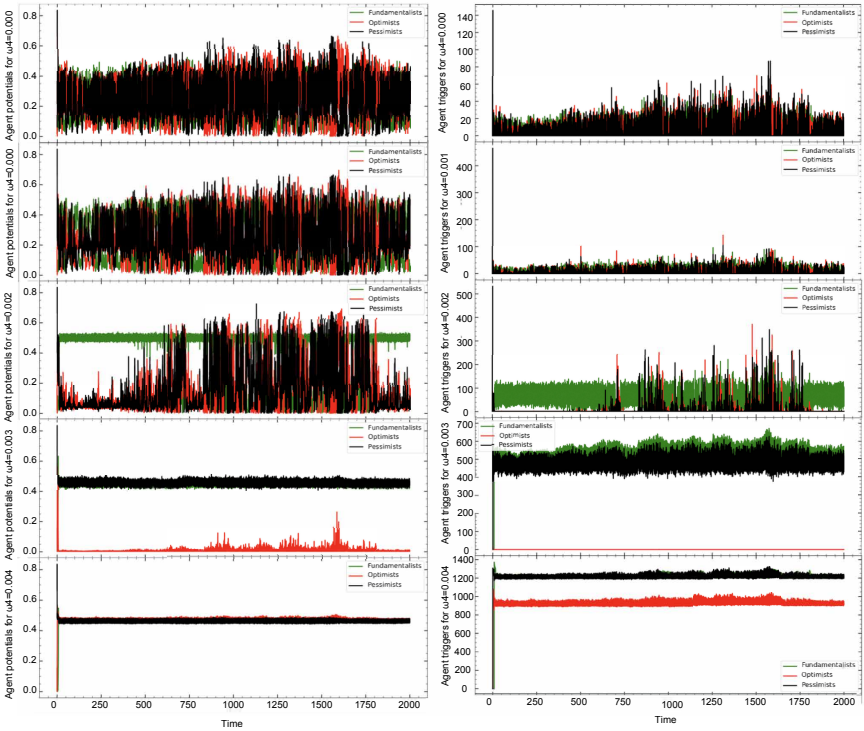


FIGURE 4.5: Agent potentials (left) and triggers (right) for ω_4=0, 0.001,0.002,0.003 and 0.004 from top to bottom, respectively. The peak in optimism around $t = 1500$ follows a long period of increasing fundamental values (see Fig. 4.3).

The potentials also show the transitions observed for $\omega_4 = 0.003$ and $\omega_4 = 0.004$, respectively. In the fourth row, the market clearly collapses to a state with almost no optimistic potential (even though small spikes can be observed), while for the strong hierarchy case, all opinions have the same potential.

Analysing the triggers shows a transition at $\omega_4 = 0.002$: We transition from a regime where all triggers spike from baseline of 0 to a regime where fundamentalist triggers fluctuate around 50. In contrast, the chartist triggers spike only during certain periods. These periods are the ones where we observe the strong volatility spikes and clusters in Fig. 4.3 and the rapid changes in trader occupation in Fig. 4.4. For $\omega_4 = 0.003$, we observe no op-

timist activity, with fundamentalists and pessimists both exhibiting triggers fluctuating around 450. For the strong hierarchy case $\omega_4 = 0.004$, we see high activity for all trader groups, which corresponds to the fact that the trader distributions between opinions are almost uniform.

So far, we have relied on the results of one simulation setup with varying hierarchal influence to study the model behaviour. We have made a couple of fascinating observations: Firstly, we can clearly relate the introduction of hierarchies to the model with the emergence of volatility clustering. Secondly, the model undergoes multiple transitions that affect all of the measurable outputs of the model: We can observe the emergence of speculation, opinion convergence towards fundamentalism with brief spikes in chartism, precursors to transitions in the potentials and more. However, to make sure that we present a comprehensive analysis of the model, we need to investigate different setups of the model with different parameters.

4.2.2 *Phase diagrams*

Phase diagrams have already proven useful in Chap. 3, allowing us to spot patterns in model behaviour driven by parameter choices. Since we want to investigate the effects of hierarchical information, we vary both ω_2, which controls the effect of the individual-level network and ω_4, which controls the effect of the cluster-level network on the currents that agents receive. We run simulations and compute multiple measures for each simulation. This allows us to create phase diagrams, which clearly show the model undergoing multiple transitions.

The phase diagrams in Fig. 4.6 show results for a grid search over a parameter range of $0.0 - 0.005$ for ω_4 and $0.0 - 0.1$ for ω_2. The first thing that becomes obvious when looking at the diagrams is that the model behaviour is largely driven by the influence of the hierarchy ω_4. For any given network influence ω_2, there are multiple transition points where we can observe drastic changes in model behaviour. The multiple different phases can be clearly observed when studying the fraction of fundamentalists in the model. Starting with approximately $\frac{1}{3}$ of trader occupation, we can clearly observe a transition towards almost complete fundamentalist dominance with increasing hierarchy strength, with a clearly defined domain wall. By further increasing ω_4, we can observe another transition towards approximately 50 percent occupation.

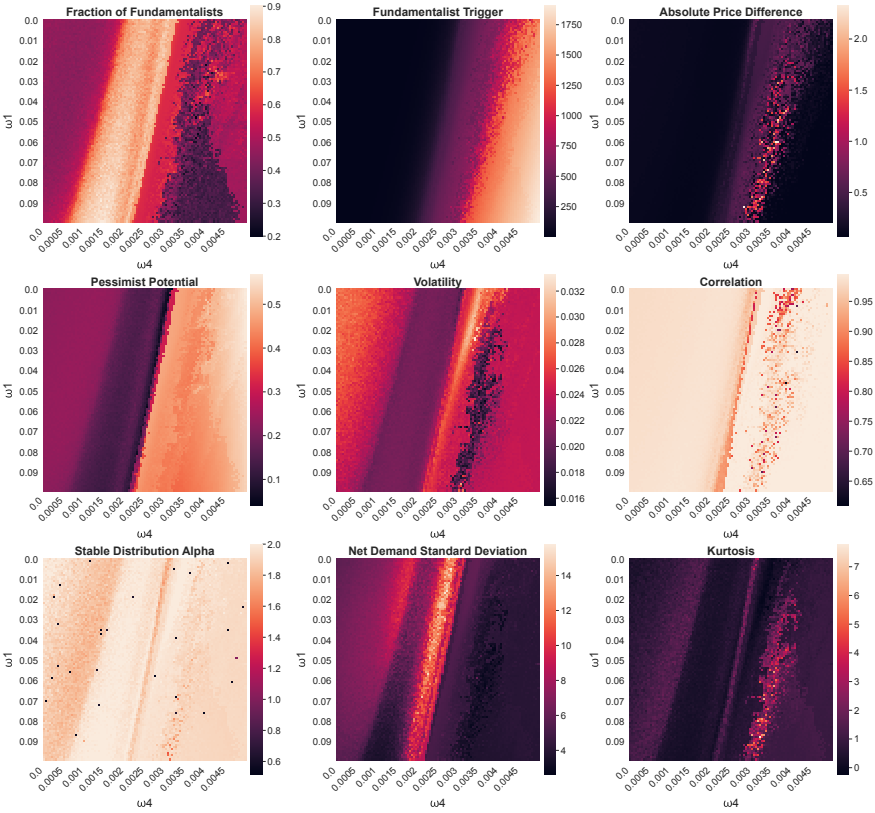


FIGURE 4.6: Phase diagrams for various measures of trader states, potentials and market price.

Similar patterns can be observed all throughout the diagrams: Both the triggers and potentials increase with the influence of the hierarchy. However, these increases do not happen steadily but instead as rapid phase changes at the same transition points as the fraction of fundamentalists. An interesting pattern can be observed for the absolute price difference: This measure describes the difference between market and fundamental value, and this is a measure of the amount of speculation in the model. There is a clear area where we can observe speculation, which agrees with the individual simulations we analyzed in the previous section.

Despite measures based on the states of the traders, we can also observe

model behaviour based on the price movements. Here, another fascinating pattern becomes apparent: In addition to the patterns observed for the measures based on agent states, there is a clear spike for all of the measures associated with the price behaviour, clearly visible by the bright line crossing the plots for volatility, correlation, stable distribution α, the standard deviation of the net demands and the kurtosis. To gain a deeper understanding of what makes this area special, we once again rely on entropy, a key measure that helped us understand model behaviour in the previous chapter.

4.2.3 *Entropy*

As introduced in the previous section, entropy is a very useful measure to understand the interplay between market behaviour and agent states. We use the same definition for entropy as defined in Sec. 3.5, so our analysis is based on the individual trader states. On the left-hand side of Fig. 4.7, the entropy for the different values of ω_4 is depicted. At first glance, the figure seems almost incomprehensible. However, a clearer pattern becomes apparent once we realise that there are two main mechanics at work here: Firstly, entropy stays almost flat for all values of ω_4, so it is practically independent of the market price deviation. This is in stark contrast to the non-hierarchical model, where entropy was strongly driven by price deviations. We can observe that entropy remains unchanged when we increase ω_4 from 0 to 0.001, which is in line with the observation that there is no fundamental change in behaviour within that range.

For $\omega_4 = 0.003$, however we can observe much higher entropy than previously observed, indicating a clear transition in behaviour. This corresponds to the simulation, where we can observe clear market price deviations. When we further increase the hierarchical influence, entropy drops again, which matches the expected result given that bubbles disappear in the strong hierarchy regime.

The most fascinating value, however, is $\omega_2 = 0.002$, which corresponds to the bright strip observed in the phase diagrams. To better understand the movement of this entropy graph, we plot the entropy normed by the value for the lowest observed market price deviation - depicted on the right side of Fig. 4.7.

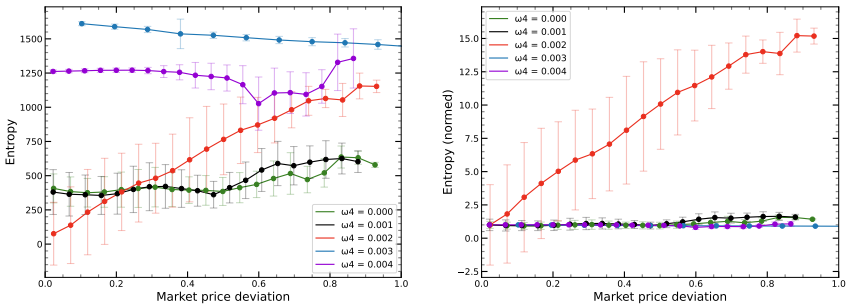


FIGURE 4.7: Left: Entropy for different values of ω_4, Right: Same plot but normed by the minimum entropy for each value of ω_4.

Here, the pattern becomes much clearer: Entropy is approximately constant regardless of market price deviations, except for the $\omega_4 = 0.002$ case, where we see a strong increase in entropy with market price deviations. This corresponds to the fact that the more disorder we observe, the higher the speculation is in this scenario. Going back to the simulation plots for $\omega_4 = 0.002$, we can explain why.

As discussed earlier and depicted in Fig. 4.4, in the $\omega_4 = 0.002$ regime, the market defaults to full fundamentalist dominance with spontaneous emergence of speculation. Interestingly, the system goes in and out of speculation, so it changes order spontaneously, resulting in the corresponding entropy behaviour. If we conjecture the same behaviour for entropy as in Chap. 3, we conclude that the market is continuously driven towards a state of equilibrium, however, the hierarchical influences drive it back towards a state of order by imposing a tendency towards a strong majority.

4.3 DISCUSSION AND CONCLUSION

In this chapter, we introduced hierarchical structures into the integrate-and-fire trader model to investigate how asymmetries access to information shapes market outcomes. Drawing from literature in complexity economics and econophysics [70, 146, 147], as well as from Mandelbrot's foundational work on fractals and hierarchical clustering of financial data [47, 146], we examined how embedding agents in a hierarchical structure can generate stylized facts and systemic instabilities that are not fully explained by

individual-level decision-making processes.

Our findings expand on the previous chapter, where we showed that incorporating neuro-inspired decision biases such as inhibition can help explain stylized facts like the emergence of bubbles and fat tails. Here, a similar dynamic emerges from the structural dimension. Hierarchical and asymmetrical information flows, conceptualized as "information silos", give certain nodes (organizations or institutions) disproportionate influence over market sentiment, as also suggested by Mantegna and Stanley's analysis of stock return cross-correlations [147]. Just as inhibition led to persistent biases and regime shifts in agent opinion distributions, hierarchical structures can amplify minor shifts in opinions, thus pushing the market away from equilibrium and into states of high speculation, which manifests itself as volatility clustering. Asymmetric information and influence can result in the dominance of one or more opinions.

A key result is that even small hierarchical influences can induce dramatic changes in market dynamics. Without hierarchical influences, fundamentalist, optimist, and pessimist opinions balance each other out. This baseline scenario reflects a high-entropy market environment where prices behave relatively efficiently. This corresponds to classical economical models that often assume equal access to information and rational agents. However, introducing hierarchical influences quickly destabilizes this equilibrium. Information can accumulate in certain clusters and as their sentiment shifts, the entire market can follow suit, leading to sudden spikes in volatility, persistent deviations of the market from the fundamental prices, and cascades of speculative behaviour.

These hierarchical-induced dynamics bear resemblance to phenomena studied extensively in econophysics: volatility clustering, heavy-tailed return distributions, and pronounced bubbles and crashes [70, 148]. When hierarchical structures are strong, the market can collapse into regimes where one type of opinion (often the fundamentalist view) dominates cluster-level decisions. In such phases, short-lived speculative bursts can emerge and vanish abruptly. When we investigated the model's internal variables, such as agent potentials and triggers, we found that these measures can help predict shifts in market regimes. This mirrors the findings from the last chapter: there, inhibition-driven potentials provided early-warning signals of impending transitions. A look at the agent potentials and firing patterns

allows us to draw insights that traditional, price-focused analyses might miss.

Investigating these states through entropy and phase diagrams revealed multiple phases of market organization. When hierarchical influence is absent or weak, the market's "natural" equilibrium state is one of high entropy and near alignment with fundamental values — mirroring efficient market conditions. As hierarchical influence strengthens, however, we enter regimes where mispriced states display high entropy, indicating that persistent departures from the fundamental value can become stable equilibria under structural asymmetries. This finding complements the picture from the previous chapter where cognitive biases shifted the market's equilibrium away from fundamentals; now we see that hierarchical structures can have a similar systemic effect.

The presence of hierarchical structures in combination with moderate network-level influence parameters can trigger volatility clustering and generate return distributions with fat tails. As dicussed previously, these well-know stylized facts have been associated with structural factors like hierarchies in real-world markets. By explicitly modelling hierarchical information channels, we create a mechanistic understanding of how structural asymmetries might contribute to these complex market behaviours, extending beyond the purely cognitive explanations offered previously. The results presented here resonate strongly with previous work that combines the Lux–Marchesi model with hierarchical structures [74], where it was shown that combining the Lux–Marchesi model with hierarchical influences leads to the spontaneous emergence of speculation, bubbles, and crashes. In that earlier study, the model was modified to incorporate scale-invariant, asymmetric information flows, ultimately reproducing known characteristics of financial markets, including volatility clustering and fat-tailed distributions of returns. Crucially, it revealed that for certain hierarchy strengths, markets evolve into regimes characterized by uniform opinions and sudden mood swings, which give rise to sharp market movements and speculative behaviour. Our current integrate-and-fire hierarchical model reproduces many of these patterns: it exhibits abrupt shifts from stable "fundamentalist" dominance to bursts of speculative sentiment and back again. Just as in the hierarchical Lux-Marchesi framework, these transitions can be triggered internally, without the need for exogenous shocks.

Comparing these two approaches highlights a key point: hierarchical structures themselves, whether applied to a traditional agent-based framework or to a neuro-inspired model of decision-making, are capable of producing systemic instabilities and stylized facts reminiscent of real financial markets. In the earlier hierarchical model, the emergence of speculation and sudden crashes could be traced to critical parameters defining the strength of the hierarchy's influence. In our integrate-and-fire model, we also found distinct phases of market behaviour, with abrupt transitions that depend on both hierarchical connectivity and influence parameters. Both models emphasize that structure and especially hierarchies are a central driver of market outcomes, which needs to be considered alongside to decision processes.

This chapter shows that the complexity and inefficiency of real financial markets cannot be reduced solely to the cognitive biases of individual traders. While the previous chapter highlighted how neuro-inspired inhibition and decay parameters can induce substantial deviations from rational benchmarks, the current chapter shows that macro-level hierarchical structures play an equally important role. Hierarchical information flows can independently induce market fragility, volatility clustering, and systemic mispricing, aligning with the growing body of literature that emphasizes the structural organization of markets as a key determinant of their instability [70, 148].

This concludes our investigation of the single-stock model. In the next chapter, we will generalize the integrate-and-fire ABM into a general framework and introduce another crucial determinant of agent behaviour: constraints.

5

A NEW FRAMEWORK FOR SOCIAL INTERACTION

From the previous investigation, it has become apparent that the integrate-and-fire model is indeed a suitable candidate for opinion modelling. Especially in a setting where multiple opinions are competing, it provides a way to negotiate these conflicting options and observe the infectious spread of opinion, the formation of fixed opinion clusters and even spontaneous shifts in opinion. Behavior is heavily dependent on network structure. This, combined with the variability and adaptability of the model (the model can stabilize even if put into precarious parameter spaces), is why it can be used to describe decision processes in various systems.

The original integrate-and-fire trader model introduced in Chap. 2 has a very tight integration of opinion modelling and market dynamics. This is due to the fact that a firing event corresponds to a switching of state, and this state directly influences the market through the Lux-Marchesi price dynamics. With this tight integration, many stylized facts can naturally emerge, like fat-tailed return curves, reversion to the mean, crashes and others. However, how the integration is implemented somewhat limits the systems to which this can be applied. Often, we observe very complex opinion dynamics that do not necessarily correspond to intricate price patterns.

The model can also be expanded with a crucial part of human motivation: constraints. In the case of just one asset, an argument can be made to ignore these constraints if we assume a very large market. Many single-asset agent-based models let agents work without a budget. However, in most markets, constraints are very relevant, and they will heavily influence the actions of a trader (you should not buy a racecar if you cannot afford it, no matter how much you might want it). In addition, behaviour might depend on the past actions of the trader. This can include refractory periods, which have been shown to lead to synchronization [149]. To account for constraints, we introduce the integrate-and-fire model as an opinion modelling framework in which opinion spreading and trader actions are separated. Within this framework, a wide variety of opinion-sensitive systems can be

modelled. The framework as a concept will be presented in the following, and then an application example will be discussed in Sec. 5.2.

5.1 TOWARDS A GENERALIZED FRAMEWORK

There are a couple of things to unpack when considering a generalized framework for decision processes using the integrate-and-fire approach. The first is to define the scope and the elements of the framework, how they interact with each other and what kind of systems the framework can be applied to. As described in Chap. 2, the integrate-and-fire model allows us to bring the results of behavioural- and neuroeconomics into the space of agent-based modelling. It does so by introducing an information-gathering and decision-making process that accounts for the behaviour on the neuronal level in a highly abstract way. Still, it maintains many of the essential features of decision-making.

While decision-making is certainly an important aspect of a model, as becomes evident when observing how much the integrate-and-fire trader model deviates from the Lux-Marchesi model, it only forms part of a comprehensive model. Other parts of the model include what we will call the "system", which essentially defines the playing field on which agents interact. For economic applications, this is most likely a market, but can also include governments, nature and other entities that might interact and be affected by the actions of agent behaviour. In the integrate-and-fire trader model, this role was filled by the Lux-Marchesi price dynamics, aggregating demand and supply and changing the market price accordingly. In Sec. 5.2, we will discuss a much more elaborate system, an order book.

Another crucial part is the agents that drive system behaviour. Of course, the initial model had agents, which were defined by their action potential u, their threshold ϑ and their state i. Agent's actions simply depended on their state and always acted on the system in the same way. To describe more realistic versions of the market, it is desirable to enable the possibility of more elaborate agent designs. This can include agents whose actions are under constraints and depend on state variables and others. The final part is the role of information, which can be generated by both the system and by agents' actions and that influences how agents act. In the following, all three core modules of the framework will be introduced.

5.1.1 *The role of system*

The system is defined by its state

$$\mathbf{S}_t = (v_{t,1}, \ldots, v_{t,K}), \tag{5.1}$$

which is defined by the running variables $v_{t,m}$, $m \in K$. The running variables hold all of the relevant information about the system at time t. In the case of the initial model, this corresponds to the current market prices p_t, the previous one p_{t-1}, and the fundamental values $p_{f,t}$. In a more elaborate setting, this could relate to the current state of a government's policy, seasonal variables or outstanding orders in an order book. The system's evolution depends on the actions of the N agents, which are given as A_n. Thus, the system's state at time t will depend on both the system's past states and the traders' actions at time $t-1$.

5.1.2 *Agents*

The role of actions - described in the most abstract way - is to transform the state of a system $\mathbf{S}_t$ at time t, given the information I into an output A, which will act on the system. Each agent underlies a decision-making process governed by the leaky integrate-and-fire mechanic. For a single agent faced with M possible opinions, the action potential can be described as a set of differential equations

$$\frac{du_i}{dt} = (u_i - u_{r,i})\tau_i + I_i(t) + \sum_{j \in M} \beta_{ij} u_j, \quad i = 1, 2, \ldots, M. \tag{5.2}$$

The evolution for the action potential u_i for opinion i is defined by multiple terms. The first, $(u_i - u_{r,i})\tau_i$, describes the decay of the action potential over time towards the resting potential u_r. Here, τ_i defines the timescale for the decay of potential for opinion i. The current $I_i(t)$ describes all of the information regarding opinion i that the agent is subjected to. They are added to the action potential because they are aggregated as part of the decision-making process. The final term $\sum_{j \in M} \beta_{ij} u_j$ can describe inhibitory effects like the ones found in the leaky computing accumulation model [118]. For inhibitory effects, β_{ij} will be negative $\forall i, j$, meaning that if agents are close to a decision for option i, the action potential u_i is close to the threshold ϑ_i, this will reduce the other action potentials $u_{j \neq i}$, making it less likely for agents to choose them.

See Chap. 3 for an argument explaining why this is a valuable addition.

The opinion state of an agent is described by the M-dimensional action potential $\boldsymbol{u}$, a resting potential $\boldsymbol{u_r}$ and a threshold $\boldsymbol{\vartheta}$:

$$\boldsymbol{u}_t = \begin{pmatrix} u_{t,1} \\ \vdots \\ u_{t,M} \end{pmatrix}, \boldsymbol{u_{r,t}} = \begin{pmatrix} u_{r,t,1} \\ \vdots \\ u_{r,t,M} \end{pmatrix}, \boldsymbol{\vartheta} = \begin{pmatrix} \vartheta_1 \\ \vdots \\ \vartheta_M \end{pmatrix}, \tag{5.3}$$

At each timestep t, the agent will receive a current

$$\mathbf{I}_t = \begin{pmatrix} I_{t,1} \\ \vdots \\ I_{t,M} \end{pmatrix}, \tag{5.4}$$

Like the original integrate-and-fire trader model, when u_i reaches the threshold ϑ_i, the agent will act. However, while the original model had a clear and predefined action, fire and switch state, in the generalised framework, the action can depend on constraints and state variables of both the system and the agent. This action usually consists of two parts: an interaction with the system and an information transfer to connected agents and is best illustrated by the following example:

Let's say you receive information that absolutely convinces you that a stock i will skyrocket next week. In a world without constraints, you would invest a huge sum into the stock, expecting a huge profit. However, after logging into your bank account to make the investment, you are confronted with the harsh reality that you are a broke PhD student unable to invest in anything except basic necessities. In this case, your constraints prevented you from engaging the market based on your information. You might, however, still complain to your friend about your situation, who, having saved all of his hard-earned cash from teaching classes, might take the opportunity to invest in the stock. So, while you did not act on the system, you still transferred some of the information to other agents, which is why both actions can happen independently in the new framework.

So the process is clear: When agents are driven to act, they first check their constraints and their personal state variables D_t and other system

variables, and then act on the system and transfer information to other traders. Their action output

$$A(n,t) = \begin{cases} 0 & u < \vartheta_i \forall i \in M \\ f(D_{n,0}, \ldots, D_{n,t-1}, S_0, \ldots, S_t, \{i\}) & \exists i \in M : u_i >= \vartheta_i \end{cases} \tag{5.5}$$

thus depends on the relevant opinions i, as well as their state variables D and the system variables S. Their information output, the trigger T will depend on the action they choose:

$$\mathbf{T}_{n,t} = f(A_{n,t}). \tag{5.6}$$

Based on their action $A_{n,t}$, they will update their own state variables

$$D_{n,t+1} = f(D_{n,0}, \ldots, D_{n,t}, A_{n,t}), \tag{5.7}$$

as well as those ofthe system

$$S_{t+1} = f(S_0, \ldots, S_t, A_{1,t}, \ldots, A_{N,t}). \tag{5.8}$$

Finally, regardless of the action, the action potential u_i will be reset to $u_{i,r}$ after firing.

5.1.3 *Information*

The model assumes the existence of M different options with which agents are faced. These can refer to the choice of M different stocks, M different political opinions and many others. The parameter M determines the dimensionality of the model. Each opinion i has a corresponding dimension, which can also be considered as an information channel. The currents are the drivers of the agent dynamics, and they are typically made up of three parts: Changes in the system state, actions of other traders and random information.

$$\mathbf{I}_t = \mathbf{I}_{t,sys} + \mathbf{I}_{t,net} + \mathbf{I}_{t,rand}. \tag{5.9}$$

$\mathbf{I}_{t,sys}$ is the information that a trader receives from observing the system's state. This can be market movements, sales numbers or other external information.

$$\mathbf{I}_{t,sys} = f(\mathbf{S}_0, \ldots, \mathbf{S}_t). \tag{5.10}$$

$\mathbf{I}_{t,net}$ is the information that a trader receives directly from other agents. Agents are connected through a network with an adjacency matrix $\mathbf{K}$, a $N \times N$ matrix for N agents. The output of each agent $\boldsymbol{T}$ is a function of their action $\boldsymbol{A}$, and these outputs are then transmitted to the other agents in the network. Regardless of how the triggers are calculated, agents receive the triggers from all of the traders they are connected to. For a normal network, the network-induced current for agent n is given as

$$\mathbf{I}_{net}, n = \sum_{j=1}^{N} T_{j,t} k_{jn} = \sum_{j=1}^{N} f(A_{j,t}) k_{jn}. \tag{5.11}$$

In exceptional cases, like system structure incorporating hierarchies, the information transfer can take different forms.

5.1.4 *Simulation Steps*

A simulation is set up as follows:

1. Set initial conditions: Initial system state $\mathbf{S}$, initial action potentials $\mathbf{u}_i$
2. Calculate currents based on system state
3. Apply currents to agent, evaluate triggers
4. Agents evaluate actions if they are triggered
5. Calculate agent output based on actions
6. Let actions act on the system
7. Go to step 2 and repeat

The difference to the simulation in the original model is that agents if triggered, will choose the action they take based on constraints and then act instead of simply switching states and firing.

The general introduction to the framework was very abstract, so we will now try to fill it with life by applying it to a multi-stock market.

5.2 APPLICATION: EXTENDED MODEL

The original version of the integrate-and-fire trader model describes just one actively traded asset. This can be an aggregate of the global stock market, a

currency-exchange pair or something else. Within that framework, it can be used to evaluate a wide range of phenomena. However, examining not just the opinion of traders concerning a single asset but their behaviour when they are confronted with a choice of different stocks can bring even further insights into market dynamics.

To gain insight into this kind of behaviour, the model needs to be rebuilt from the ground up to account for the very different dynamics. The integrate-and-fire framework allows us to build a more ambitious model that follows the same intuition. There are a couple of core market features that will be incorporated into this new model. The first, as mentioned before, is that the model now allows traders to choose between different stocks. The second is that agents now hold a portfolio of stocks in which they allocate portions of their wealth between different stocks. Another change is that agents now exchange stocks via an order book, allowing for a much more microscopic view of market behaviour. All of these changes make the model more realistic without adding unnecessary complexity. Indeed, as we will see, some of the additions even make the model more straightforward.

5.2.1 *Model basics*

The model deals with N agents trading M different stocks in an order book exchange. Despite the significant changes, we will keep two of the Lux-Marchesi dynamics: Stocks still have underlying fundamental values and traders' opinions are still driven by optimism, pessimism and fundamentalism. However, these influences will now be combined into a single dimension. We will go through each part of the new model in detail to explain the changes and the new mechanics.

5.2.2 *General: Multi-Stock*

Upgrading the model from a single-stock variant to a multi-stock version requires redefining some of the fundamental entities of the model. The first is that the notions of state and opinion do not apply anymore for a system of multiple stocks. Agents do not have general views of the economy, they only have views regarding specific stocks. That is why each dimension of

the model now corresponds to the opinion about a single stock. Information is still communicated with multidimensional currents

$$\mathbf{I} = \begin{pmatrix} I_i \\ \vdots \\ I_M, \end{pmatrix} \tag{5.12}$$

where $i = 1 \dots M$ are the different stocks between agents can choose. The model still works under the assumption that traders underly both chartist and fundamentalist influences. However, they are accounted for in a different way in this new model. Since there is now just a single dimension for each asset, all available information must be incorporated in a one-dimensional manner. This poses a challenge: How can we incorporate the different influences (optimism, pessimism, fundamentalism) into a single dimension? For the single-stock model, the opinions about the future of the stocks were described by the current

$$\mathbf{I}_{price} = \begin{pmatrix} I_f \\ I_+ \\ I_- \end{pmatrix} \propto \begin{pmatrix} \text{abs}(p - p_f) \\ \max(0, \frac{dp}{dt}) \\ \max(0, -\frac{dp}{dt}) \end{pmatrix}. \tag{5.13}$$

The information flows captured by these currents still apply to any of the individual stocks, but we need to reduce the dimensionality. We help ourselves by simply adding the influences of each dimension if they convey a positive opinion about the stock or subtracting them if they convey a negative opinion. The new currents are essentially the norm of the original current. However, appropriate signs are added to reflect the different directions in which they might push a trader's opinions about a stock. From the definition of optimists, pessimists and fundamentalists, it is clear how these influence a trader. Optimistic influences will make a trader more likely to buy a stock, while pessimistic influences will discourage him from buying a stock. Fundamentalist influences will encourage traders to buy a stock if the market price p is below the fundamental value p_f and discourage them if $p > p_f$. This is why the market influence for stock i is given as

$$I_{market,i} = \alpha_1(p_i - p_{f,i}) + \alpha_2 \max(0, \frac{dp_i}{dt}) - \alpha_3 \max(0, -\frac{dp_i}{dt}). \tag{5.14}$$

Here, dp_i/dt is the derivative of the price of stock i, and $p_i - p_{f,i}$ refers to the difference between the market price p_i and fundamental value $p_{f,i}$ of stock i.

Replacing the previous dimensions with single-stock dimensions may seem like an enormous deviation from the original model. However, it indeed follows the exact same intuition and the dynamics of the multi-dimensional integrate-and-fire model are utilized in the same manner. In the original model, the three dimensions refer to the different possible opinions about a stock, and the dynamics of the model provide a description of how these compete against each other. Traders are constantly subjected to information that might push them to make decisions based on fundamentalist or chartist opinions, and these influences are balanced through competition in different dimensions. Similarly, in the new model, a trader is constantly facing the choice of allocating its wealth between different options. Traders receive information that might lead them to buy a range of different stocks, and all these influences compete with each other via multidimensional currents. So, in both the original and the stock model, the different dimensions correspond to different options from which the trader must choose and which are competing for attention.

After getting the initial definitions out of the way, we will now take a look at the system in which the agents operate: a market with an order book.

5.2.3 *System: Order book*

The Lux-Marchesi model employs a simple yet powerful approach to modelling the effect of traders' actions on market price traders: By using a macroeconomic aggregation of demand and supply, prices rise when there is excess demand and fall when there is excess supply. It is very fundamental way of modelling market forces, but a lot of information is lost in the process. There are multiple reasons why a different way of setting prices is essential.

In the Lux-Marchesi model, every trader has the same average influence on the market. Stock price changes simply depend on the number of agents in each group. However, in real market settings, there is significant heterogeneity between actors in financial markets, with some taking billion-dollar speculative positions and others investing their leftover cash at the end of the month. This is an important distinction since large individual trades can have a significant impact on markets. A famous example is the large

effect that some individuals can have on cryptocurrency markets [150].

Another reason is that to understand price patterns on short timescales, market microstructure is important. The way traders exchange assets and information shapes the reactions of markets to news, the emergence of speculation and others. The original integrate-and-fire model is very suited to model the spread of information between agents. Still, many of the dynamics on short timescales are lost because of the highly aggregated price mechanics. Introducing an order book allows us to capture the entirety of agent behaviour and should leave us with a much richer description of the market.

A limit order book is an electronic list of buy and sell orders for a financial instrument like a stock, organized by price level. It is used to match buyers and sellers in financial markets. The book lists all outstanding limit orders, so orders are placed to buy or sell at a specified price or better. Orders are then sorted from best to worst based on the limit price. While buy orders are sorted from highest price to lowest (best bids at the top), sell orders are the reverse, with the lowest asks (offer prices) at the top. When a new order comes in to buy or sell, it gets matched to the best price available on the opposite side of the book. For example, a buy order gets matched with the lowest ask. If an order exceeds the volume available at the best price, it keeps getting matched at the next best price, moving down the list. If orders that are not immediately matched enter the market, they get added to the book on the appropriate side and ranked by limit price. Limit order books provide liquidity to markets and reflect dynamic supply and demand. Market movements are shaped by changing volumes and depths.

In this new agent-based model, agents can register orders when they decide to increase or decrease any of the positions in their portfolio. Orders have a lifespan of five timesteps, meaning that they are automatically cancelled if they have not been filled after that time period. While many order books enforce a minimum trading quantity, for the sake of simplicity, we refrain from introducing this into the new model. Orders can be partially matched if the order volumes differ.

stock id	1	2	3	4	5	risk-less
price	0.1	0.4	0.2	0.1	0.4	1
holdings	3	2	6	3	3	12
wealth in asset	0.3	0.8	1.2	0.3	1.2	1.2

TABLE 5.1: Portfolio of trader n before wealth reallocation

5.2.4 *Agents: Portfolio*

As part of introducing heterogeneity into the model, agents are now subject to constraints. While they effectively had infinite budgets in the initial model, they now allocate limited wealth between different assets. That is why we introduce portfolios for agents. In this version of the model, the trader must choose to allocate his wealth between different stocks i and a risk-less asset. The riskless asset has a fixed price and provides guaranteed interest r. Each agent holds its wealth in a portfolio with holdings h_i of different assets. Each asset has a price p_i, which results in a wealth invested in asset i of

$$q_i = h_i \times p_i. \tag{5.15}$$

An example of an agent's portfolio is shown in Tab. 5.1.

In addition, traders have a memory in that they remember the order in which they decided to buy the assets in their portfolio. Desirability is ranked based on when an asset was bought last. The ranked portfolio from Tab. 5.1 can be seen in Tab. 5.2.

When traders decide to increase investment in an asset, they must give up some other investment in exchange. Investments in stocks can either be financed by cash, so holdings in the riskless asset, which can always be liquidated immediately, or by selling positions in other stocks. Typically, traders will try to execute their new investment choices as soon as possible, meaning that they will try to finance the new investment using the riskless asset and then try to sell other stocks later.

Since the model now incorporates an order book, transactions are not necessarily executed immediately. Traders cannot freely exchange different assets with each other but can only enter offers into the order book. Once these orders are completed, so their orders could be matched with other

rank	stock id	wealth in holding	fraction in wealth
1	2	0.8	0.16
2	5	1.2	0.24
3	1	0.3	0.06
4	3	1.2	0.24
5	4	0.3	0.06
-	risk-less	1.2	0.24

TABLE 5.2: Holdings of trader n ranked by preference

orders, changes in portfolio manifest.

Let's say our previously introduced trader n decides to invest wealth of 1.5 into stock 3. This would lead to the entering of at least three orders into the order book: The first would be an immediate "bid", or buy order over 1.2, so the entirety of immediately available cash, to buy stock 3. The next one would be an order to sell the position of 4, since it has the lowest desirability of any of the available stocks. The last one would be another bid order of for stock 3 that becomes activate after a successful sale of the positon in stock 4. The reason for this is that agents cannot go into debt – by only allowing the bid order to go active after the ask has cleared agents are prevented from buying stocks without the funds. To illustrate this, we will look at the portfolio after every order. For simplicity of this illustration, we will assume that stocks are traded at market price, even though this is not the case in the actual model dynamics, as described in Sec. 5.2.5.

Let's start by looking at the portfolio immediately after the orders were entered into the order book. This is depicted in Tab. 5.3. As visible in Tab. 5.3b, the three orders are now recorded in the order book: trader n tries to buy stock 3, sell stock 4, and then buy more stock 3 once the sale of 4 provides additional liquidity. Outstanding orders in an order book block the agents from trading the corresponding assets, as seen in Tab. 5.3a. This is implemented so that agents cannot offer more than they own. So far, the agent has merely registered orders, and no transactions have taken place yet. This will only change once other agents register additional orders.

Let's say a different trader, trader m, registers an ask order for stock 3

(a) trader portfolio

stock id	1	2	3	4	5	risk-less
price	0.1	0.4	0.2	0.1	0.4	1
holdings	3	2	6	0	3	0
wealth in asset	0.3	0.8	1.2	0	1.2	0
holdings (blocked)	0	0	0	3	0	1.2
wealth in asset (blocked)	0	0	0	0.3	0	1.2

(b) order book

kind	issuer	stock id	quantity	price	active	filled
bid	trader n	3	6	0.2	yes	no
ask	trader n	4	3	0.1	yes	no
bid	trader n	3	1.5	0.2	no	no

TABLE 5.3: Trader portfolio and order book immediately after trader n decides to invest into stock 3.

over 6 units, matching the bid order of trader n. In the order book, these orders are automatically matched - leading to yet another portfolio change as depicted in Tab. 5.4.

Now, the investment wish of n finally becomes a reality: The holding of stock 3 increases by 6, and both orders become inactive because they have been filled. To illustrate the last mechanic, let trader m register another order, matching the ask order for stock 4. It is, once again, immediately matched because the prices of the orders are identical.

Agent n has now successfully liquidated the position in stock 4, providing the necessary cash to activate the bid for stock 4 that was previously inactive. Tab. 5.5 shows the corresponding changes in the portfolio and order book. In the actual model dynamics, price setting for bid and ask orders vastly differs and depends upon the information available to the traders.

(a) trader portfolio

stock id	1	2	3	4	5	risk-less
price	0.1	0.4	0.2	0.1	0.4	1
holdings	3	2	12	0	3	0
wealth in asset	0.3	0.8	2.4	0	1.2	0
holdings (blocked)	0	0	0	3	0	0
wealth in asset (blocked)	0	0	0	0.3	0	0

(b) order book

kind	issuer	stock id	quantity	price	active	filled
bid	trader n	3	6	0.2	no	yes
ask	trader n	4	3	0.1	yes	no
bid	trader n	3	1.5	0.2	no	no
ask	trader m	3	6	0.2	no	yes

TABLE 5.4: Trader portfolio and order book after trader m issues an ask order for stock 3.

5.2.5 *Agents: Order Price*

We have discussed how the model handles traders' decisions to buy a particular stock for a specific price. However, a description of the process by which the agent decides which stock, at what quantity, what price and when to trade is still missing.

Agents act as a consequence of a firing event. Regarding the opinion model, the firing event in this model is very similar to the original model. However, the market implications are very different. This is due to the multi-asset nature of this updated model. The firing time for each stock i is similarly defined:

$$u(t_{f,i,n}) = \vartheta_{i,n}, \tag{5.16}$$

which corresponds to the fact that a trader fires current pointing into the direction of the information it received after that amount of information

(a) trader portfolio

stock id	1	2	3	4	5	risk-less
price	0.1	0.4	0.2	0.1	0.4	1
holdings	3	2	12	0	3	0
wealth in asset	0.3	0.8	2.4	0	1.2	0
holdings (blocked)	0	0	0	0	0	0.3
wealth in asset (blocked)	0	0	0	0	0	0.3

(b) order book

kind	issuer	stock id	quantity	price	active	filled
bid	trader n	3	6	0.2	no	yes
ask	trader n	4	3	0.1	no	yes
bid	trader n	3	1.5	0.2	yes	no
ask	trader m	3	6	0.2	no	yes
bid	trader m	4	4	0.1	no	yes

TABLE 5.5: Trader portfolio and order book after trader orders play out

reaches a threshold. After firing, the action potential regarding the particular opinion is reset according to

$$\lim_{\delta \to 0^+} u(t_{n,i} + \delta) = u_{n,i,r} \,. \tag{5.17}$$

When a trader fires in the information dimension i, the channel regarding stock i, orders are issued to reallocate a fraction of wealth into stock i. Many models that work with order books set the price of these orders to fluctuate randomly around ask, bid or mid prices. Since we have a model that captures the decision process of traders and also the information they receive, we can utilize this to capture the irrationality of agent decisions. We define a new quantity we call excitement that affects both price and volume.

$$\Gamma(u_i) = \frac{2}{1 + \exp(-u_i + u_{i,r})} - 1 \Bigg|_{t=t_{f,i}} \tag{5.18}$$

The excitement of a trader is a measure that is defined only at the firing time and measures the surplus of information that a trader receives when the threshold is reached. Since we work in discrete time, and agents receive triggers from other agents all at once, usually the threshold is not exactly hit, but overshot. This "overflow" of information means the agent wants to trade more urgently and in an order book setting, this requires price adjustment. Γ follows a sigmoid function, as depicted in Fig. 5.1.

The excitement captures a couple of effects that may impact price setting: The first is that - at the time of decision - some traders might have more positive information than others and thus will have an increased desire to invest because their expected payoff will be higher. In addition, it captures trader motivation: If I get a big push of information that pushes me over the edge, I will act more impulsively than if I only received a small additional amount of information. Lastly, it allows the element of panic in markets because avalanches of information will drive price movements.

Agents have access to order book bid and ask prices and set their individual prices based by weighting these using their own excitement.

$$bid(t,n) = \Gamma(t,n)ask(t) + (1 - \Gamma(t,n))bid(t) \tag{5.19}$$

$$ask(t,n) = \Gamma(t,n)bid(t) + (1 - \Gamma(t,n))ask(t). \tag{5.20}$$

Given their internal valuation, order prices are then set based on

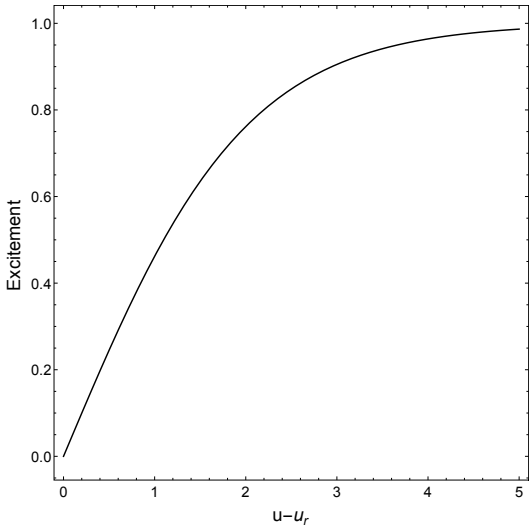


FIGURE 5.1: The excitement of a trader scales with surplus information and has a sigmoid shape.

$$p_{bid}(t,n) = N(bid(t,n,\sigma)) \tag{5.21}$$

and

$$p_{ask}(t,n) = N(ask(t,n,\sigma)) \tag{5.22}$$

for any order issued by trader n at t.

5.2.6 *Information: Network and Cross-Influences*

Like before, agents receive information in the form of current

$$\mathbf{I}_t = \mathbf{I}_{t,sys} + \mathbf{I}_{t,net} + \mathbf{I}_{t,rand}. \tag{5.23}$$

The system part of the current $I_{t,sys}$ is given by the previously introduced price current:

$$I_{sys}(t) = I_{price}(t), \tag{5.24}$$

which we repeat for completeness

$$\mathbf{I}_{price} = \begin{pmatrix} I_f \\ I_+ \\ I_-. \end{pmatrix} \propto \begin{pmatrix} \text{abs}(p - p_f) \\ \max(0, \frac{dp}{dt}) \\ \max(0, -\frac{dp}{dt}) \end{pmatrix}. \tag{5.13}$$

Prices in the order book are determined by the mid prices, meaning the average of bid and ask.

Network

In addition to market movement, agents also receive information from other traders. When a trader decides to buy a stock i, this information is shared with all the traders that are connected to that trader. Thus, the network part of the current is given as

$$\mathbf{I}_{net,i,n=} \sum_{j=1...N} T_{j,n} k_{jn} \tag{5.25}$$

Similar to the original model, $\mathbf{T}_i$ is the trigger and K is the connection matrix, also referred to as $(k)_{ij}$. The trigger $\mathbf{T}_{i,n}$ is the output of trader n with regard to stock i into the network at each timestep. The last current $I_{rand,i,n}$ is the random information a trader n receives about stock i at any given time.

Cross-Influence

As previously described, action potentials in the integrate-and-fire framework follow a differential equation

$$\frac{du_i}{dt} = (u_i - u_{r,i})\tau_i + I_i(t) + \sum_{j \in M} \beta_{ij} u_j, \tag{5.2}$$

where the last part describes possible inhibitory behaviour between different options. In the stock market, however, cross-influences are not necessarily negative. It is well known that certain stocks are correlated, and these correlations will often inform trading decisions. That is why agents, when integrating information, will include these cross-influences, resulting in the following differential equation

$$\frac{du_i}{dt} = (u_i - u_{r,i})\tau_i + \sum_j \rho_{ij} I_i(t) + \sum_{j \in M} \beta_{ij} u_j, \tag{5.26}$$

with ρ_{ij} being the cross-influence of information regarding stock j on the perception of stock i. Agents construct these cross-influences themselves by observing correlations of prices, but since agents have limited processing capabilities, they will only selectively populate their internal correlation matrix with the other entries being zero. However, since information is shared between traders at firing events, expected correlations between stocks still influence global market phenomena.

As before, I_{rand} simulates the arrival of random news and follows a normal distribution.

5.3 RESULTS

In the following, results of the updated model are presented. Since the definition of states is not as straightforward anymore, we concentrate on some of the stylized facts produced by the model. In Fig. 5.2, some of the stock prices produced by the model are shown. Clearly, the market price follows the fundamental value, but it can deviate significantly, with sharp drops in fundamental value often inducing even larger adjustments in fundamental price.

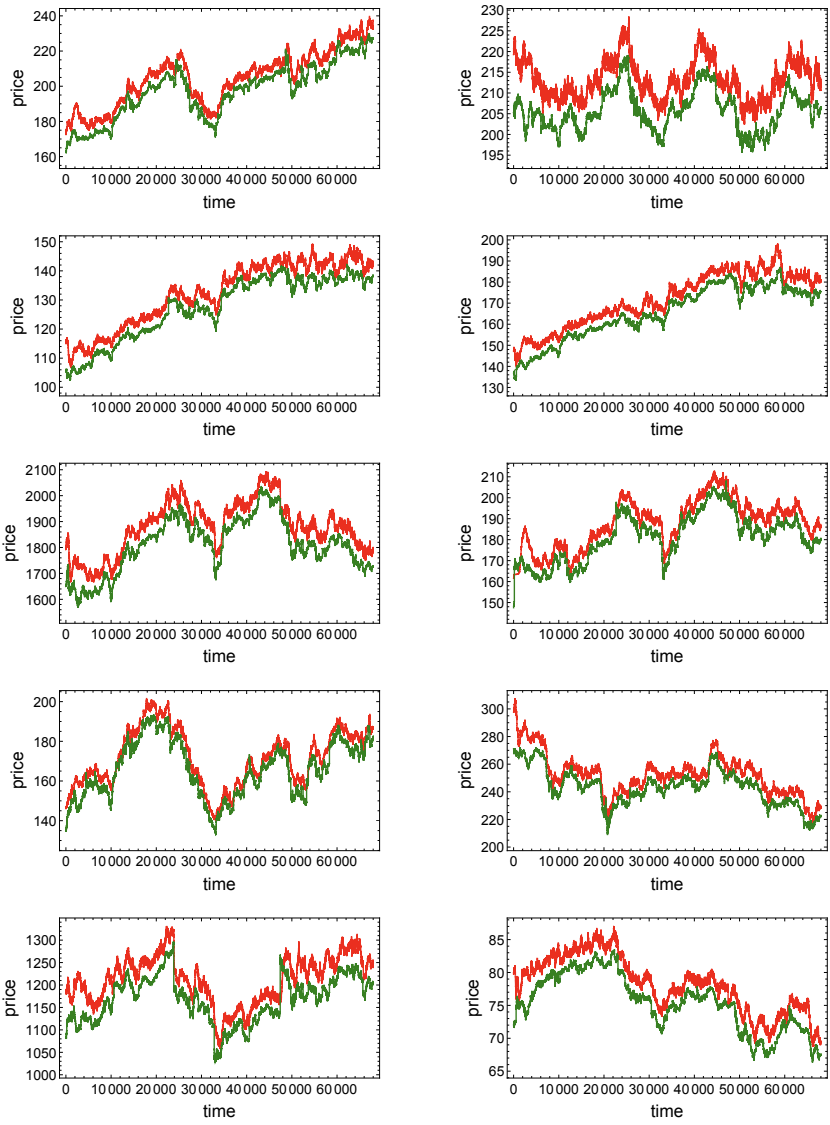


FIGURE 5.2: Plots of all 10 timeseries for one simulation of the multi-stock order book model. The fundamental values, which have been generated by fitting a geometric brownian motion to the top 100 stocks by market capitalization are plotted in green, while the market prices are red.

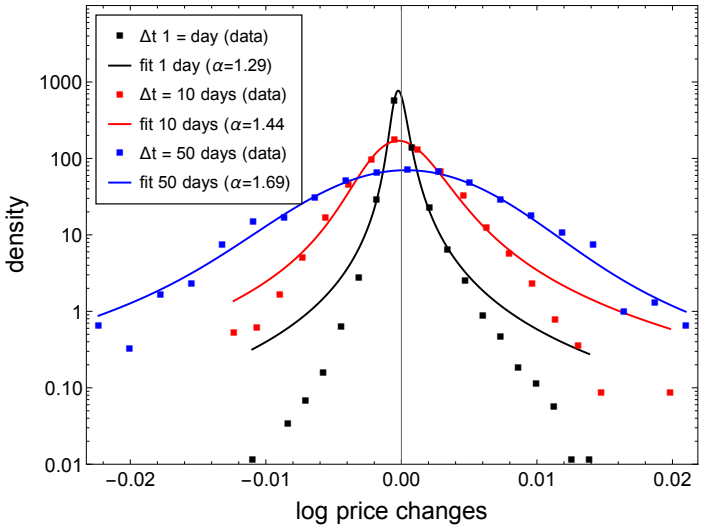


FIGURE 5.3: Log returns for a single stock for $\Delta t = 1, 10$ and 50 days. While the fit agreement is not perfect, It is till apparent that returns get consistently more-fat tailed on shorter timescales.

Looking at the market and fundamental prices is a good and necessary sanity check, but the real challenge of the model is to reproduce the stylized facts of markets. Their relevance was introduced in Sec. 1.3 and we will now investigate them. Probably the most prominent stylized fact is the fat-tailed nature of stock returns. To test whether this is the case for our model, a histogram of log-returns for a single stock is depicted in Fig. 5.3. There it becomes apparent that on short timescales, returns are indeed fat-tailed, while they tend towards a normal distribution over larger period. This matches actual market behaviour.

While single stocks are usually fat-tailed because of their high vulnerability to external factors, stock market indices tend to be less heavy on the tails since the weighted average reduces the overall exposure to risk. Figure 5.4 confirms this is also true in our model. A holding-weighted average of stocks is indeed less fat-tailed than the individual stocks.

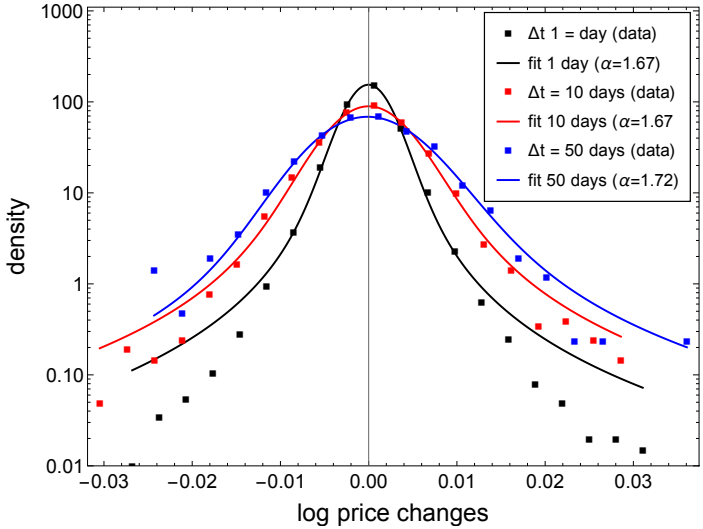


FIGURE 5.4: Log returns a weighted sum of all of the stocks by holding.

To further investigate model behaviour, we take a look at the development of total wealth of all traders over time. Since the total wealth is essentially a holding-weighted sum of all the stock prices, it serves the same rule a stock market index would in a typical exchange. Notable when examining the wealth over time is the sharp drop in wealth around timestep 35000, which mirrors the patterns observed in the fundamental price.

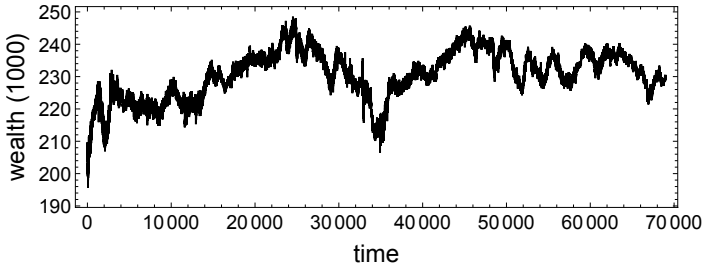


FIGURE 5.5: Total wealth of agents over time

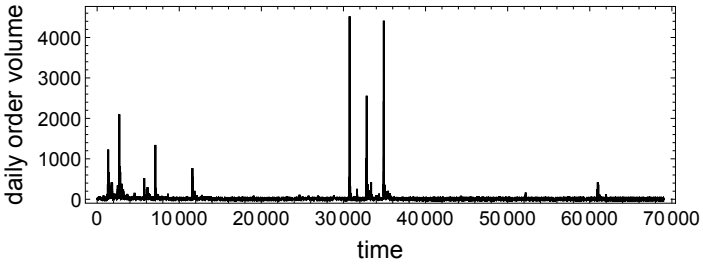


FIGURE 5.6: Timeseries of daily order volumes

The interesting model behaviour lies in the microstructure of the market, the order book, so we investigate the order volume to see whether the drop has an effect on orders. As Fig. 5.6 shows, this is indeed the case. During the crash, there is a significant increase in order book activity.

To further investigate, we look at the returns and the order book during the period of the crash in Fig. 5.6 showing the same period.

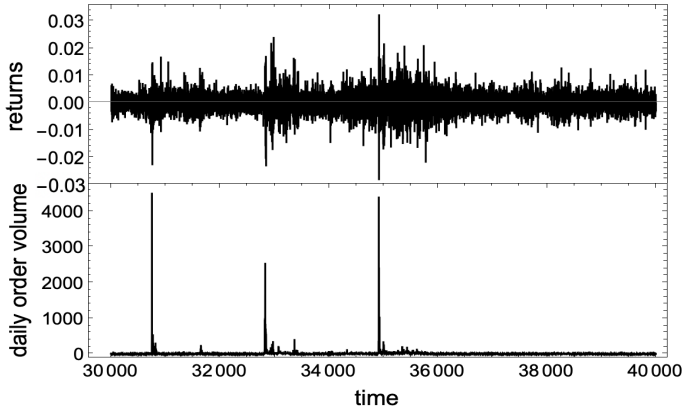


FIGURE 5.7: Top: Time series of returns between $t = 30000$ and $t = 40000$. Bottom: excerpt of the time-series depicted in Fig. 5.6. The time series of returns shows volatility clustering, and the order volume shows increased orders during volatility spikes

Here, we discover two key things: The first observation is that the model exhibits volatility clustering, another important stylized fact to reproduce (see Sec. 1.3). Second, when looking at both the returns and the order book, we confirm that the big spikes in orders occur during periods of increased volatility, another stylised fact (see Sec. 1.3).

The next step of the investigation is to look at some of the distributions the order book produces. By looking at the results, we see a common theme: Distributions tend to be fat-tailed, with extreme events exceeding even the likelihood of power law distributions. As we will see, this mirrors some of the extreme distributions found in real markets as well. We begin with a look at the daily order volumes in Fig. 5.8. As shown in Fig. 5.7, extreme shifts in market price are accompanied by extreme volumes in orders.

This has effects on the distribution of daily orders, with a clear outlier on order activity at the crash previously investigated. Next, we turn our attention to agent heterogeneity. A big argument for introducing microstructure into the model is that it allows us to better account for differences between individual traders. Fig. 5.9 shows why this is a feature worth exploring. While most of the distribution follows a pareto distribution, some extreme outliers emerge. These outliers are agents that are far more active in the market than others, effectively taking on the role of market makers. Market makers fill a crucial role in markets in that they provide liquidity. It is fascinating to see such a behaviour naturally emerge in this model, given that traders start with similar initial conditions.

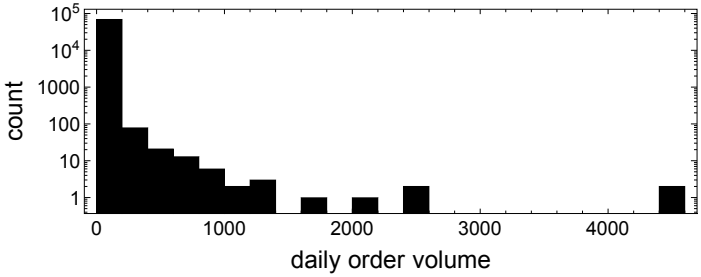


FIGURE 5.8: Histogram of daily order volumes over a simulation run.

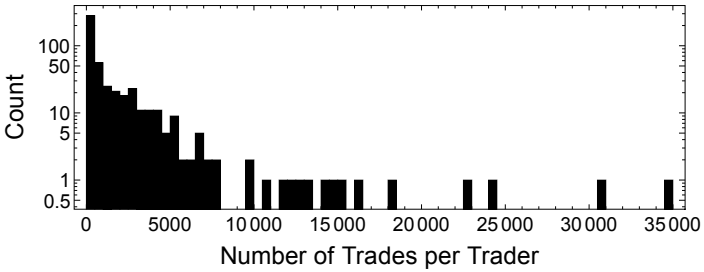


FIGURE 5.9: Histogram over the number of trades per trader. Notice that the distribution follows a pareto distribution, with the exception of a cluster of mega-active traders

The emergence of market makers is accompanied by the emergence of mega-rich individuals. Figure 5.10 shows a distribution of total agent wealth at the end of a simulation. Again, there are clear outliers that accumulate vastly more wealth then others. This is a nice illustration of Paretos 80-20 rule that we discussed in Sec. 1.1, which states that a few people control the wast amount of wealth.

A look at real markets reveals that this mirrors the wealth distribution of high net-worth individuals. By looking at the wealth of the Forbes 400, shown in Fig. 5.11, we find that the distributions look remarkably similar.

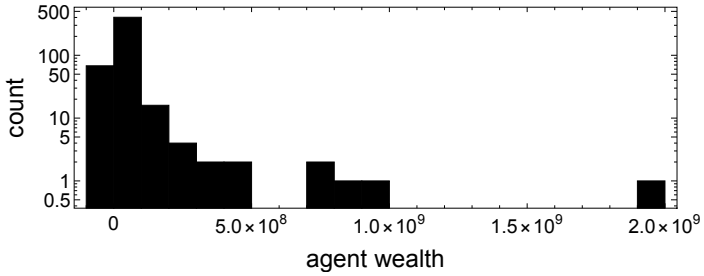


FIGURE 5.10: Histogram over the wealth of traders. Notice that the distribution follows a pareto distribution, with the exception of a cluster of mega-rich traders

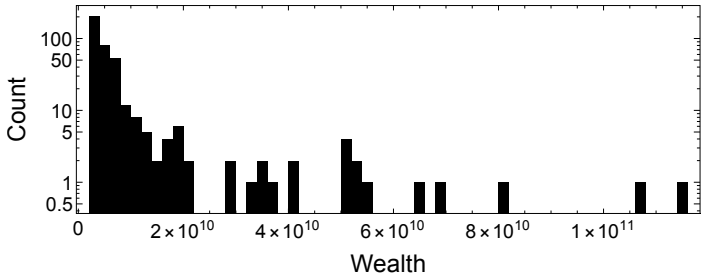


FIGURE 5.11: Distribution of wealth of the Forbes top 400 richest people. Data taken from [14].

As the results show, the introduction of an order book and agent heterogeneity make the model more realistic, and the implementations reproduce stylised facts.

5.4 CONCLUSION

In this chapter, we introduced the integrate-and-fire framework and introduced an order book model utilizing the new features. The framework describes the model as an interplay of agents and the system, and defines the role information plays as the key driver of system dynamics. By introducing an extended version of the differential equation describing the action potential, decision-making processes can now be modelled with higher fidelity.

These new definitions are put to use in an agent-based model that extends the original integrate-and-fire trader model with multiple stocks, agent portfolios and order books. Information is still the driver of decision-making, but actions are more complicated now because agents are heterogeneous, have budgetary constraints and have more options to choose from. Agents also incorporate their own observations of past market behaviour in their decision-making process. This is done by introducing cross-influences based on stock market correlations.

The resulting model can produce many of the stylized features of financial markets, such as fat tails, volatility clustering and increased order volume during volatility spikes. Market makers and inequalities emerge naturally from a population of initially very similar agents. While these results are

promising, there is room for follow-up investigations. Some areas to explore are:

5.4.1 *Order book microstructure*

We have detailed data about the orders that agents enter. It would be worthwhile to take a closer look at the micro-behaviour of these dynamics and comparing to the established theory of order book behaviour. Finding a way of relating information processing in the model to order book structure would be interesting

5.4.2 *Network behaviour*

Finding network indicators for wealth and activity in the order book would certainly be worthwhile and could answer questions about which factors drive success in markets. Are the agents with a better understanding of market correlations advantaged? The well-connected agents? Or do some people get lucky and can then use their wealth in a self-fulfilling prophecy of success?

In addition, information flow measures could be implemented to investigate the effect of shocks, which can be implemented by supplying agents with negative information. The model dynamics suggest that some individuals have more influence than others.

6

PRACTICAL APPLICATIONS

In the last chapters, we have introduced the integrate-and-fire model as an adequate description of decision processes in financial markets. However, most of the findings in the past chapters have leaned toward a qualitative description of markets. By replicating stylised facts, we were able to show that the model can indeed describe many of the processes in financial markets while accounting for the neuronal basis of decision-making. In the previous chapter, we formalised the model and introduced the integrate-and-fire framework as a general framework to model the interactions of many interacting humans.

In this chapter, we will take a step back from focusing the model entirely on financial markets and present work that shows that the integrate-and-fire model can be used to describe real-world phenomena with high accuracy, especially contagion processes. This chapter is included to show a working example of the model applied to real data.

6.1 INTEGRATE-AND-FIRE MODEL OF DISEASE TRANSMISSION

The following work presented resulted from a co-supervision of a Master's thesis (Shahbanno A. Hussain) by both my supervisor, Dimitri Vvedensky, and me. I am presenting it here because, in addition to co-supervising, the concept for the model stemmed from my work on the integrate-and-fire model, and I was the one who conceived and wrote the original project proposal. However, the majority of the credit for this work belongs to Shahbanno, as she is the one who implemented the model, ran the simulations and wrote the first draft for the academic paper published in Physical Review E [151].

The original inspiration for this model came from a simple insight: when modelling the spread of diseases, the interaction of infected humans with the outside world can be condensed to a simple process: integration and firing. We integrate by absorbing infectious matter, and when a critical

threshold is reached, our immune system becomes overwhelmed, and we begin to "fire", i.e. cough, thus spreading the infection to the ones in close proximity to us. While this might be a very abstract view of the process, it is able to capture and reproduce essential features of pandemics, as will be demonstrated in the results section later.

6.1.1 *Introduction*

From its early beginnings, epidemiological modelling has been a crucial tool to understand the spread of infectious diseases. The recent SARS-CoV-2 (COVID-19) pandemic has revived global interest in this field to never-before-seen levels. It clearly demonstrated the need for accurate models that can forecast disease transmission across populations and inform governmental policies [152]. One of the foundational models in epidemiology is the susceptible-infected-recovered (SIR) model. Introduced by Kermack and McKendrick in 1927 [153], the SIR model disguinishes the SIR model disguinishes three possible states of agents:

- **Susceptible (S)**: Individuals who are vulnerable to contracting the disease but have not yet been infected.
- **Infected (I)**: Individuals who are currently infected with the disease and are capable of transmitting it to susceptible individuals.
- **Recovered (R)**: Individuals who have recovered from the disease and are assumed to have gained immunity.

The model assumes a closed population. This means that there are no births or deaths unrelated to the disease, and that people neither enter nor leave the observed population otherwise. Disease spreads through contact between susceptible and infected individuals. The dynamics of the model are governed by a set of ordinary differential equations:

$$\frac{dS}{dt} = -\beta SI, \qquad \frac{dI}{dt} = \beta SI - \gamma I, \qquad \frac{dR}{dt} = \gamma I, \tag{6.1}$$

where: β is the transmission rate, representing the probability of disease transmission per contact times the average number of contacts per individual per time step and γ is the recovery rate, representing the proportion of infected individuals who recover per time step.

The SIR model is governed by a key parameter, the basic reproduction number $R_0 = \beta/\gamma$. It describes the average number of secondary infections

produced by one infected individual in a fully susceptible population. If $R_0 > 1$, the infection will spread through the population; if $R_0 < 1$, the infection will eventually die out. The SIR model, in its basic form, assumes homogeneous mixing, meaning that in a large enough population, everyone is equally likely to come into contact with one another [154]. Considering that this is a thesis in the field of complexity science, it must come as no big surprise to the reader that we find this description unsatisfactory. In reality, humans are much more likely to interact within certain groups, such as family members, friends, or coworkers. These social structures can significantly influence the spread of disease and must be considered for more accurate modelling, which is why a network description of social structures is necessary to provide accurate modelling. As we have seen in the previous chapters, the structure of a network can have significant effects on model behaviour, which is why this fact should not be neglected in the pursuit of a simple model.

Network structures have been implemented into SIR models [155]. Usually, nodes represent individuals that are connected by edges, which represent social connections; these edges can become occupied when infection transmissions occur [156]. Thresholding has been implemented in spiking networks, where secondary infections can occur [157]. However, these approaches fail to integrate over multiple interactions, leaving the system without any memory. The COVID pandemic, however, has made the importance of memory, e.g. acquired immunity in the population, abundantly clear [158]. This is one reason why models that allow for reinfection have proven useful in describing pandemics. One such model is the SIRS (Susceptible-Infected-Recovered-Susceptible).

The SIRS model allows agents to lose immunity over time and then become susceptible to infection again [159]. They allow for the introduction of cross-immunity matrices, which account for the partial immunity acquired by infection with one strain of a disease to others [160]. Due to natural competition, this will lead to the most infectious strains being dominant over time [161]. However, these models still often fail to properly incorporate memory and spatial resolution.

6.1.2 *Applying the integrate-and-fire model to infections*

At the core of the model's intuition is the realisation that the processes describing information transmission in neurons are strikingly similar to the transmission of infection in individuals. In the integrate-and-fire model, neurons receive currents from other neurons until a threshold is reached. At this point, the neuron fires, transmitting current to other connected neurons. It enters a refractory period, in which it becomes inactive until it starts receiving currents again. Similarly, individuals moving through society are exposed to infectious diseases all the time. However, only when enough infectious matter is absorbed (integrated) so that the immune system is overwhelmed do we become infected. An infection will usually lead to the host replicating the virus or bacteria that caused the infection, and the host will start to transmit the disease to others, in some cases literally firing out the disease (coughing). Given enough time, the host will recover and usually be immune to receiving further infections of the same kind for a while before becoming susceptible again.

Compared to the previous chapters, the mapping is very straightforward for this application and requires only limited adjustments. We use the mathematical description of a simplified integrate-and-fire model. We model the social structure of the population using an undirected network $G = (V, E)$, where V is the set of nodes (agents) and E is the set of edges (social connections) between agents. The network follows a power-law degree distribution with exponent 3, making it a scale-free network [162]. Many real social networks follow exponents around 3 [84]. Each agent $i \in V$ has a degree k_i, representing the number of connections (neighbours).

6.2 AGENT STATES AND VARIABLES

Each agent i at time t is characterized by the following variables: Its state $S_i(t)$, its infectious matter $I_i(t)$, its infection threshold θ_i and its transmission rate $\beta_i(t)$ and, if infected, its time since infection $T_i(t)$. We will now explain all of them.

The state of agent i at time t is given as.

$$S_i(t) = \begin{cases} 0, & \text{if agent } i \text{ is } \textbf{susceptible} \text{ at time } t \\ 1, & \text{if agent } i \text{ is } \textbf{infected} \text{ at time } t \end{cases}. \tag{6.2}$$

The infectious matter $I_i(t) \geq 0$ represents the amount of infectious matter accumulated by agent i at time step t. Once the infectious matter surpasses the infection threshold θ_i, the agent becomes infected. This threshold is initialised as $\theta_i(0) = \theta_0$. After an agent recovers, the threshold is increased to account for acquired immunity.

The transmission rate $\beta_i(t)$ determines the fraction of infection transmitted to neighbours in each given time step. When the threshold is updated, we also update β to offset the increased threshold, so that the amount transmitted to neighbours remains unchanged. If an agent does become infected, the time since infection $T_i(t)$ tracks the duration of disease.

6.3 MODEL DYNAMICS

The model progresses in discrete time steps Δt. At each time step, every agent i is subject to the following processes:

6.3.1 *Transmission of Infectious Matter*

Each agent transmits infectious matter to its neighbours. The amount transmitted from agent i to neighbour j at time t is:

$$T_{i \to j}(t) = \begin{cases} \beta_i(t) \cdot \dfrac{I_i(t)}{k_i}, & \text{if } S_i(t) \in \{0, 1\} \text{ and } S_j(t) = 0, \\ 0, & \text{if } S_j(t) = 1. \end{cases} \tag{6.3}$$

Here, k_i is the degree of agent i and $\beta_i(t)$ and $I_i(t)$ are the transmission rate and infectious matter of agent i at time t, respectively. We rescale the transmission rate by k_i to distribute infectious matter equally among neighbours. This scaling is implemented to account for the limited time agents have to spend with others.

6.3.2 *Reception and Accumulation of Infectious Matter*

For each susceptible agent i ($S_i(t) = 0$), the total incoming infectious matter at time t is:

$$I_i^{\text{in}}(t) = \sum_{j \in N_i} T_{j \to i}(t), \tag{6.4}$$

where N_i is the set of neighbours of agent i.

The agent updates its infectious matter:

$$I_i(t + \Delta t) = I_i(t) \cdot e^{-\Delta t / \tau_S} + I_i^{\text{in}}(t), \tag{6.5}$$

where τ_S is the decay time constant for susceptible agents. Infection naturally decays over time as the body fights against the infection.

Infected agents ($S_i(t) = 1$) do not accumulate additional infectious matter from neighbours, but instead follow

$$I_i(t + \Delta t) = I_i(t) \cdot e^{-\Delta t / \tau_I}, \tag{6.6}$$

where τ_I is the decay time constant for infected agents.

6.3.3 *Infection Event (Fire)*

A susceptible agent i becomes infected if its infectious matter exceeds its infection threshold:

$$\text{If } S_i(t) = 0 \text{ and } I_i(t) \geq \theta_i(t), \text{ then } S_i(t + \Delta t) = 1. \tag{6.7}$$

At the time of infection, we initialise the time since infection: $T_i(t + \Delta t) = 0$, which tracks infection duration.

6.3.4 *Recovery and Immunity Update*

Infected agents recover after a fixed recovery time τ_R (refractory period):

$$\text{If } S_i(t) = 1 \text{ and } T_i(t) \geq \tau_R, \text{ then } S_i(t + \Delta t) = 0. \tag{6.8}$$

Recovery has four effects on agents:

1. **Reset Infectious Matter**: Since the body has successfully fought the disease, we assume all traces of infection have been removed.

$$I_i(t + \Delta t) = 0.$$

2. **Increase Infection Threshold (Immunity)**: Recovery comes with the benefit of acquired immunity, which we can model easily by adjusting the threshold using the **immunity threshold factor** $f > 1$.

$$\theta_i(t + \Delta t) = \theta_i(t) \cdot f,$$

where $f > 1$ is the **immunity threshold factor** representing the increase in the threshold due to acquired immunity.

3. **Adjust Transmission Rate (Output Scaling)**: As discussed previously, we rescale the transmission rate to offset the increased threshold.

$$\beta_i(t + \Delta t) = \beta_0 \cdot \prod_{n=1}^{n_i} \frac{1}{f}.$$

Here, β_0 is the base transmission rate and n_i is the number of times agent i has been infected. As can be inferred from the product, we rescale β_0 by $\frac{1}{f}$ for each past infection.

4. **Reset Time Since Infection**: Naturally, we have to reset the infected time to zero since the agent is no longer infected.

$$T_i(t + \Delta t) = 0.$$

6.3.5 *Immunity Decay*

The increased infection threshold decays exponentially back to the base value θ_0 over time, representing loss of immunity:

$$\theta_i(t + \Delta t) = \max\left(\theta_0, \theta_i(t) \cdot e^{-\Delta t / \tau_\theta}\right), \tag{6.9}$$

where τ_θ is the immunity decay constant.

6.4 PARAMETER DEFINITIONS

The model behaviour is governed by 8 key parameters, which we list here for completeness:

- The total number of agents N,
- the base infection threshold θ_0,
- the base transmission rate β_0,
- the decay time constant of infectious matter for susceptible agents τ_S,
- the decay time constant of infectious matter for infected agents τ_I,
- the recovery time (refractory period) for infected agents τ_R,
- the immunity decay constant for the infection threshold τ_θ and
- f, the immunity threshold factor, which represents the increase in infection threshold upon recovery.

6.5 INITIAL CONDITIONS

To set up the model, we initialise a small fraction of agents with non-zero infectious matter to simulate initial infections. Except for the initially infected agents, all agent states are set to susceptible. All agents start with the base infection threshold: $\theta_i(0) = \theta_0$ and transmission rates are initialised to the base rate: $\beta_i(0) = \beta_0$.

6.6 MODEL SUMMARY

The model simulates the spread of an infectious disease through a heterogeneous population represented by a scale-free network. Agents interact by transmitting infectious matter to their neighbours, accumulating exposure over time. The integrate-and-fire mechanism allows agents to become infected when their accumulated infectious matter exceeds a threshold. The inclusion of memory through the accumulation of infectious matter and adaptive infection thresholds accounts for varying susceptibility and immunity in the population.

6.7 KEY FEATURES AND EFFECTS OF PARAMETERS

To allow for a better understanding of the model, we list some of the key dynamics of the model and how they are influenced by parameters.

- **Transmission Dynamics**:

  - Higher transmission rates (β_0) lead to faster spread.
  - We scale transmission inversely with agent degree (k_i) to prevent high-degree nodes from disproportionately influencing the spread.
- **Memory of Interactions**:
  - Previously accumulated infectious matter $I_i(t)$ provides memory of past exposures.
  - $I_i(t)$ decays exponentially to model the body's natural defence mechanisms.
- **Adaptive Immunity**:
  - The infection threshold $\theta_i(t)$ is increased after recovery to model acquired immunity.
  - This increased threshold is subject to immunity decay (τ_θ), which allows for the possibility of reinfection over time.
- **Output Scaling**:
  - We adjust transmission rates $\beta_i(t)$ after recovery to prevent agents with higher thresholds from disproportionately transmitting infectious matter.
  - The output scaling factor is the reciprocal of the product of all threshold increases, ensuring balance in the model.
- **Recovery Dynamics**:
  - The recovery time τ_R determines how long agents remain infected.
  - Longer τ_R can lead to prolonged periods of infectivity.

6.8 CALIBRATION AND RESULTS

6.8.1 *UK Covid Data*

To validate our model, we utilise real-life COVID-19 case data from England [163], when all legal COVID-19 restrictions in England officially ended [164]. This timeframe allows us to observe the natural progression of the virus without the confounding effects of government-imposed

interventions. During this period, the Omicron variant was the predominant strain in circulation [163]. This is advantageous for our modelling purposes, as it permits the use of a single-strain version of the integrate-and-fire model, simplifying the analysis while still capturing the essential dynamics of transmission.

The data on raw case numbers show daily fluctuations due to variations in testing rates across the week, with a noticeable decline towards the weekend. To address this, we use weekly averages of positive tests - this enables us to calibrate the model based on consistent and representative case numbers.

To guide our parameter selection, we adopt a structured and informed approach. We start with a baseline model that includes only the most essential agent parameters: the infection threshold and recovery time. At this stage, we set both the transmission rate and the threshold increase factor to 1, and we exclude any decay processes to keep the model as simple as possible. From this starting point, we then iteratively introduce and adjust the optional parameters, focusing on those that have the most significant impact on case numbers. This helps reduce the computational burden while still optimising the parameters critical for replicating real-life dynamics. We initialise each simulation with a small seed amount of infectious matter, so that the infection can spread organically through the simulated population. This approach avoids imposing arbitrary initial conditions that may not align with reality.

Once a preliminary set of parameters yields results that align closely with the real-life data, we perform a grid search to fine-tune each parameter. The objective is to minimise the root-mean-squared error (RMSE) between the simulation output and the actual case data. The outcome of the calibration is summarised in Tab. 6.1.

Network Parameters	
Number of agents N	2,500
Power-law degree distribution exponent m	3
Agent Parameters	
Infection threshold θ_0	100
Recovery time τ_R	3
Transmission rate β_0	1
Infectious matter decay constant (infected agents) τ_I	6
Threshold factor (immunity) f	1.6
Immunity decay constant τ_θ	35
Initial Conditions	
Initial percentage of agents with infectious matter	5%
Percentage of threshold initially injected into agents	0.01%

TABLE 6.1: Optimised model parameters for matching COVID-19 case data in England. Note that all agents share the same parameters, reflecting a homogeneous population in the model.

We run the simulation on a scale-free network with 2,500 agents, which balances the need for sufficient case numbers (reducing statistical errors) and manageable computational time. The power-law degree distribution exponent of 3 is chosen to reflect the heterogeneous connectivity observed in real social networks [162]. The large infection threshold value is chosen to reduce computational rounding errors during the summation of infectious matter. To account for stochastic variations in the network generation and initial infection placement, we average the results over 100 simulation runs. This ensures that our findings are robust and not unduly influenced by specific random configurations. This is especially important considering the observation in earlier chapters that model behaviour can change dramatically with small changes in network structure.

6.8.2 *Simulation Results and Analysis*

Temporal Dynamics of New Cases

Figure 6.1 illustrates the number of new COVID-19 cases over time, comparing the simulation results with the actual data from England [163]. The model demonstrates a remarkable alignment with the real-world data, achieving an RMSE of 0.014%. This high degree of accuracy is noteworthy, especially considering that the model assumes homogeneity among agents regarding parameters such as recovery time and infection threshold.

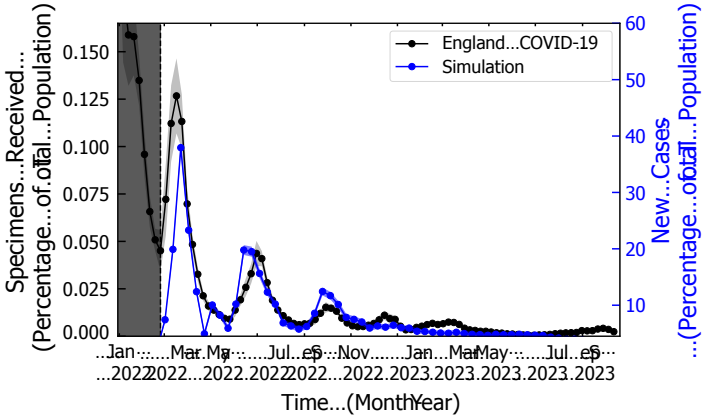


FIGURE 6.1: Comparison of real-life COVID-19 cases in England (black curve) with new cases from the integrate-and-fire model simulation (blue curve), expressed as a percentage of the population. The simulation parameters are detailed in Tab. 6.1. Shaded regions represent standard deviations from averaging (light-grey for data, blue for simulation). The dark-grey area indicates the period when government restrictions were active, ending on February 24, 2022 [164].

As depicted in Figure 6.1, the simulation successfully captures the secondary and tertiary peaks observed in the real data. This shows that the model has the ability to incorporate memory effects related to previous interactions and immunity. Traditional epidemiological models often fail to reproduce such multiple peaks without external interventions. A key advantage of the integrate-and-fire approach is that these memory effects come naturally with the model, even in this very simple form. While the overall fit is excellent, we observe a few discrepancies between model and data:

INITIAL PEAK WIDTH The initial peak in the real data is broader than in the simulation. This may be attributed to delayed network connectivity as people gradually resume normal activities post-restrictions.

LATER PEAKS The fourth and fifth peaks are less accurately captured. Possible reasons include the model's fixed population size (not accounting for births) and the lack of consideration for infections introduced via international travel.

CASE PROPORTIONS The simulation indicates a fraction of new cases approximately 100 times larger than observed. This suggests a higher transmission rate within the model, implying that while the model captures the general trends, absolute values may need scaling. With a larger number of agents, we would expect to be able to limit the necessity for rescaling. However, this would have required computational efforts outside of the scope of the project.

To align the simulation timeline with the real data, we scale the simulation time steps by a factor of 10.33, effectively mapping each simulation time step to approximately 10.33 days. Given the homogeneity of agent parameters, these scaled values can be interpreted as average characteristics of the population, as shown in Tab. 6.2.

Parameter	**Susceptible Agents**	**Infected Agents**
Recovery time	–	31.00 days
Infectious matter decay constant	–	62.00 days
Immunity decay constant	361.67 days	361.67 days

TABLE 6.2: Scaled time parameters in the COVID-19 simulation, where 1 time step equates to 10.33 days.

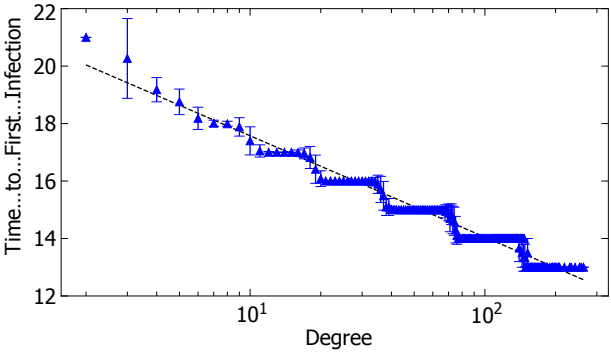


FIGURE 6.2: Semilog plot of time to first infection versus degree. Each point represents the mean and standard deviation from 100 simulation runs. The dashed line is a linear fit, indicating an inverse relationship between degree and time to infection.

Comparing these parameters to biological data, we find that our fit produces results close to reality. When comparing the recovery time, reported values for COVID-19 range from 5 to 30 days [165–169], which is close to our model's 31-day recovery period.

For the Immunity Decay, studies suggest that immunity protection decreases exponentially, with a $1/e$ reduction over approximately 275 days [170]. Our model's immunity decay constant of 361.67 days is within a reasonable range.

These comparisons validate the plausibility of the model parameters and suggest that the integrate-and-fire model can be tuned using biological data to predict case numbers.

Infection Spread and Network Degree

An important aspect of our analysis is the relationship between an agent's network degree (number of connections) and the time to first infection. Figure 6.2 presents a semilog plot illustrating this relationship.

The analysis reveals that agents with higher degrees become infected earlier. This is not surprising, since we would expect that highly socialising individuals play a significant role in propagating the infection. The linear regression on the semilog plot yields:

$$\text{Time to Infection} \approx \exp(3.0982) \cdot k^{-0.0995} \tag{6.10}$$

where k is the agent's degree. Scaling by the time step factor, we can predict infection times based on connectivity.

Visualisation of Infection Spread

To further illustrate the dynamics of infection spread, Fig. 6.3 provides snapshots of the network at different time steps. Nodes are sized proportional to their degree, and colour-coded to represent their infection status.

The visualisation confirms that high-degree nodes are among the first to become infected and act as hubs facilitating the spread to less-connected agents.

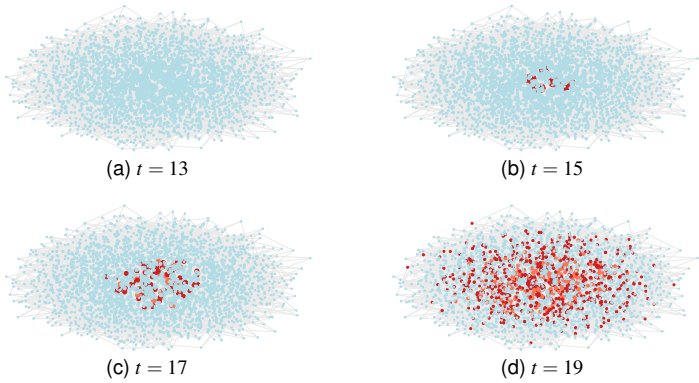


FIGURE 6.3: Snapshots of infection spreading through the network over time. Nodes represent agents, with size proportional to degree. Blue nodes are susceptible, while shades of red indicate infected agents, with darker shades representing more recently infected individuals.

6.8.3 *Impact of Interventions*

Methodology for Intervention Analysis

To evaluate potential control measures, we introduce interventions into the model based on the parameters established in Tab. 6.1. We implement interventions starting at time step 16, simulating policy responses to rising infection rates. Recognising that interventions are typically temporary, we stop them after time step 65. This is where one of the core strengths of our model shines: the capacity to remember previous interactions and immunity levels allows us to study not only the immediate effects of interventions but also the consequences upon their removal.

Quarantine Strategies

We consider two levels of quarantine:

ISOLATING INFECTED AGENTS Upon becoming infected, agents are immediately isolated, rendering them inactive in terms of transmission. They remain isolated for the duration of their recovery time.

ISOLATING INFECTED AGENTS AND THEIR NEIGHBOURS Extending the quarantine to include an infected agent's immediate contacts, akin to "test and trace" programs [171], effectively removes clusters from the network temporarily.

Figure 6.4 illustrates these quarantine strategies within a network segment.

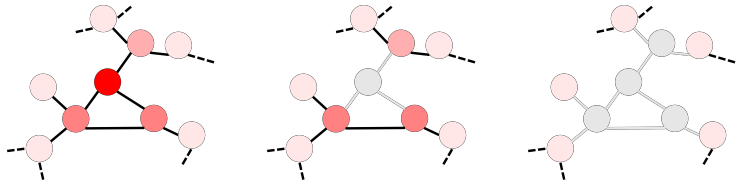


FIGURE 6.4: Visualisation of quarantine interventions: (a) No intervention; (b) Isolation of infected agents; (c) Isolation of infected agents and their immediate neighbours. Nodes represent agents, with red shades indicating levels of infectious matter.

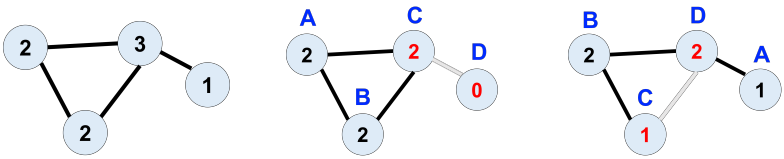


FIGURE 6.5: Effect of gathering restrictions on network connectivity: (a) No intervention; (b) and (c) Gathering restriction limiting interactions to 2 contacts. Nodes are marked in red if their degree is reduced due to the restriction. Grey edges represent inactive connections.

Gathering Restrictions

We also model social gathering restrictions, such as limiting the number of individuals in a meeting (e.g., the "rule of six" in England [172]). This is implemented by capping the number of contacts an agent can interact with at each time step. The algorithm ensures that both agents in a potential interaction do not exceed their allowed number of contacts.

Randomisation is introduced in the order of agent interactions to prevent biases. Figure 6.5 demonstrates the effect of gathering restrictions on network connectivity.

Results of Interventions

Figure 6.6 presents the impact of different interventions on the number of new cases over time.

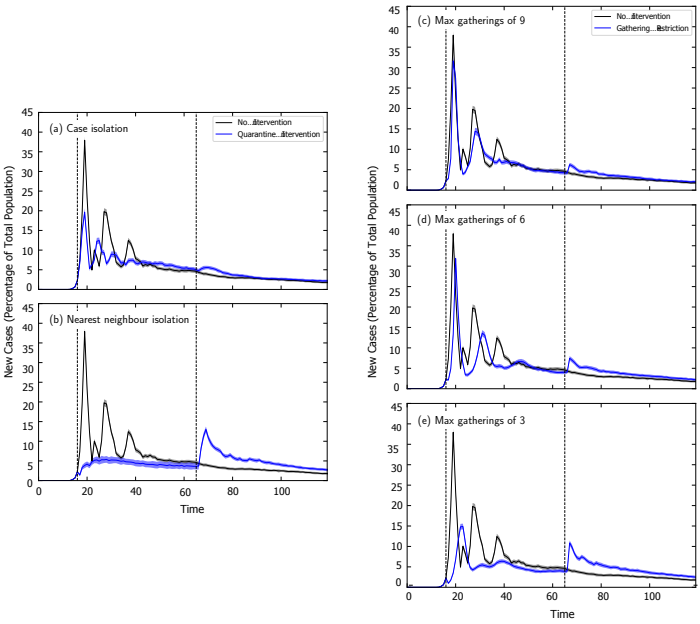


FIGURE 6.6: Comparison of new COVID-19 cases over time under different intervention scenarios.

Key observations include:

- **Flattening the Curve**: All interventions reduce the height of the initial peak, demonstrating their effectiveness in slowing the spread.
- **Delayed Herd Immunity**: Stronger interventions lead to higher secondary peaks after the interventions are lifted, indicating a delay in achieving herd immunity.
- **Cumulative Cases**: The total number of cases converges to similar values across scenarios, suggesting that interventions redistribute cases over time rather than reducing the overall number.
- **Delay of Peak**: Gathering restrictions notably delay the timing of the peak, as the reduced network connectivity slows down transmission.

These results underscore the importance of considering both short-term and long-term effects when designing public health interventions.

6.9 MODELLING MULTIPLE STRAINS WITH THE INTEGRATE-AND-FIRE MODEL

6.9.1 *Introduction*

The single-strain integrate-and-fire model, as previously discussed, provides valuable insights into the spread of infectious diseases like COVID-19. However, many real-world scenarios involve multiple strains of a pathogen co-circulating within a population. These strains may interact, compete, and influence each other's spread due to factors like cross-immunity and differential transmissibility. To capture these complexities, it is necessary to extend the single-strain model into a multi-strain framework. This allows us to study the interplay and competition of multiple strains.

6.9.2 *Model Implementation*

Extension to Multiple Strains

The multi-strain model considers two strains, which we will call Strain A and Strain B, though the framework is generalisable to any number of strains. Each strain propagates independently through the network, with agents accumulating infectious matter for each strain separately. This means that each agent collects, or "integrates" the infectious matter received from neighbours for each strain independently, similar to different opinions in the integrate-and-fire trader model.

Strain-Specific Parameters

Each strain possesses its own set of parameters:

- **Infection Threshold** (θ_i)
- **Recovery Time** (T_i)
- **Transmission Rate** (β_i)
- **Infectious Matter Decay Constant** (γ_i)

These parameters allow for modelling differences in transmissibility, virulence, and recovery between strains.

Co-infection and Infection Dynamics

In biological systems, simultaneous co-infection with multiple strains is relatively rare [173]. To reflect this, when an agent's infectious matter exceeds the threshold for any strain, the agent becomes infected with that strain and ceases to accumulate infectious matter for all strains. This cessation represents the physiological changes that prevent concurrent infections. If thresholds for multiple strains are exceeded simultaneously – a rare event – we designate the agent as co-infected. In this case, we assign the longer recovery time of the two strains to avoid bias. Upon recovery, the agent's infectious matter for all strains is reset to zero.

Cross-Immunity Mechanism

Recovery from one strain can allow agents to gain partial immunity to other strains, a phenomenon known as cross-immunity [174]. To model this, we introduce a threshold increase factor matrix F, which modifies the infection thresholds post-recovery:

$$F = \begin{pmatrix} F_{AA} & F_{AB} \\ F_{BA} & F_{BB} \end{pmatrix}, \tag{6.11}$$

where F_{ij} represents the factor by which the infection threshold for strain j increases after recovery from strain i. For example, an agent recovering from Strain A will have its threshold for Strain A increased by F_{AA} and for Strain B by F_{AB}. This mechanism ensures that agents return to a single susceptible state but with modified vulnerability to each strain, depending on their infection history.

6.9.3 *Application to Influenza Spread in Canada*

Selection of Real-life Data

To validate the multi-strain model, we apply it to the spread of influenza in Canada. Influenza is an ideal case study due to its transmission via close human contact [175] and the co-circulation of two primary strains: Influenza A and Influenza B [176]. Influenza A is generally more prevalent and is known to mutate rapidly, leading to annual infection cycles [177].

We utilise positive influenza specimen data from Canada's FluWatch program [178], focusing on the period before 2020 to avoid anomalies introduced by the COVID-19 pandemic and related government interventions.

Model Calibration and Parameter Tuning

INITIAL PARAMETER SETTINGS We begin by assigning identical parameters to both strains, leveraging insights from the single-strain model. The only initial difference is the transmission rate, with Strain A being more contagious and having a longer recovery time to reflect its higher prevalence and virulence.

THRESHOLD INCREASE FACTOR MATRIX For simplicity, we adopt a symmetric threshold increase factor matrix:

$$F = \begin{pmatrix} F_{AA} & F_{AB} \\ F_{BA} & F_{BB} \end{pmatrix} = \begin{pmatrix} 5 & 3 \\ 3 & 5 \end{pmatrix}. \tag{6.12}$$

This matrix models moderate cross-immunity between the strains, with stronger immunity to the same strain post-recovery.

INITIAL CONDITIONS To test the model's ability to replicate the dominance of Strain A over Strain B, we intentionally introduce a higher initial amount of infectious matter for Strain B. This setup challenges the model to simulate Strain A overtaking Strain B in prevalence, consistent with real-world observations.

PARAMETER OPTIMIZATION We refine the parameters through iterative visual comparisons and grid search optimisation, minimising the root-mean-squared error between the simulation outputs and the real-life data. The finalised parameters are listed in Tab. 6.3.

Network Parameters		
Number of agents N	2,500	
Power-law degree distribution exponent m	3	
Agent Parameters	**Strain A**	**Strain B**
Infection threshold θ_0	100	100
Recovery time τ_R	40	20
Transmission rate β_0	1.00	0.965
Infectious matter decay constant τ_I	4	4
Threshold factor (immunity) f	$\begin{pmatrix} 5 & 3 \\ 3 & 5 \end{pmatrix}$	
Immunity decay constant τ_θ	52	52
Initial Conditions	**Strain A**	**Strain B**
Initial percentage of agents with infectious matter	0.5%	0.5%
Percentage of threshold initially injected	1%	10%

TABLE 6.3: Optimised parameters for modelling Influenza Strain A and Strain B in Canada. All agents share the same parameters, and certain parameters are identical for both strains for consistency.

TIME SCALING To align the simulation timeline with the real-world data, we scale each simulation time step to represent approximately 10.60 days.

INTERPRETATION OF RECOVERY TIME The recovery times in the model (424 days for Strain A and 212 days for Strain B after scaling) are significantly longer than the typical 3–7 days observed clinically [179]. This discrepancy arises because the model's infectious matter decay constant is much shorter than the recovery time, causing the agent's infectious matter to diminish rapidly – analogous to clinical recovery – even though the agent remains in the "infected" state in the model. This is probably a calibration artefact, so that both clinical recovery and temporary immunity are covered by this recovery time, when it is actually just the time during which reinfection is impossible.

6.9.4 *Results and Discussion*

Comparison with Real-life Data

Figure 6.7 presents the simulation results alongside the actual influenza case data from Canada.

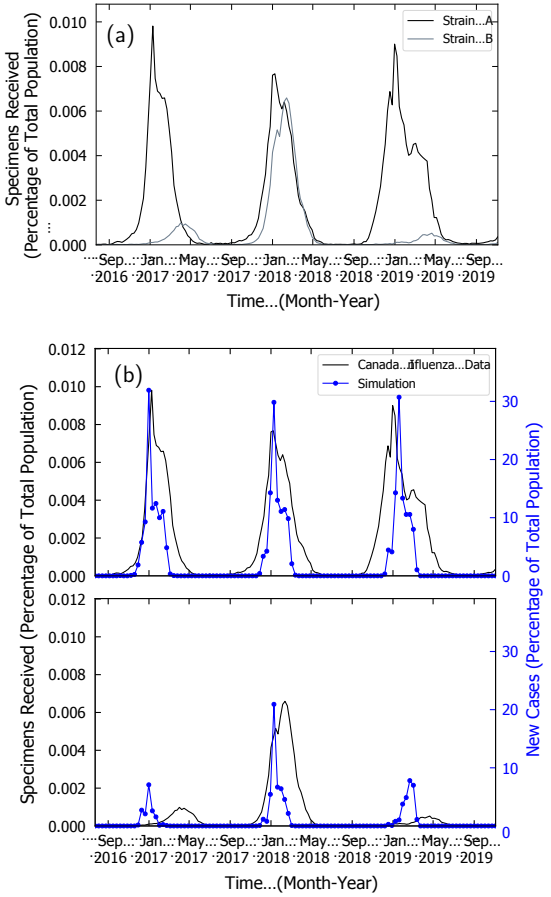


FIGURE 6.7: (a) Observed influenza cases in Canada for Strain A (black curve) and Strain B (grey curve) [178]. (b) Simulation results for Strain A (blue curve, top graph) and Strain B (blue curve, bottom graph) using parameters from Tab. 6.3. The simulation curves are extracted from a single run aligned with the real-life data. The scaling factor is the same for both strains.

The model successfully replicates the alternating peaks of Strain A and Strain B, with RMSE values of 0.002% and 0.0012%, respectively. The synchronisation of peaks and the dominance of Strain A align with empirical observations.

Analysis of Discrepancies

PEAK WIDTHS The simulation peaks are narrower than those observed in the real data. This suggests that the simulated network is more densely connected than the actual social network in Canada, leading to faster spread. In Sec. 6.8.3, we noted that sparser networks result in broader peaks due to slower transmission.

POPULATION DENSITY CONSIDERATIONS Canada's population distribution is highly uneven [180], with a much lower overall population density (4 people/km² [181]) compared to England (400 people/km² [182]). This disparity suggests that the power-law degree distribution exponent used may not accurately reflect Canada's social network structure. Alternative network models, such as random partition graphs, might better capture regional variations and connectivity patterns.

PEAK TIMING OFFSET In the real data, Strain B peaks after Strain A, whereas the simulation shows both strains peaking simultaneously. This synchronisation likely results from both strains following identical transmission dynamics and reaching highly connected nodes simultaneously. Additionally, the assumption that agents stop accumulating infectious matter for all strains when infected with any strain may introduce unwanted coupling between strains.

Potential Model Enhancements

To address the identified discrepancies, several modifications could be considered:

- **Adjusting Network Structure**: Utilising a network model that accounts for regional population densities and connectivity variations may produce more accurate peak widths and transmission rates.
- **Staggered Strain Introduction**: Introducing Strain A into the system after Strain B has begun spreading could recreate the observed offset in peak timings.

- **Modifying Co-infection Dynamics**: Allowing agents to continue accumulating infectious matter for other strains after infection with one strain, or retaining a fraction of infectious matter upon recovery, could reduce unwanted coupling and better represent independent strain dynamics.

In-depth Analysis of Agent Dynamics

To further understand the model's behaviour, we examine the infection progression in a randomly selected agent and its immediate neighbours (Fig. 6.8).

The agent consistently becomes infected with Strain A before its neighbours, which means that it is a potential super-spreader within its cluster. The steeper slope for Strain A reflects its higher transmission rate. The similarity in the rise and fall patterns for both strains emphasises the synchronisation issue previously noted.

Average Infection Probability

By averaging the ratio of infectious matter to threshold across all agents, we obtain a measure analogous to the average infection probability $\bar{I}(t)$ (Fig. 6.9).

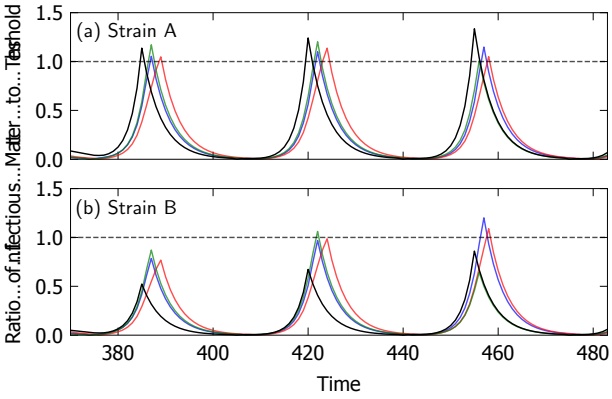


FIGURE 6.8: Ratio of infectious matter to infection threshold over time for a randomly selected agent (black curve) and its nearest neighbours (red, blue, and green curves). An agent becomes infected when the ratio exceeds unity.

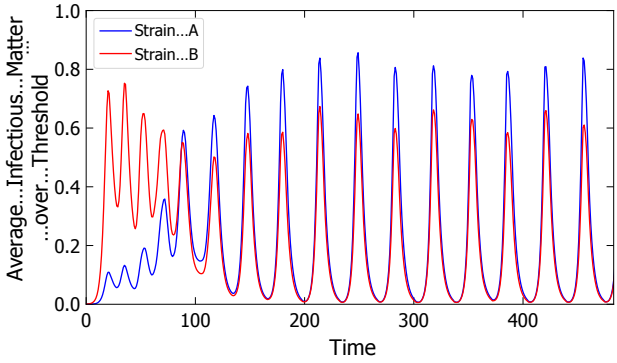


FIGURE 6.9: Average ratio of infectious matter to infection threshold across the population over time for Strain A (blue curve) and Strain B (red curve).

Initially, agents are more likely to be infected with Strain B due to its higher initial presence. However, Strain A rapidly becomes dominant, consistent with its higher transmissibility and prevalence in real life [176]. The simultaneous peaking of both strains suggests that modifying the timing of strain introductions could improve the model's alignment with observed data.

6.10 CONCLUSION AND FUTURE WORK

Building on the integrated framework for social interaction introduced in the previous chapter, this work extended the integrate-and-fire approach to epidemiological modelling. By treating infection transmission as a threshold-based process – similar to neuronal firing – we captured memory effects in disease spread. Agents accumulate "infectious matter" over time, become infected upon exceeding an infection threshold, and develop immunity by effectively raising this threshold after recovery. This approach brings to epidemiology some of the same advantages observed in opinion and market models: the natural emergence of non-trivial dynamics, such as secondary peaks in infection curves, without invoking external intervention measures.

Through systematic simulations, the model reproduced real-world disease patterns, including secondary waves of COVID-19 in England and seasonal influenza dynamics in Canada. We showed that the model is able to fit actual case data to a low margin of error. This underscores the

wide applicability of the integrate-and-fire model. Network structure, a key element of the general framework introduced in Chap. 5 also plays a crucial role in this model. Highly connected individuals (compare with the highly connected and influential traders in Chap. 5), became early infection hubs. This, once again, demonstrates the importance of network topology in shaping complex dynamics.

Moreover, the model's reaction to parameter changes revealed how certain biological or behavioural rates, such as infectious matter decay or recovery time, govern the system's steady states. Notably, constraints and memory – elements we previously discussed in the context of opinion formation and trading – proved equally essential here. Just as financial constraints and past actions influenced trader decisions, the accumulation of infectious matter and rising thresholds capture a memory process in epidemiology. The model's capacity to reflect how interventions can suppress immediate outbreaks but inadvertently lead to larger subsequent waves further highlights the importance of considering underlying memory and structural factors.

Despite its strengths, several limitations remain. We observed synchronisation of multiple pathogen strains, which contradicts typical epidemiological patterns, and the current network assumptions may not fully capture the complexities of real-world social interactions. The model also lacks consideration of demographic changes, heterogeneous agent parameters, and open-system dynamics, all of which warrant further exploration. Future work should introduce hospitalisation states, variable population sizes, travel-related infections, and vaccination strategies. Similarly, refining multiple-strain interactions, varying agent heterogeneity, and exploring alternative network topologies would enhance the model's applicability and alignment with empirical data.

In conclusion, this integrate-and-fire epidemiological model represents a significant step toward bridging insights from complex decision-making and opinion formation to the spread of infectious diseases. Just as memory and constraints were central to explaining non-linear opinion dynamics or market instability, they also enable more accurate and nuanced disease modelling. By integrating these features with network structures and multiple strains, the model moves closer to the inherently complex, interconnected reality of public health challenges. The approach presents a promising

tool for researchers and policymakers seeking to comprehend and manage infectious diseases in a highly interconnected world.

7

SUMMARY, CONCLUSION AND OUTLOOK

For centuries, the collective mindset of the general population has presented a significant challenge to scholars, political analysts, and rulers alike. Despite advancements in social and political science, economics, and increased data collection and analysis capacity, public opinion still mostly remains a mystery. There have been many instances in history, where political leaders were blind-sighted by sudden, unexpected shifts in popular sentiment. Many seemingly stable regimes, confident in their ability to understand and control their citizens, found themselves confronted with rapidly escalating protests, which sometimes turned into full-scale rebellions. In a similar fashion, underlying cultural tensions can suddenly erupt into sweeping social revolutions.

These episodes of seemingly abrupt societal upset highlight the intricate and often concealed nature of mass sentiment. Economic hardship, cultural tensions, and long-standing grievances can simmer beneath the surface for years, largely unnoticed or willfully ignored by those in power, until they suddenly ignite. In these pivotal instances, long-held assumptions about public opinion implode, leaving leaders struggling to regain their footing in the face of forces they neither anticipated nor fully understood.

This uncertainty has plagued not just politicians who had to contend with rapidly changing sentiment, but economists and financial researchers as well. Opinions in markets are not understood to this day, with sudden mood swings surprising many experienced professionals. Research in complexity science has shown that crashes in markets can be likened to critical points in statistical mechanics, thus marking a strong argument for the involvement of physicists in the study of market opinion. Physicists have spent centuries developing methods to model physical systems, and the idea of viewing systems through the lens of correlations, critical points and agent-based simulations have been tried and tested for many applications. Many of the previous approaches model traders akin to particles that are subject to physical forces and often inhibit imitative behaviour. These approaches

range from the diffusion of opinion to mean-field-like models with discrete states. However, these approaches usually have one thing in common: they take an outside view of the trader in their attempt to find an appropriate description for the traders behavioural changes.

These approaches have shown success in replicating stylized facts of financial markets – however they are inherently limited due to their approach: Centuries in research in physics has demonstrated that for truly accurate prediciton and understanding it is necessary to understand the inner working of the system. These inner workings can then be abstracted again towards a bigger picture.

Economists have understood this early, which is why most economic models start from an understanding of the micro-scale and the individual and then generalize these findings. By assuming rational behaviour, economic models of equilibrium can be built where the supposedly accurate description of the individual can relate exactly to large-scale behaviour. Unfortunately, an entirely rational description of individual behaviour is not accurate, even if we were to ignore the problem of information availability and processing capabilities.

As soon as we try to extend the economic models built on the assumption of rational agent behaviour with irrationality, these models stop being exact immediately. Only in the case of a purely rational actor can we rely on the accurate description of individual motivations to explain macroscopic behaviour. By adding external terms to correct for this, we immediately take an outside perspective of traders and try to amend an inherently inaccurate model by adding correction terms.

What is needed instead is a description of behaviour that takes an inside perspective. We need to understand the decision-making processes of the individual in order to even dream of coming towards an accurate description of market behaviour. Studies in neuroscience and behavioural economics have demonstrated the enormous effect that information processing on a neuronal level has on decision processes. If we want to find good descriptions of trader behaviour, we need to start with a truly bottom-up approach, which does not start with the individual trader in relation to other traders, but with the decision process within that trader, where different opinions and options compete with each other.

This is what was attempted in this thesis. Starting from a comprehensive review of the fruitful relationship that economics and physics had over the centuries, we concluded that many of the purely statistical descriptions of markets, which assume homogeneity, rationality and perfect information, are inadequate. Agent-based models have been put forward as a valuable tool for modelling bottom-up approaches and can alleviate many of the issues with classical economics. However, a closer look at the decision-making processes and neuroeconomic research has led us towards a new approach: Incorporating neuronal-level decision-making into an agent and then building agent-based models on top of that.

The model that we developed combined economic intuition from established agent-based models like the Lux-Marchesi models with a model for neurons, the integrate-and-fire model. By adapting the model towards a model of opinions, we were able to show that it allows us to apply the Leaky Competing Accumulator model developed by Usher and McClelland into an agent-based framework. After developing the model, we first analysed its dependence on network and parameter values and showed that it was able to replicate a wide range of behaviour. However, the truly fascinating insights came when we started to incorporate some of the neuronal limits on decision-making into the model: By introducing inhibition and decay, we could directly show how irrational behaviour on a neuronal scale translates towards stylized facts of the financial markets like the emergence of bubbles and fat-tailed distributions. This is due to the inherent incapability of humans to make completely accurate decisions. This approach allows us to refute the efficient market hypothesis with a model built directly on neural-level decision processes.

In a similar fashion, we were able to show that the asymmetric access to information leads towards another stylized fact, volatility clustering. The core of the integrate-and-fire model is the transfer and processing of information, which is why it is so suitable to investigate the effects of information hierarchies. We can incorporate these into the model without having to change anything about the fundamental agent behaviour; we simply limit the information that agents have access to.

After demonstrating that a single-stock version of the integrate and fire trader model allows us to draw insights and analyze the root causes of

stylized facts found in financial markets, we then generalized the model into a general framework of opinion modelling. This general framework allows us to apply the integrate and fire model towards any group of interacting individuals and can model their opinion dynamics. Building on this framework, we developed a multi-asset model including an order book and showed that the model in this more elaborate form can reproduce many stylized facts of financial time series, in addition to wealth and order size distributions.

Finally, in the last chapter, we were able to show that the integrate-and-fire model can be applied to real-world systems and model infectious processes with high accuracy. We were able to reproduce the case numbers of the UK COVID-19 pandemic with almost no error, even matching the biological parameters. This is especially remarkable given that the model used was relatively simple and has only a few parameters. We were also able to capture patterns like secondary peaks, as well as the delay of peaks due to interventions.

7.1 FUTURE WORK

It is important to emphasize that this thesis is only the first attempt to introduce a neuronal level of decision-making into agent-based modelling. Despite us being able to reproduce many of the stylized facts found in financial markets, work has just begun in turning this into a full description of financial systems. Instead, this work should be understood as a proof of concept on which future work can build. Enhancements to this work are needed both in the fundamental understanding of markets and decision processes, as well as on the applications, where the model can be adapted to draw new insights.

7.1.1 *Fundamentals*

There are various areas where the fundamental aspect of the model could be improved. We have spent considerable time analysing the behaviour of the single-stock model and have been able to relate stylized facts to certain aspects of neuronal decision-making. Additional work could go deeper into the interplay of inhibition and hierarchies and try to find even more accurate descriptions of both.

Besides that, we might be able to relate even more stylized facts to certain aspects of neuronal decision-making, which might prove additional insights. We have observed how action potentials allow us an inside-view of traders decision processes, essentially providing an x-ray vision of the market. When combined with additional analysis tools, even deeper insights might be gained.

An even more fundamental, albeit very unusual, approach would be to investigate whether the concept of consciousness could be applied to traders in financial markets. Frameworks like the integrated information theory of consciousness have introduced tools to measure the emergence of consciousness in systems. These measures usually rely on measuring the integration of a system versus the outside world. While these tools would usually negate the existence of consciousness between different humans, we might be able to make an argument for why it could be applied to markets.

Traders in financial markets interact in real time and share their information and perception of the world with others through their trades. Collectively, the market encompasses the sum of all of the beliefs and information of traders and facilitates communication of these entities. Since it is impossible to directly observe the states of traders in the real world, applying the Phi-measure of Integrated Information Theory to the integrate-and-fire trader model might allow us to investigate whether any concept of shared consciousness might exist in markets. The Phi-measure calculated the maximum amount of integrated information in a system and is used as proxy to measure the emergence of conciousness [183].

7.1.2 *Calibration*

So far, most of our analysis for financial markets has focused on replicating stylized facts. This was mainly due to data availability, and calibrating the model could provide meaningful insights into the markets. We have seen in Chap. 6 that the integrate-and-fire model can calibrate to real-world systems with high accuracy, so it should also be possible for markets, given enough data.

Especially the use of machine learning has the potential to be very beneficial for model calibration. Concepts like surrogate models [184] could help speed up calibration runs, especially if we train neural networks and utilize

backpropagation. In addition to that, machine learning could prove useful in another part of model calibration: Natural language processing capabilities allow the real-time conversion of news and social media posts into measures of sentiment. These sentiments can easily be converted to currents, which means that the model could both be calibrated, as well as run on real-time market information, which would greatly improve its applicability.

This calibration would be particularly useful when applied to the multi-stock model, which can account for the complex relationships between multiple assets within a market. The multi-stock model, in general, could benefit from further analysis, especially in understanding order book dynamics and the role of traders within a network on their influence. Investigating the information flows within the model, especially when combined with hierarchical information structures, could also help uncover the effects of insider trading by adding edges between different information silos and observing market reactions.

7.1.3 *Additonal applications*

There are many additional applications of this model, in addition to just looking at traders. We have seen the applicability of the integrate-and-fire model to infectious processes for pandemics. However, it can also be applied to another infectious process: banking risk.

In the following, I present a brief overview of some work in progress that was inspired by a discussion with my friend Nishanth Kumar, who mentioned to me the fact that central bankers and financial regulators are now increasingly looking into infectious disease models like the Susceptible-Infectious-Susceptible (SIS) model to better understand and mitigate banking risks. This shift stems from the recognition that financial systems exhibit dynamics comparable to the spread of infectious diseases, where distress can rapidly propagate through interconnected institutions, leading to systemic crises. During our discussion, it occurred to us that since the integrate-and-fire model was able to be tuned to infectious diseases very effectively, similar success might be possible in the banking sector.

Just as infectious diseases spread through a host population, financial distress can propagate through networks of interconnected financial institutions, potentially causing systemic crises. Traditional risk models, often

reliant on equilibrium assumptions and siloed treatments of individual institutions, have proven insufficient in predicting and mitigating financial meltdowns. The global financial crisis of 2008 and earlier episodes, such as the 1997 Asian Financial Crisis, have illustrated that one institution's failure can quickly cascade across the entire system, much like a pathogen infecting successive hosts.

Models such as the SIS (Susceptible–Infected–Susceptible) allow us to apply insights from epidemiology to abstract concepts in finance: shocks to a single node (bank or market participant) can spread to its neighbours through credit and funding linkages. As we have seen in the previous chapter, the integrate-and-fire model can very accurately model such processes, which makes it a natural candidate to model systemic risk in global banking networks.

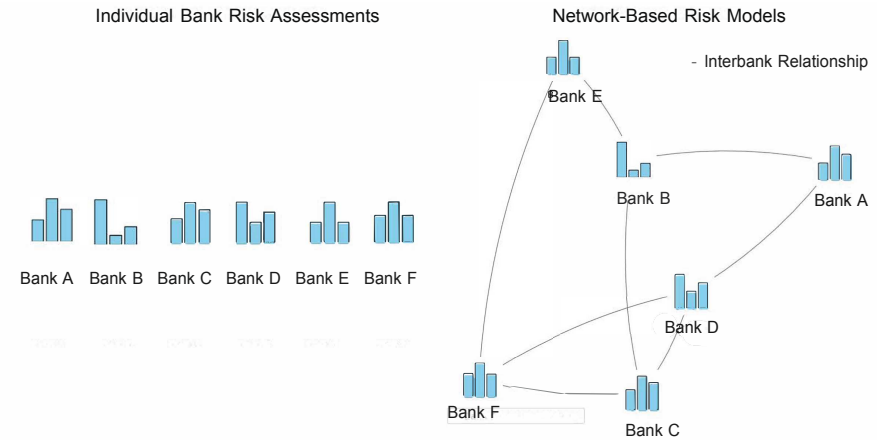


FIGURE 7.1: Illustration of risk modelling for banks as individual entities (left), or within the context of interconnected liabilities (right)

From Epidemiology to Finance: The Integrate-and-Fire Paradigm

In the integrate-and-fire model, neurons integrate incoming signals until a threshold is reached, at which point they fire. Translated into a banking context, each financial institution (analogous to a neuron) integrates (absorbs) incoming losses. Losses can be due to market shocks, counterparty defaults, and other risk factors. Once these accumulated losses exceed the bank's loss-absorbing capacity, the bank "fires" by defaulting. This default, in turn, causes losses to other banks that had exposures to it, potentially pushing them closer to their own thresholds and causing a contagion effect.

There are three key features that explain why the integrate-and-fire dynamic is useful for banking risk:

MEMORY AND ACCUMULATION Instead of treating each failure event as instantaneous and independent, the model tracks the gradual buildup of stress over time.

THRESHOLD DYNAMICS Banks fail only after their capital drops below a critical threshold, analogous to neurons firing after exceeding membrane potential thresholds.

NETWORK-DRIVEN CONTAGION The timing and magnitude of failures depend on the structure and weights of interbank exposures, allowing for heterogeneous outcomes dependent on network topology.

Determining loss thresholds

Defining the loss-absorption capacity is crucial to ensure the applicability of the model to the real work. This step is still in progress and will require collaboration with experts. We can find a starting point within the Basel-3 framework, which lines out capital requirements for banks. One example of this is the Total Loss-Absorbing Capacity (TLAC) requirement, which was established to ensure that global systemically important banks (G-SIBs) have sufficient capital and liabilities that can be written down or converted to equity in a resolution scenario.

In the integrate-and-fire model, TLAC (or analogous measures based on capital requirements) can give guidance on how to define the thresholds at which a bank fails. TLAC thus serves as a concrete parameter that calibrates

the model's firing thresholds. By adjusting TLAC and other regulatory parameters, such as minimum capital ratios under Basel III or risk-weighted asset calculations, we can simulate how different regulatory regimes influence systemic robustness. For example, a higher TLAC requirement generally increases an institution's capacity to withstand shocks, reducing the probability of cascading failures. In the proposed model, a higher TLAC would correspond to a higher loss absorption threshold, thus requiring more "distress" for the shock to propagate.

Reconstructing the Global Banking Network from BIS Data

A crucial step in this modelling framework is constructing a realistic representation of the global banking network. We draw on data from the Bank for International Settlements (BIS), which provides aggregated international banking statistics, including cross-border claims and liabilities. While BIS data is typically at the country level rather than at the bank level, we use the assumption of scale-invariance to reconstruct a typical inter-bank network.

We start with a country-level network, which is directly based on the BIS-reported exposures between countries. The network represents claims and liabilities as a directed weighted edge in a network of countries. We incorporate domestic (internal) claims and liabilities as self-loops. Missing values are estimated using statistical regressions to ensure a comprehensive data set.

Now, we disaggregate the single country node into multiple bank nodes. These bank-level subnetworks are constructed using a directed configuration model with heavy-tailed degree distributions to reflect empirical findings that a few large, highly connected banks coexist with many smaller institutions. Bank-level exposures are assigned so that the sum of interbank links within a country matches the country's domestic activity, and cross-border exposures are distributed according to statistical and topological criteria, maintaining the global claims structure from BIS. The reconstructed network can be seen in Fig. 7.2.

The resulting multi-layered network replicates known empirical features of financial systems, including scale-free degree distributions and heterogeneity in bank sizes and connectivity.

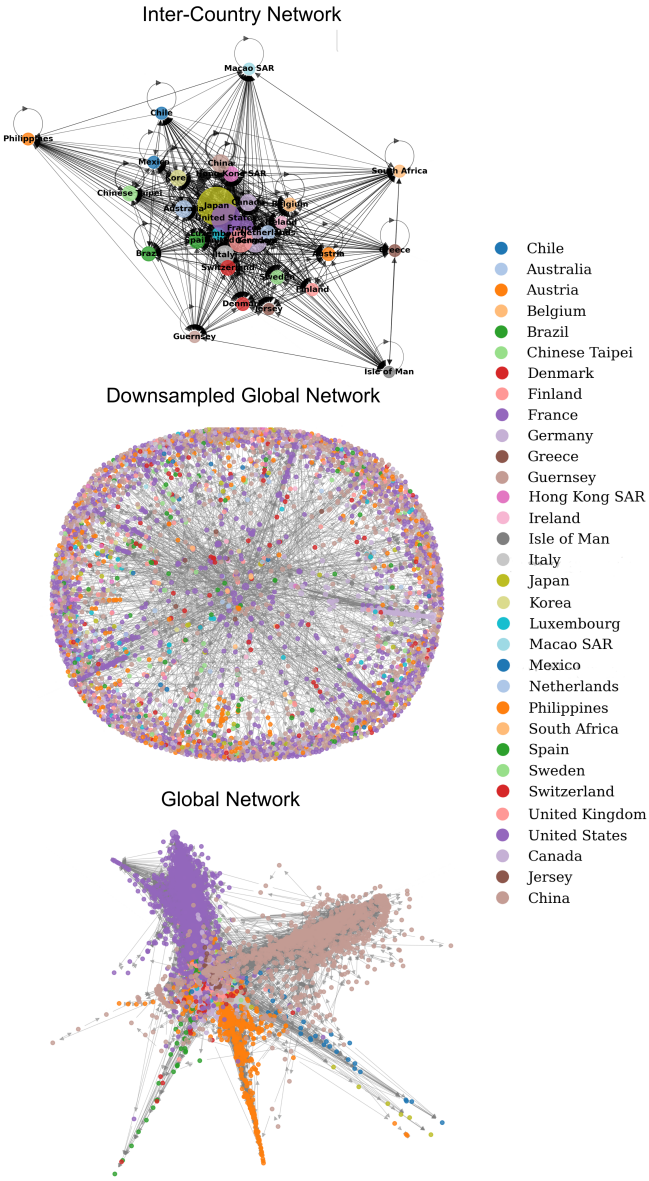


FIGURE 7.2: Top: Inter-country network of claims and liabilities created from BIS data. Middle: Upsampled country of indvidual banks, with only the top edges and nodes shown, Bottom: Plot of the same network, but now with all nodes and edges.

Modeling Shocks, Propagation, and Systemic Outcomes

Once the network is established, we simulate the evolution of bank capital and the propagation of shocks over time. The process can be summarized as follows: At each time step, we apply shock to a randomly selected subset of banks. We define risk ranking for banks, which include implied failure rates. These shocks might represent sudden loan defaults, market losses, or operational failures. Between shocks, banks adjust their capital toward their required levels, reflecting regulatory compliance (such as Basel III requirements) or the slow accumulation of retained earnings. If a banks loss-absorbing threshold is surpassed, the bank is considered defaulted.

When a bank fails, it triggers losses at other banks that are its creditors, as some or all of the failed bank's liabilities become worthless. This reduces the capital of those creditor banks, potentially causing a chain reaction of further defaults. The model thus captures second-round effects and longer cascades of failures. A range of these simulation runs for different risk requirements is depicted in Fig. 7.3.

By running the simulation multiple times with varying parameters, we can measure systemic risk indicators, such as failure rates and avalanche sizes, which measure the number of banks failing in a chain reaction. We can also create country-level risk heat maps. By aggregating all banks within a country, we can visualize which countries' banking sectors are most vulnerable to external shocks or how different regulatory settings shift the locus of risk. See Fig. 7.4 for an example.

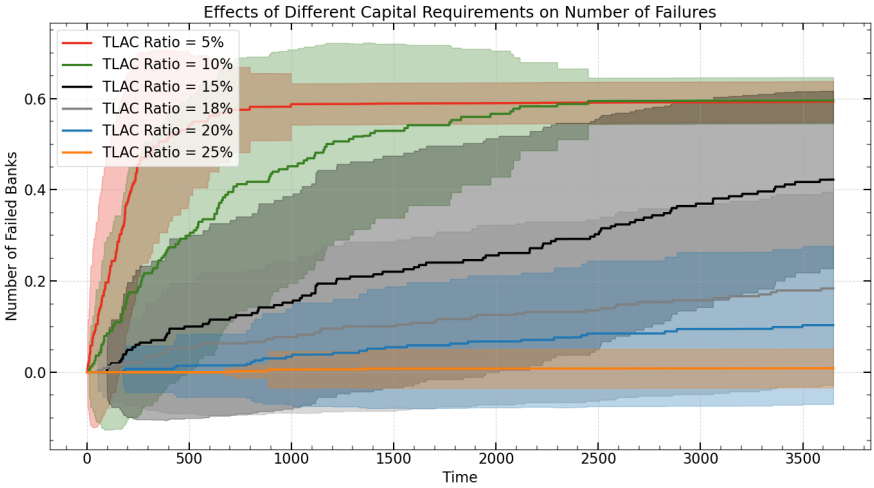


FIGURE 7.3: Fraction of failures over time for different total loss absorption capacities (TLACs). The shades represent the error bars.

Policy Analysis and Regulatory Experiments

A key advantage of this modelling framework is the ability to conduct counterfactual experiments and policy analyses. By adjusting regulatory parameters, such as CRR or TLAC levels, leverage ratios, or liquidity requirements, and re-running the simulation, we can observe how these changes affect systemic stability.

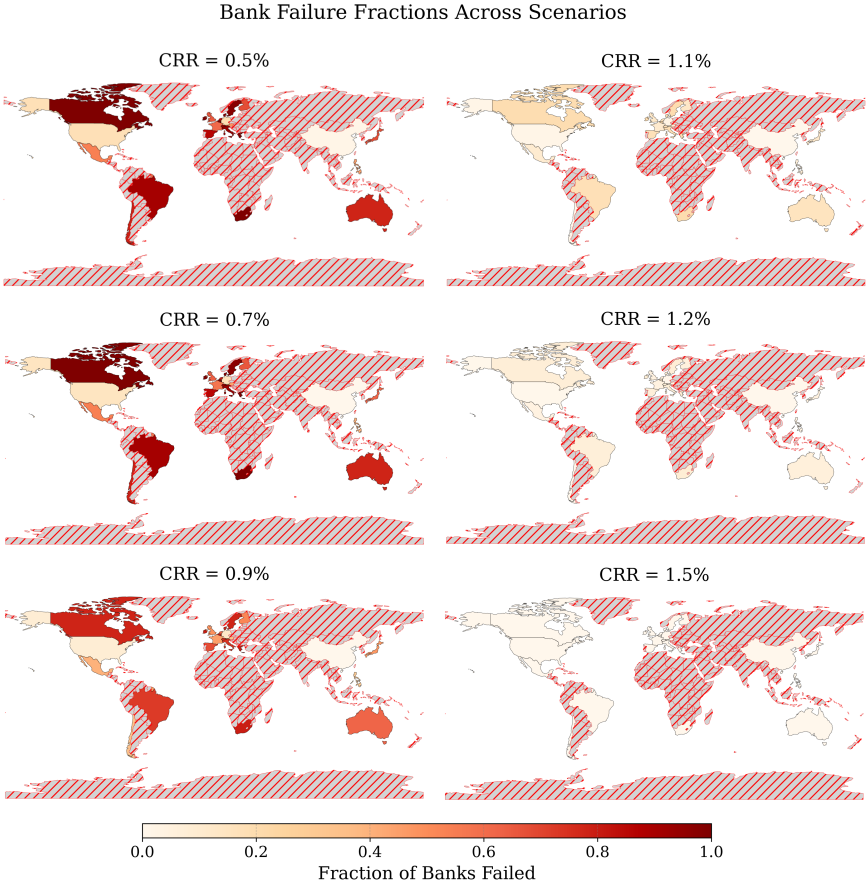


FIGURE 7.4: Heatmap of failure rates for different weighted Capital Requirement Ratios (CRR).

7.1.4 *Final comments*

The integrate-and-fire banking risk model work serves as a concept for future work based on the integrate-and-fire framework. As described in the previous chapters, the integrate-and-fire model can replicate an incredible variability of human behaviour and is applicable to real-world systems ranging from financial markets to epidemics, banking risks and others. This thesis marks just the first step into many possible future applications. When

properly calibrated, this model could serve as x-ray vision into the market, uncovering decision processes and hidden sentiments in markets and other systems.

BIBLIOGRAPHY

1. Smith, A. in. *The Glasgow Edition of the Works and Correspondence of Adam Smith, Vol. 2: An Inquiry into the Nature and Causes of the Wealth of Nations, Vol. 1* (ed Todd, W. B.) (Oxford University Press, 1776).
2. Smith, A. & Stewart, D. *The Glasgow Edition of the Works and Correspondence of Adam Smith, Vol. 3: Essays on Philosophical Subjects with Dugald Stewart's Account of Adam Smith* (eds Wightman, W. P. D., Bryce, J. C. & Ross, I. S.) (Oxford University Press, 1980).
3. Diemer, A. & Guillemin, H. L'économie politique au miroir de la physique : Adam Smith et Isaac Newton. *Revue d'histoire des sciences* **64**, 5 (2011).
4. Kim, K. Adam Smith's 'History of Astronomy' and View of Science. *Cambridge Journal of Economics* **36**, 799 (2012).
5. Newton, I. *Philosophiae naturalis principia mathematica* (Jussu Societatis Regiae ac Typis Josephi Streater. Prostat apud plures bibliopolas, 1687).
6. Bassett, G. W. The St. Petersburg Paradox and Bounded Utility. *History of Political Economy* **19**, 517 (1987).
7. Savage, L. J. The Theory of Statistical Decision. *Journal of the American Statistical Association* **46**, 55 (1951).
8. Bradley, R. Ramsey's Representation Theorem. *Dialectica* **58**, 483 (2005).
9. Maxwell, J. C. IV. On the dynamical theory of gases. *Philosophical Transactions of the Royal Society of London* **157**, 49 (1867).
10. Boltzmann, L. in Brush, S. G. *Kinetische Theorie II* 115 (Vieweg+Teubner Verlag, Wiesbaden, 1970).
11. Edgeworth, F. Y. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences* **10** (CK Paul, 1881).
12. Marshall, A. *Principles of Economics* 319 pp. (Prometheus Books, Amherst, NY, 1997).
13. Amoroso, L. Vilfredo Pareto. *Econometrica* **6**, 1 (1938).

14. Pfeffer, F. T. & Dvir-Djerassi, A. The U.S. Wealth Distribution: Off the Charts. *Socius: Sociological Research for a Dynamic World* **8**, 237802312211439 (2022).

15. Ekström, G., Nettles, M. & Dziewoński, A. The global CMT project 2004–2010: Centroid-moment tensors for 13,017 earthquakes. *Physics of the Earth and Planetary Interiors* **200–201**, 1 (2012).

16. Dziewonski, A. M., Chou, T.-A. & Woodhouse, J. H. Determination of earthquake source parameters from waveform data for studies of global and regional seismicity. *Journal of Geophysical Research: Solid Earth* **86**, 2825 (1981).

17. Office, I. M. *Whole Island - Earthquakes 28.11-29.11-2023* ().

18. Friedrich, J. M., Macke, R. J., Wignarajah, D. P., Rivers, M. L., Britt, D. T. & Ebel, D. S. Pore Size Distribution in an Uncompacted Equilibrated Ordinary Chondrite. *Planetary and Space Science* **56**, 895 (2008).

19. Marquet, P. A., Quiñones, R. A., Abades, S., Labra, F., Tognelli, M., Arim, M. & Rivadeneira, M. Scaling and Power-Laws in Ecological Systems. *Journal of Experimental Biology* **208**, 1749 (2005).

20. Albert, R. & Barabási, A.-L. Statistical Mechanics of Complex Networks. *Reviews of Modern Physics* **74**, 47 (2002).

21. Thompson, K. J., Bagnato, G. G., Telles, G. D., Caracanhas, M. A., Dos Santos, F. E. A. & Bagnato, V. S. Evidence of Power Law Behavior in the Momentum Distribution of a Turbulent Trapped Bose–Einstein Condensate. *Laser Physics Letters* **11**, 015501 (2014).

22. Einstein, A. Über Die von Der Molekularkinetischen Theorie Der Wärme Geforderte Bewegung von in Ruhenden Flüssigkeiten Suspendierten Teilchen. *Annalen der Physik* **322**, 549 (1905).

23. Brown, R. XXVII. *A Brief Account of Microscopical Observations Made in the Months of June, July and August* 1827, *on the Particles Contained in the Pollen of Plants; and on the General Existence of Active Molecules in Organic and Inorganic Bodies*. *The Philosophical Magazine* **4**, 161 (1828).

24. Bachelier, L. Théorie de La Spéculation. *Annales scientifiques de l'École normale supérieure* **17**, 21 (1900).

25. Black, F. & Scholes, M. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* **81**, 637 (1973).

26. Merton, R. C. Theory of Rational Option Pricing. *The Bell Journal of Economics and Management Science* **4**, 141 (Spr. 1973).

27. *Coping with the Complexity of Economics* (eds Faggini, M. & Lux, T.) (Springer Milan, Milano, 2009).

28. Jensen, H. J. *Complexity Science: The Study of Emergence* (Cambridge University Press, Cambridge, United Kingdom ; New York, NY, 2023).

29. Waldrop, M. M. *Complexity: The Emerging Science at the Edge of Order and Chaos* (Simon and Schuster, 1993).

30. Mas-Colell, A., Whinston, M. D. & Green, J. R. *Microeconomic Theory* 981 pp. (Oxford Univ. Press, New York, NY, 1995).

31. Hayek, F. A. von. The Pretence of Knowledge (Nobel Lecture). *American Economic Review* **79**, 3 (1989).

32. Karl, R. *Microprocessor Trend Data* Github:karlrupp/microprocessor-trend-data.

33. De Wolf, R. The Potential Impact of Quantum Computers on Society. *Ethics and Information Technology* **19**, 271 (2017).

34. Birhane, A., Kasirzadeh, A., Leslie, D. & Wachter, S. Science in the Age of Large Language Models. *Nature Reviews Physics* **5**, 277 (2023).

35. Hey, T., Tansley, S., Tolle, K. & Gray, J. *The Fourth Paradigm: Data-intensive Scientific Discovery* (Microsoft Research, 2009).

36. Einav, L. & Levin, J. Economics in the Age of Big Data. *Science* **346**, 1243089 (2014).

37. López, E., Guerrero, O. A. & Axtell, R. L. A Network Theory of Inter-Firm Labor Flows. *EPJ Data Science* **9**, 33 (2020).

38. Einav, L. & Levin, J. *The Data Revolution and Economic Analysis* w19035 (National Bureau of Economic Research, Cambridge, MA, 2013), w19035.

39. Athey, S. Beyond Prediction: Using Big Data for Policy Problems. *Science* **355**, 483 (2017).

40. Fama, E. F. The Behavior of Stock-Market Prices. *The Journal of Business* **38**, 34 (1965).

41. Samuelson, P. A. in Malliaris, A. G. & Ziemba, W. T. *World Scientific Handbook in Financial Economics Series* 25 (WORLD SCIENTIFIC, 2015).

42. Sewell, M. *History of the Efficient Market Hypothesis* Research Note 11(04) (London School of Economics, London, UK, 2011).

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

44. Ţiţan, A. G. The Efficient Market Hypothesis: Review of Specialized Literature and Empirical Research. *Procedia Economics and Finance* **32**, 442 (2015).
45. Malkiel, B. G. The Efficient Market Hypothesis and Its Critics. *Journal of Economic Perspectives* **17**, 59 (2003).
46. Pesaran, M. H. & Timmermann, A. Predictability of Stock Returns: Robustness and Economic Significance. *The Journal of Finance* **50**, 1201 (1995).
47. Mandelbrot, B. The Variation of Certain Speculative Prices. *The Journal of Business* **36**, 394 (1963).
48. Stanley, H. Econophysics: Can Physicists Contribute to the Science of Economics? *Computing in Science & Engineering* **1**, 74 (1999).
49. Drăgulescu, A. A. & Yakovenko, V. M. Probability Distribution of Returns in the Heston Model with Stochastic Volatility*. *Quantitative Finance* **2**, 443 (2002).
50. Tversky, A. & Kahneman, D. in *Choices, Values, and Frames* (eds Kahneman, D. & Tversky, A.) 1st ed., 44 (Cambridge University Press, Cambridge, United Kingdom, 2000).
51. Briggs, R. A. in *The Stanford Encyclopedia of Philosophy* (eds Zalta, E. N. & Nodelman, U.) Fall 2023 (Metaphysics Research Lab, Stanford University, 2023).
52. Kahneman, D. & Tversky, A. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* **47**, 263 (1979).
53. Ploeg, A. P. C. van der. *Stochastic Volatility and the Pricing of Financial Derivatives Tinbergen Institute Research Series* **366**. 340 pp. (Thela Thesis, Amsterdam, 2006).
54. Wheeler, G. in *The Stanford Encyclopedia of Philosophy* (ed Zalta, E. N.) Fall 2020 (Metaphysics Research Lab, Stanford University, 2020).
55. Andersen, J. V. & Sornette, D. A Mechanism for Pockets of Predictability in Complex Adaptive Systems. *Europhysics Letters (EPL)* **70**, 697 (2005).
56. Conlisk, J. Why Bounded Rationality? *Journal of Economic Literature* **34**, 669 (1996).
57. Farhadi, A. & Vvedensky, D. Risk, Randomness, Crashes and Quants. *Contemporary Physics* **44**, 237 (2003).

58. Jorion, P. Risk Management Lessons from the Credit Crisis. *European Financial Management* **15**, 923 (2009).

59. Sornette, D. Physics and Financial Economics (1776–2014): Puzzles, Ising and Agent-Based Models. *Reports on Progress in Physics* **77**, 062001 (2014).

60. Dessalles, J. L., Ferber, J. & Phan, D. in *Intelligent Complex Adaptive Systems* (eds Yang, A. & Shan, Y.) 255 (IGI Global, 2008).

61. Reynolds, C. W. *Flocks, Herds and Schools: A Distributed Behavioral Model* in *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques* SIGGRAPH '87: Computer Graphics and Interactive Techniques (ACM, 1987), 25.

62. Reynolds, C. W. *Flocks, Herds and Schools: A Distributed Behavioral Model* in *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques* SIGGRAPH '87: Computer Graphics and Interactive Techniques (ACM, 1987), 25.

63. Orcutt, G. H. A New Type of Socio-Economic System. *The Review of Economics and Statistics* **39**, 116 (1957).

64. Simon, H. A. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics* **69**, 99 (1955).

65. Schelling, T. C. Dynamic Models of Segregation†. *The Journal of Mathematical Sociology* **1**, 143 (1971).

66. Ulam, S. *et al. Random Processes and Transformations* in *Proceedings of the International Congress on Mathematics* **2** (Citeseer, 1952), 264.

67. Von Neumann, J., Burks, A. W., *et al.* Theory of Self-Reproducing Automata. *IEEE Transactions on Neural Networks* **5**, 3 (1966).

68. Kirman, A. Ants, Rationality, and Recruitment. *The Quarterly Journal of Economics* **108**, 137 (1993).

69. Arthur, W. B., Holland, J. H., LeBaron, B. D., Palmer, R. G. & Tayler, P. Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. *SSRN Electronic Journal* (1997).

70. Lux, T. & Marchesi, M. Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market. *Nature* **397**, 498 (1999).

71. Arthur, W. B. Inductive Reasoning and Bounded Rationality. *The American Economic Review* **84**, 406 (1994).

72. Bardoscia, M., Barucca, P., Battiston, S., Caccioli, F., Cimini, G., Garlaschelli, D., Saracco, F., Squartini, T. & Caldarelli, G. The Physics of Financial Networks. *Nature Reviews Physics* **3**, 490 (2021).

73. Lux, T. & Marchesi, M. Volatility Clustering in Financial Markets: A Microsimulation of Interacting Agents. *International Journal of Theoretical and Applied Finance* **03**, 675 (2000).

74. Meine, D. C. & Vvedensky, D. D. Emergence of Speculation in a Hierarchical Agent-Based Model. *Physica A: Statistical Mechanics and its Applications* **620**, 128619 (2023).

75. Chen, S., Tien, J. & Spotton Visano, B. in *New Frontiers in Applied Artificial Intelligence* (eds Nguyen, N. T., Borzemski, L., Grzech, A. & Ali, M.) 531 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008).

76. Castellano, C., Fortunato, S. & Loreto, V. Statistical Physics of Social Dynamics. *Reviews of Modern Physics* **81**, 591 (2009).

77. Sornette, D. & Johansen, A. A Hierarchical Model of Financial Crashes. *Physica A: Statistical Mechanics and its Applications* **261**, 581 (1998).

78. Christensen, K. & Moloney, N. R. *Complexity and Criticality* (Imperial College Press, 2005).

79. Westphal, R. & Sornette, D. Market Impact and Performance of Arbitrageurs of Financial Bubbles in an Agent-Based Model. *Journal of Economic Behavior and Organization* **171**, 1 (2020).

80. Dunne, J. A., Williams, R. J. & Martinez, N. D. Network Structure and Biodiversity Loss in Food Webs: Robustness Increases with Connectance. *Ecology Letters* **5**, 558 (2002).

81. Pastor-Satorras, R. & Vespignani, A. Epidemic Spreading in Scale-Free Networks. *Physical Review Letters* **86**, 3200 (2001).

82. Haldane, A. G. & May, R. M. Systemic Risk in Banking Ecosystems. *Nature* **469**, 351 (2011).

83. The Network Origins of Aggregate Fluctuations. *Econometrica* **80**, 1977 (2012).

84. Newman, M. E. J. The Structure and Function of Complex Networks. *SIAM Review* **45**, 167 (2003).

85. Shields, R. Cultural Topology: The Seven Bridges of Königsburg, 1736. *Theory, Culture & Society* **29**, 43 (2012).

86. Moreno, J. L. *Who Shall Survive?: A New Approach to the Problem of Human Interrelations.* (Nervous and Mental Disease Publishing Co, Washington, 1934).

87. Erdős, P. & Rényi, A. On Random Graphs. I. *Publicationes Mathematicae Debrecen* **6**, 290 (2022).

88. Erdos, P., Rényi, A., *et al.* On the Evolution of Random Graphs. *Publ. math. inst. hung. acad. sci* **5**, 17 (1960).

89. Newman, M. *Networks* (Oxford University Press, 2018).

90. Milgram, S. The Small World Problem. *Psychology today* **2**, 60 (1967).

91. Leskovec, J. & Horvitz, E. *Planetary-Scale Views on a Large Instant-Messaging Network* in *Proceedings of the 17th International Conference on World Wide Web* WWW '08: The 17th International World Wide Web Conference (ACM, Beijing China, 2008), 915.

92. Watts, D. J. & Strogatz, S. H. Collective Dynamics of 'Small-World' Networks. *Nature* **393**, 440 (1998).

93. Barabási, A.-L. & Albert, R. Emergence of Scaling in Random Networks. *Science* **286**, 509 (1999).

94. Allen, F. & Gale, D. Financial Contagion. *The Journal of Political Economy* **108**, 1 (2000).

95. Glasserman, P. & Young, H. P. Contagion in Financial Networks. *Journal of Economic Literature* **54**, 779 (2016).

96. Glasserman, P. & Young, H. P. Contagion in Financial Networks. *Journal of Economic Literature* **54**, 779 (2016).

97. Boss, M., Elsinger, H., Summer, M. & Thurner, S. *The Network Topology of the Interbank Market* `http://arxiv.org/abs/cond-mat/0309582` (2024). Pre-published.

98. Acemoglu, D., Ozdaglar, A. & Tahbaz-Salehi, A. Systemic Risk and Stability in Financial Networks. *American Economic Review* **105**, 564 (2015).

99. Freixas, X., Parigi, B. M. & Rochet, J.-C. Systemic Risk, Interbank Relations, and Liquidity Provision by the Central Bank. *Journal of Money, Credit and Banking* **32**, 611 (2000).

100. Battiston, S., Gatti, D. D., Gallegati, M., Greenwald, B. & Stiglitz, J. *Liaisons Dangereuses: Increasing Connectivity, Risk Sharing, and Systemic Risk* w15611 (National Bureau of Economic Research, Cambridge, MA, 2009), w15611.

101. Onnela, J.-P., Chakraborti, A., Kaski, K., Kertesz, J. & Kanto, A. Dynamics of Market Correlations: Taxonomy and Portfolio Analysis. *Physical Review E* **68**, 056110 (2003).

102. Tumminello, M., Lillo, F., Piilo, J. & Mantegna, R. N. Identification of Clusters of Investors from Their Real Trading Activity in a Financial Market. *New Journal of Physics* **14**, 013041 (2012).

103. Carvalho, V. M. From Micro to Macro via Production Networks. *Journal of Economic Perspectives* **28**, 23 (2014).

104. HERSKOVIC, B. Networks in Production: Asset Pricing Implications. *The Journal of Finance* **73**, 1785 (2018).

105. Marti, G., Nielsen, F., Bińkowski, M. & Donnat, P. in *Progress in Information Geometry* (ed Nielsen, F.) 245 (Springer International Publishing, Cham, 2021).

106. Baddeley, M. *Behavioural Economics: A Very Short Introduction* (Oxford University Press, 2017).

107. McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G. & Cohen, J. D. Time Discounting for Primary Rewards. *The Journal of Neuroscience* **27**, 5796 (2007).

108. Rodrigue, J.-P. *The Geography of Transport Systems* 5th ed. (Routledge, Fifth edition. | Abingdon, Oxon ; New York, NY : Routledge, 2020., 2020).

109. Asch, S. E. in *Organizational Influence Processes* 295 (Routledge, 2016).

110. Kilner, J. & Lemon, R. What We Know Currently about Mirror Neurons. *Current Biology* **23**, R1057 (2013).

111. *Neuroeconomics* (Elsevier, 2014).

112. Gabay, A. S., Radua, J., Kempton, M. J. & Mehta, M. A. The Ultimatum Game and the Brain: A Meta-Analysis of Neuroimaging Studies. *Neuroscience & Biobehavioral Reviews* **47**, 549 (2014).

113. Knutson, B. & Greer, S. M. Anticipatory Affect: Neural Correlates and Consequences for Choice. *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**, 3771 (2008).

114. Glimcher, P. W., Camerer, C. F., Fehr, E. & Poldrack, R. A. in *Neuroeconomics* 1 (Elsevier, 2009).

115. *Expected Utility Hypotheses and the Allais Paradox* (eds Allais, M. & Hagen, O.) (Springer Netherlands, Dordrecht, 1979).

116. Ellsberg, D. Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics* **75**, 643 (1961).

117. Newsome, W. T., Britten, K. H. & Movshon, J. A. Neuronal Correlates of a Perceptual Decision. *Nature* **341**, 52 (1989).

118. Usher, M. & McClelland, J. L. The Time Course of Perceptual Choice: The Leaky, Competing Accumulator Model. *Psychological Review* **108**, 550 (2001).

119. Busemeyer, J. R. & Townsend, J. T. Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment. *Psychological Review* **100**, 432 (1993).

120. McClelland, J. L. On the Time Relations of Mental Processes: An Examination of Systems of Processes in Cascade. *Psychological Review* **86**, 287 (1979).

121. Abbott, L. Lapicque's Introduction of the Integrate-and-Fire Model Neuron (1907). *Brain Research Bulletin* **50**, 303 (1999).

122. Stevens, C. F. The Neuron. *Scientific American* **241**, 54 (1979).

123. Gerstner, W., Kistler, W. M., Naud, R. & Paninski, L. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition* 1st ed. (Cambridge University Press, Cambridge, United Kingdom, 2014).

124. *Sketch of a brain neuron* `https://commons.wikimedia.org/wiki/File:Sketch_of_a_brain_neuron.png`.

125. Herculano-Houzel, S. The Remarkable, yet Not Extraordinary, Human Brain as a Scaled-up Primate Brain and Its Associated Cost. *Proceedings of the National Academy of Sciences* **109**, 10661 (supplement_1 2012).

126. Lymperopoulos, I. N. & Ioannou, G. D. Online Social Contagion Modeling through the Dynamics of Integrate-and-Fire Neurons. *Information Sciences* **320**, 26 (2015).

127. Brandes, U., Gaertler, M. & Wagner, D. in *Algorithms - ESA 2003* (eds Di Battista, G. & Zwick, U.) red. by Goos, G., Hartmanis, J. & Van Leeuwen, J., 568 (Springer Berlin Heidelberg, Berlin, 2003).

128. Stephenson, K. & Zelen, M. Rethinking Centrality: Methods and Examples. *Social Networks* **11**, 1 (1989).

129. Christensen, K. & Moloney, N. R. *Complexity and Criticality Imperial College Press Advanced Physics Texts* **1**. 392 pp. (Imperial College Press, London, 2005).

130. Hurst, H. E. Long-Term Storage Capacity of Reservoirs. *Transactions of the American Society of Civil Engineers* **116**, 770 (1951).

131. Mandelbrot, B. B. & Wallis, J. R. Robustness of the Rescaled Range R/S in the Measurement of Noncyclic Long Run Statistical Dependence. *Water Resources Research* **5**, 967 (1969).

132. Garofalo, M., Botta, A. & Ventre, G. Astrophysics and Big Data: Challenges, Methods, and Tools. *Proceedings of the International Astronomical Union* **12**, 345 (2016).

133. Collins, J. H., Howe, K. & Nachman, B. Extending the Search for New Resonances with Machine Learning. *Physical Review D* **99**, 014038 (2019).

134. Takahashi, S., Chen, Y. & Tanaka-Ishii, K. Modeling Financial Time-Series with Generative Adversarial Networks. *Physica A: Statistical Mechanics and its Applications* **527**, 121261 (2019).

135. *Advances in Behavioral Economics* (eds Camerer, C. F., Loewenstein, G. & Rabin, M.) (Princeton University Press, Princeton, New Jersey, 2004).

136. Carandini, M. & Heeger, D. J. Normalization as a Canonical Neural Computation. *Nature Reviews Neuroscience* **13**, 51 (2012).

137. Ratcliff, R. & Smith, P. L. A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychological Review* **111**, 333 (2004).

138. Bogacz, R. Optimal Decision-Making Theories: Linking Neurobiology with Behaviour. *Trends in Cognitive Sciences* **11**, 118 (2007).

139. Arthur, W. B. Complexity and the Economy. *Science* **284**, 107 (1999).

140. Hogarth, R. M. & Einhorn, H. J. Order Effects in Belief Updating: The Belief-Adjustment Model. *Cognitive Psychology* **24**, 1 (1992).

141. Tversky, A. & Kahneman, D. Judgment under Uncertainty: Heuristics and Biases: Biases in Judgments Reveal Some Heuristics of Thinking under Uncertainty. *Science* **185**, 1124 (1974).

142. Shiller, R. J. *Irrational Exuberance: (Second Edition)* (Princeton University Press, Princeton, New Jersey, 2005).

143. Shannon, C. E. A Mathematical Theory of Communication. *Bell System Technical Journal* **27**, 379 (1948).

144. Weller, A., Tang, K., Sontag, D. & Jebara, T. *Understanding the Bethe Approximation: When and How Can It Go Wrong?* in *30th Conference on Uncertainty in Artificial Intelligence, UAI 2014* (AUAI Press, Quebec City, Canada, 2014), 868.

145. Statistical Theory of Superlattices. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences* **150**, 552 (1935).

146. Mandelbrot, B. B. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk. Selecta Volume E* 551 pp. (Springer, New York, NY, 1997).

147. Mantegna, R. N. & Stanley, H. E. *Introduction to Econophysics: Correlations and Complexity in Finance* 1st ed. (Cambridge University Press, Cambridge, United Kingdom, 1999).

148. Bouchaud, J.-P. & Potters, M. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management* 2nd ed. (Cambridge University Press, Cambridge, United Kingdom, 2003).

149. Pereira, T., Eroglu, D., Bagci, G. B., Tirnakli, U. & Jensen, H. J. Connectivity-Driven Coherence in Complex Networks. *Physical Review Letters* **110**, 234103 (2013).

150. Ante, L. How Elon Musk's Twitter Activity Moves Cryptocurrency Markets. *Technological Forecasting and Social Change* **186**, 122112 (2023).

151. Hussain, S. A., Meine, D. C. A. & Vvedensky, D. D. Integrate-and-Fire Model of Disease Transmission. *Physical Review E* **110**, 014305 (2024).

152. McAdams, D. *et al.* Nash SIR: An Economic-Epidemiological Model of Strategic Behavior during a Viral Epidemic. *Covid Economics (forthcoming)* (2020).

153. Kermack, W. O. & McKendrick, A. G. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* **115**, 700 (1927).

154. Weiss, H. H. The SIR Model and the Foundations of Public Health. *Materials matematics* **2013**, 1 (2013).

155. Meyers, L. A., Newman, M. E. J., Martin, M. & Schrag, S. Applying Network Theory to Epidemics: Control Measures for Mycoplasma Pneumoniae Outbreaks. *Emerging Infectious Diseases* **9**, 204 (2003).

156. Newman, M. E. Spread of Epidemic Disease on Networks. *Physical Review E* **66**, 016128 (2002).

157. Hamilton, K., Date, P., Kay, B. & Schuman D, C. *Modeling Epidemic Spread with Spike-Based Models* in *International Conference on Neuromorphic Systems 2020* (2020), 1.

158. Sette, A. & Crotty, S. Adaptive Immunity to SARS-CoV-2 and COVID-19. *Cell* **184**, 861 (2021).

159. McAdams, D. The Blossoming of Economic Epidemiology. *Annual Review of Economics* **13**, 539 (2021).

160. Kamo, M. & Sasaki, A. The Effect of Cross-Immunity and Seasonal Forcing in a Multi-Strain Epidemic Model. *Physica D: Nonlinear Phenomena* **165**, 228 (2002).

161. Li, C., Wang, J., Xu, J. & Rong, Y. The Global Dynamics of a SIR Model Considering Competitions among Multiple Strains in Patchy Environments. *Mathematical Biosciences and Engineering* **19**, 4690 (2022).

162. Albert-László, B. & Márton, P. *Network Science* (Cambridge University Press, Cambridge, United Kingdom, 2017).

163. *UK Coronavirus Dashboard - Data* `https://coronavirus.data.gov.uk/details/cases?areaType=nation&areaName=England`.

164. *Prime Minister Sets out Plan for Living with COVID* `https://www.gov.uk/government/news/prime-minister-sets-out-plan-for-living-with-covid`.

165. Ibitoye, O. S., Olasunkanmi, Y. A., Olowolafe, T. A., Olabode, A. T., Salawu, M. M. & Afolabi, R. F. Predictors and Time to Recovery from COVID-19 among Patients Attended at the Treatment Centers in Ekiti State, South West, Nigeria. *The Pan African Medical Journal* **42** (2022).

166. Kaso, A. W., Hareru, H. E., Kaso, T. & Agero, G. Time to Recovery from Covid-19 and Its Associated Factors among Patients Hospitalized to the Treatment Center in South Central Ethiopia. *Environmental Challenges* **6**, 100428 (2022).

167. Tsegaye, S., Bekele, F., Lulu, Y., Debele, G. R., Bekana, T., Tolesa, L. D. & Bidira, K. Time to Recovery and Its Predictors among COVID-19 Positive Patients Admitted to Treatment Centers of Southwestern Ethiopian Hospitals. A Multicenter Retrospective Cohort Study. *Annals of Medicine and Surgery* **84**, 104917 (2022).

168. Churiso, G., Diriba, K., Girma, H. & Tafere, S. Clinical Features and Time to Recovery of Admitted COVID-19 Cases at Dilla University Referral Hospital Treatment Center, South Ethiopia. *Infection and drug Resistance*, 795 (2022).

169. Liu, B., Jayasundara, D., Pye, V., Dobbins, T., Dore, G. J., Matthews, G., Kaldor, J. & Spokes, P. Whole of Population-Based Cohort Study of Recovery Time from COVID-19 in New South Wales Australia. *The Lancet Regional Health–Western Pacific* **12** (2021).

170. Stein, C., Nassereldine, H., Sorensen, R. J., Amlag, J. O., Bisignano, C., Byrne, S., Castro, E., Coberly, K., Collins, J. K., Dalos, J., *et al.* Past SARS-CoV-2 Infection Protection against Re-Infection: A Systematic Review and Meta-Analysis. *The Lancet* **401**, 833 (2023).

171. *Test and Trace: What to Do If You Are Contacted* https://www.gov.uk/guidance/nhs-test-and-trace-how-it-works. 2020.

172. Home Office. *Rule of Six Comes into Effect to Tackle Coronavirus* https://www.gov.uk/government/news/rule-of-six-comes-into-effect-to-tackle-coronavirus. GOV.UK, 2020.

173. Sobel Leonard, A., McClain, M. T., Smith, G. J., Wentworth, D. E., Halpin, R. A., Lin, X., Ransier, A., Stockwell, T. B., Das, S. R. & Gilbert, A. S. et al. The Effective Rate of Influenza Reassortment Is Limited during Human Infection. *PLOS Pathogens* **13** (2017).

174. Andreasen, V. & Sasaki, A. Shaping the Phylogenetic Tree of Influenza by Cross-Immunity. *Theoretical Population Biology* **70**, 164 (2006).

175. *How Flu Spreads* https://www.cdc.gov/flu/about/disease/spread.htm. Centers for Disease Control and Prevention, 2022.

176. *Types of Influenza Viruses* https://www.cdc.gov/flu/about/viruses/types.htm. Centers for Disease Control and Prevention, 2023.

177. Nobusawa, E. & Sato, K. Comparison of the Mutation Rates of Human Influenza A and B Viruses. *Journal of Virology* **80**, 3675 (2006).

178. *FluNet* https://www.who.int/tools/flunet, organization = World Health Organization.

179. Public Health Agency of Canada. *Flu (Influenza): For Health Professionals* https://www.canada.ca/en/public-health/services/diseases/flu-influenza/health-professionals.html. / Gouvernement du Canada, 2023.

180. Government of Canada, Statistics Canada. https://www150.statcan.gc.ca/n1/pub/11-402-x/2008/3867/ceb3867_000-eng.htm. 2008.

181. *Population Density (People per Sq. Km of Land Area) - Canada* https://data.worldbank.org/indicator/EN.POP.DNST?locations=CA.

182. *Population Density - Office for National Statistics* https://www.ons.gov.uk/datasets/TS006/editions/2021/versions/4.

183. Albantakis, L., Barbosa, L., Findlay, G., Grasso, M., Haun, A. M., Marshall, W., Mayner, W. G. P., Zaeemzadeh, A., Boly, M., Juel, B. E., Sasai, S., Fujii, K., David, I., Hendren, J., Lang, J. P. & Tononi, G. Integrated Information Theory (IIT) 4.0: Formulating the Properties of Phenomenal Existence in Physical Terms. *PLOS Computational Biology* **19** (ed Graham, L. J.) e1011465 (2023).

184. Angione, C., Silverman, E. & Yaneske, E. Using Machine Learning as a Surrogate Model for Agent-Based Simulations. *PLOS ONE* **17** (ed Bouffanais, R.) e0263150 (2022).