# Complexity and Criticality in Financial Markets: Systemic Risk across Frequencies and Cross Sections

*Jeremy D. Turiel*

A dissertation submitted in partial fulfillment

of the requirements for the degree of

**Doctor of Philosophy**

of

**University College London**.

Department of Computer Science

University College London

SID: 15058723

November 27, 2022

I, Jeremy D. Turiel, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

# Abstract

Extreme market events and systemic collapses cause most of the popular attention to finance and financial markets. Extreme phenomena and the dynamics of connected/interacting systems have been the subject of financial modeling since early derivatives modeling, exposure risk modeling and portfolio construction. In the present work we discuss how traditional methods have for the most part failed to properly model the interconnected global financial and economic system. This led to systemic risk events and simplistic regulation which does not properly account for its implications. Analogously, we discuss how from as early as Mandelbrot's works on financial prices and fat tails, academics, practitioners and regulators alike were warned of fat tails in financial modeling and in particular market making and derivatives pricing. The improper modeling or dismissal of these lies at the centre of financial downturns ranging from LTCM's collapse to the quant downturn of August 2007.

The solution I promote in this thesis is that of complexity and criticality. In line with this we propose two lines of work. The former analyses markets as complex networks and their structure through to practical takeaways including a proof of concept for portfolio construction. The latter instead focuses on extreme events in high frequency markets with results for both tail modeling and systemic events and practical insights from those. Recent events have shown how retail investors and their savings are now heavily involved in financial markets. We hope that our contribution of methods of practical use for proper risk modeling will encourage their adoption by practitioners and regulators with the outcome of a more stable and efficient financial system.

# Acknowledgements

In choosing my PhD and supervisor I was looking not only for an academic guide, but also for a life mentor and friend to guide me and help me improve in all aspects of my academic, professional and even personal life. I deeply thank my PhD supervisor Prof. Tomaso Aste for being all this and for having the patience and understanding to bear with me and help me succeed throughout my PhD. Covid-19 interrupted in person collaboration and conferences, but we managed to adapt and thrive, publishing in major journals and winning a best paper award when back at the first conference.

I would like to acknowledge all members of the Financial Computing & Analytics group at UCL for the support, guidance and interesting discussions over lunch, coffee or just in the corridor. In particular I had the pleasure to co-author with my co-supervisor P. Barucca and with D. Fernandez-Reyes who was also my transfer viva examiner and always a go to person for experienced guidance and encouragement. I have often bothered G. Livan, F. Caccioli and G. Germano for questions and advice and their warm help was a motivational booster. C. Phelan has been my go to Bloomberg person for way more support than she expected, but has always kindly helped and been instrumental to much of my work. Outside of UCL I had the opportunity of chats and encouragement from G. Caldarelli at conferences and over email. Covid shortened geographical distances and I am humbled by the opportunity to co-author works with P. Kolm and N. Westray. P. Kolm has been an amazing academic guidance, source of experience and advice and in a way a third supervisor. N. Westray has become a good friend and a life and career mentor, on top of all the technical knowledge and thorough discussions I have enjoyed with him.

My family and close friends have been my rock throughout difficult moments and have shared all my joys and successes. My parents in particular have showed much needed patience, understanding and an open mindedness that I discovered to be exceptional, rare and precious. As my brother went on to further studies and started sharing my passion for quantitative finance, the excitement and complicity of our discussions over dinner has really made this a family effort - I look forward to rejoicing of his achievements and seeing his thesis soon enough!

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Context for complexity and criticality in the thesis

Technological and methodological advances have allowed researchers to gain unprecedented insights from the use of big data and advanced computational methods. This has allowed to validate pre-existing complex theories and create new ones.

Many natural and anthropic systems involve the interaction of a multitude of agents which results in global characteristics of the system. It was found that, as one is able to track agents and their interactions in a system, many phenomena and characteristics can be explained as emerging from simple rules. With the availability of computing and growth of AI-like systems, the field of agent-based modeling is rapidly expanding. From a physicist's perspective though, such systems are at the basis of statistical mechanics, condensed matter theory and more. From these fields, the discipline and term "complexity" have emerged, defined in [1] as the repeated application of simple rules in systems with many degrees of freedom that gives rise to emergent behaviour not encoded in the rules themselves.

Criticality follows nicely from this concept and is better defined amongst statistical physicists. This refers to the behaviour of complex and extended systems at a phase transition where observables become scale free, i.e. with no characteristic size. Analogously to complexity, at a phase transition, the interacting components of the systems give rise to emergent macroscopic phenomena which cannot be intuitively explained from the simple laws governing the microscopic components of the system. Criticality highlights the importance of the interaction between the microscopic components in the system to explain emergent phenomena and can be viewed as the study of complex systems at a critical point of transition (hence the name). Phase transitions arise in classic statistical mechanics systems in physics and are well-studied. Examples of these are the Ising model and percolation [2, 3, 4].

Many systems in today's world are complicated and present systemic breakdowns, but not all are complex and critical. The crucial distinction arises in that complex critical systems are often in between non-equilibrium states where the system's reactions to a "small" energy release depend strongly on the system's configuration. Small changes in the system's configuration can determine wildly different, highly non-linear reactions to small energy releases in the system.

Complexity and criticality, and related methodologies, have become increasingly relevant in the study of quantitative finance, but due to the contextual understanding and the advanced mathematical knowledge they require, they are often overlooked.

This is indeed history repeating itself, as already Mandelbrot had studied financial time series of price returns and suggested that "Large price changes are more frequent than predicted". In his early work on the topic [5], which was overlooked for a long time, he suggests that the random walks used to model financial time series were wrong on a crucial point. These walks assumed prices to evolve as a Brownian motion with Gaussian-like steps, while Mandelbrot highlighted the presence of fat-tails in the distribution of returns and suggested to model this as a Levy distribution [5, 6, 7, 8]. In the 1990s the seminal works of Mantegna and Stanley started the community and literature in Econophysics, this was largely based on the fat-tailes described by Mandelbrot. Still, Mandelbrot's work was largely revisited by mainstream academia and practitioners only after the quant meltdown of August 2007. It is now widely accepted, but neither the method nor the observation were accepted at the time of publication.

Normality assumptions on the price process allow access to a broader suite of widely known modeling techniques, which is perhaps acceptable when looking for predictable patterns through quantitative models (Alpha research), but can prove very dangerous when used to determine and model risk.

The first notorious example of this came with the rise of the first quant hedge funds and quant groups in major investment banks in the late 1980s, as Black Monday (October $19^{th}$ 1987) [9, 10, 11] brought this reality to the world of quantitative investing.

The collapse and following bailout of Long Term Capital Management (LTCM) in 1998 is another example of this, where a fat-tailed diversion in bond prices with different issue dates due to a run to liquidity/safety caused massive losses to the fund, sending it into default. The fund was using large amounts of leverage and hedging its bets. The magnitude of the event/divergence was theoretically impossible in a Normal distribution - this shows how neglecting fat tails when dealing with risk caused one of the most systemic and potentially dangerous defaults in modern financial history.

This brief anecdotal introduction is meant to motivate the reader and the wider community not only to acknowledge fat tails, but also to deal with their implications and related modeling techniques.

Mandelbrot is mostly know in complex systems for his seminal work on fractality/fractals, a major one applied to financial markets [12].

Fractals show how certain properties are present at different scales in the same system. In a broad analogy, in this work we also consider complexity and criticality at different scales in financial systems.

Indeed most chapters will be structured in two main sections, one focuses on "traditional" financial network analysis of daily returns and dynamics while the other looks at the less notorious world of high frequency finance and its critical phenomena, which have proven far more crucial for modern quantitative finance.

The brief introduction to complex systems and their system-wide phase transitions due to synchronisation (criticality) together with the small anecdotal example regarding LTCM hint to the fact that risk, but in particular systemic risk are central to both critical systems and financial markets.

Uniting both these themes we will deal with risk and in particular systemic risk at different scales in financial markets from the perspective of complexity and criticality.

One last ancillary, but nonetheless relevant, theme to introduce is that of denoising and extracting signal in real systems which in physics is done through null models. Null models are techniques of constrained randomisation of a system and its dynamics which allow to identify non-random patterns in the real data, given the randomisation assumptions. These will be applied in the present work to show the non-randomness of our results and determine the significance of observed patterns.

## 1.1 Market structure and its estimation

In his seminal work on portfolio construction [13] Henry Markowitz provides an analytical solution for optimal portfolio allocation, the general framework which stems from this work is commonly referred to as Modern Portfolio Theory (MPT). MPT relies on two main inputs: an expectation of return for each asset and an estimate of pairwise covariance between all assets in the portfolio. The use of the covariance matrix is crucial as it shows how to optimise portfolio performance and reduce risk one cannot just account for asset-wise volatility or "beta" (correlation/covariance) to the overall market. Rather, individual relations between assets are central to reducing risk. The ensemble of pairwise relations in the form of the covariance matrix introduces us to the concept of market structure.

As per the discussion above, the study of complex systems is interested with systems of interacting components. This has sparked the development of the discipline and study of complex networks. In the words of Mark Newman, one of the main academic figures in the field, "a network is, in its simplest form, a collection of points joined together in pairs by lines" [14]. The discipline has developed a vast literature ranging from theoretical models [15] to empirical analyses of real world systems [16, 17] and everything in between.

We have now established that structure and relations between assets are crucial in portfolio applications and that complex systems deal with connected systems.

We have also established that the central item of this structural estimation is the covariance matrix of price returns.

With the critical eye of null models, complex systems and more specifically Random Matrix Theory (RMT), several authors have investigated these covariance matrices and their information content. Early works look at the eigenvalue distribution of empirical return covariances and notice how the high level of noise makes only the first few large eigenvalues and corresponding eigenvectors deviate from the null model of a random matrix [18]. In their case 94% of the spectrum could be fitted by the null model. Indeed, even most recent works [19] from some of those early authors model the covariance matrix based on its largest eigenvalue alone, as this is deemed a reasonable approximation at least for theoretical work. Other early authors in econophysics noticed this phenomenon [20] and both [20, 18] used their observations to suggest that the covariance might correspond to a random band matrix which suggests that a distance metric between companies can be defined. The authors in [20] acknowledge that such observations and suggestions could already be found in earlier work [21].

Two followup works from some of the authors cited above suggest a method based on RMT to remove the eigenvalues and eigenvectors which correspond to the noise band [22, 23]. The method was found to be very effective in reducing estimation noise in empirical covariance matrices.

Early works in covariance estimation [18] had already suggested that the $N \cdot (N-1)$ parameters of the covariance matrix could present a problem as we scale to larger covariance matrices, since the number of parameters scales quadratically with the number of assets while the number of observations scales linearly with the lookback time window $T$. Indeed we know that if we don't have $T \gg N$ the estimations will be noisy and the covariance will appear random to a large extent. On the other hand, markets are non-stationary in time so a long lookback $T$ might introduce noise as the relations slowly change with time.

A solution to this was initially proposed by [24] and then elaborated upon by [25] where the authors suggest an exponential weighting of the observations in time to prioritise more recent and relevant observations while retaining a long lookback window which helps denoising the covariance estimates.

A comprehensive review of covariance estimation and cleaning methods was later presented in [26].

RMT is indeed a mathematically sound approach to the problem with strong theoretical foundations in the field of mathematics and physics. On the other hand we have discussed above how complexity introduced works on complex networks and one can see how a covariance matrix is a mathematical object representing relations between its elements and hence a connected system. One issue remained to transition from covariace matrices to networks: sparsity. Indeed shrinkage and thresholding methods were proposed in [26] and similar works, but many methodologies and analyses in complex networks require relatively sparse connections, but with a large main connected component. The former was not achieved by shrinkage while the latter could not be achieved by thresolding.

For completeness, before presenting works on network filtering which is the focus of covariance estimation in this work, we discuss the earlier logical step of statistical validation. Statistical validation - which constitutes a generalisation of simpler thresholding methods - has been used to establish the significance of edges in correlation matrices, with applications to economics and finance as well as to other fields [27, 28, 29, 30]. Statistical validation can be implemented by comparing the empirical correlations with null-hypothesis correlations generated from time-series randomized over the time dimension to remove dependency.

In the same spirit, one of the authors of [20] proposed a method which took the “filtering” idea of [22, 23] a step further and imposed structure on the system [31]. The work by Mantegna & Stanley [32] is at the core of network filtering and suggests a recursive algorithm with selects the strongest connections from the given covariance matrix and adds them to the edge list of the network as long as the addition yields a structure that is still a tree (i.e. no loops). Beside the value added by the method, a sparse network which can be well represented visually offered a technique with great interpretability.

All methods from RMT mentioned above aim to denoise the information in the covariance estimation. These proposed methods were then tested on portfolio performance where the denoised covariance is used to build the Markowitz portfolio of some sort. In order to do this, the denoised covariance matrix needs to be invertible, which is the case for well-defined covariance matrices $T \gg N$ in general. This creates an issue with thresholding methods and the tree-filtering approach of [31] which do not yield invertible matrices.

This issue is solved by the authors in followup works which introduce more sophisticated filtering techniques that yield invertible matrices [33, 34]. Those authors then went on to propose fast construction methods for a subset of such filtered graphs [35] for which the covariances can be estimated locally (subsets of $N'$ neighboring assets), hence reducing the estimation noise as $N' \ll N$ and we only require $T \gg N'$.

## 1.2 Percolation vulnerability of network structures

We have now established from Markowitz's work, and the literature on covariance estimation that followed, that structure and its correct identification is important for portfolio construction. We have then seen that we can use filtering techniques and frame the investigation of market structures in the context of complex networks. Highlighting how the covariance, hence the structure, is what allows to minimise risk in MPT we now take a step back and understand why structure is important in the context of complex networks more broadly and review some of the general literature that will be relevant for financial risk applications.

The resilience and robustness of networks is of great interest to real world applications as it relates to telecommunication networks, electric grids, the internet, social networks and more. Below we present an overview of network robustness and its foundations based upon the structure presented in [36].

An early study of the theoretical properties of robustness in network models was conducted by Albert, Jeong and Barabasi [37]. In this work the authors consider two main categories of network degree distributions and corresponding models. The first category is characterised by a degree distribution that peaks around the average degree with exponentially decaying tails which is the case for early models of network growth such as the random graph model of Erdos and Reyini [38] and the small-world model of Watts and Strongatz [39]. Both of these early models lead to networks with fairly homogenous degree distributions, which is not what is observed in real world networks where we see the emergence of highly heterogeneous degree distributions which decay as a power law and are therefore scale-free [40, 41, 17]. This is the case, for instance, in citation networks [42] and the Internet [43, 44]. In such networks, the probability of a node with a large number of connections (with respect to the median) is much higher than in a Gaussian distribution. To obtain networks with such characteristics the authors in [37] adopt a scale-free model [40] which is based upon growth and preferential attachment [40, 41]. The main result of this seminal work is then that scale-free networks are surpringly robust on average to random error and attacks when compared to more homogenous networks. The caveat lies though in the fact that nodes with large degree (the "hubs") can be exploited in non-random targeted attacks to break the network in smaller components, as these nodes are topologically central.

The random and targeted attacks on networks discussed above merely consist in the removal of nodes in a network. This can be viewed as the failure of a server in the connected World Wide Web or a telecommunication system or the disappearance of an individual in a social network. When looking at resilience we are interested in attacks which cause failures that spread across the network or break down its connected components into significantly smaller ones. An analogy can be found in Physics in *percolation theory* [2, 45] which can simply consider an arbitrary topology where each node is occupied with probability $p$ and edges exist only between active nodes. Depending on the topology, at different probability thresholds of $p$ we observe the emergence of scale-free clusters. It can be intuitively seen how we can draw connections between clusters in percolation theory and components in networks and how $p$ is complementary to the removal probability of a node.

The difference between traditional percolation theory and the corresponding phenomenon in complex networks is that the former considers percolation on lattices in an N-dimensional space, while in complex networks any node can be connected to any other in the network without constraints, hence for a network with $N$ nodes it can be embedded in $(N-1)$ dimensions where any node can have $(N-1)$ neighbours. Taking this a step further in the limit $N \to \infty$ we are considering edge percolation in infinite dimensions. Depending on topological properties we may or may not see a giant connected component (i.e. that always contains a finite fraction of all nodes $N$) - this is analogous to the critical point where we no longer see a fragmented network, but the emergence of a component with size that scales as the number of nodes $N$. A more in-depth theoretical discussion and connection between results in percolation theory and network robustness can be found in [46, 47] with a review present in [36].

## 1.3 Systemic risk in financial networks

We now bring together the previous two sections on market structure and network vulnerability to discuss systemic risk in financial markets. We have described in Section 1.1 how correlation matrices can be "filtered" to create sparse network structures. This will be relevant for the results presented in this thesis, but by means of introduction to the topic of financial networks we highlight that so called "real" networks exist in finance and are the subject of most of the early literature on the topic.

Financial (interbank) networks are complex systems where financial institutions are interconnected in a variety of ways. The "real" network of connections originating from financial contracts between institutions or other companies is one of the most relevant. These contracts normally represent lending (mostly overnight) between financial institutions in order to optimise short-term liquidity exposures and exploit available opportunities. Financial institutions also collaborate on the origination of financial products which are often repackaged and sold between them in a chain. These networks of "direct" interdependencies have been extensively studies in the literature due to data availability and the central role they played in the 2008 crisis/credit squeeze. Allen & Gale [48, 49] discuss in two visionary articles how credit lines between financial institutions where the borrowed money is used to invest in risk assets can have dramatic effects of credit freezes and fire sales across the system. Counterparty exposures in the trading and clearing of assets also form an important financial network for systemic risk, this is the subject of the work by Eisenberg & Noe [50] who propose a computational method for optimal clearing in the system and argue that the solution is unique. A similar form of asset-based exposure in financial institutions comes from cross-holdings of each other's stock and holding the same assets, where the former is commonly used to inflate balance sheets. This network is considered by Elliott et al.[51] where the authors model the cascade of failures in financial institutions based on it and discuss benefits and failures of diversification and integration. Diversification across holdings and cross-holdings indeed reduces volatility and concentration, but also creates overlap between the balance sheets of different institutions making them interdependent if one institution has to liquidate its assets and push their price down, thereby creating systemic risk and domino effects. This leads us to consider the more nuanced correlation between balance sheet holdings of different institutions which forms a fully connected weighted network between financial institutions, where edge weights often correspond to a measure of asset holding or portfolio correlations. Allen et al.[52] investigate asset commonalities and short-term debt of banks and how the combination of concentration structure and term of the debt generate excessive systemic risk during shocks and their propagation in the system. Diebold & Mariano [53] present an empirical investigation of holdings concentration and connectedness amongst financial institutions and how this evolved during different historical shock periods. Cabrales et al.[54] investigate the tradeoff between risk sharing and contagion by finding optimal tradeoff points in relation to the size and characteristics of the shock. Results show how the optimal networks often present different levels of linkage to other firms and instances of segmentation.

The recent review of systemic risk in financial networks by Jackson and Pernaud [55] provides an overview of real financial networks and their relation to systemic risk. They organise some of the instances of systemic contagion described above as follows: Contagion through "direct externalities", which arises from direct counterparty relations between financial institutions or correlated investments and exposures more generally. The former can cause insolvency which cascades across the system of loans across institutions, while the latter is a bit more subtle. Even when a bank does not directly lend to another they might hold similar assets in their portfolios or be exposed to similar factors. When an institution enters in distress it is forced to quickly liquidate its holdings and reduce leverage. This adversely impacts market prices and the correlated portfolios of other institutions which might have to de-risk themselves, thereby inducing a negative feedback loop analogous to that of August 2007. One last aspect of these phenomena which leads onto the next class of systemic risk, i.e. feedback effects and self-fulfilling prophecies, is similarity. When a bank becomes insolvent similar institutions are negatively affected even if they don't have direct or indirect exposures just based on the fact that the "market" doubts their holdings, structure and solvency.

The other main kind of systemic event is not driven by any fundamental change in value of the institution or its assets, rather it originates merely from a change in beliefs in the market and the existence of multiple states of equilibrium in the connected system. The classic example of opinion-driven systemic risk arises via self-fulfilling behaviour. The most basic case is that of bank runs: banks transform the short-term exposure of deposits into the long-term one of illiquid investments, but this creates the risk that when too many people ask for their deposits back at once the bank is not able to service them immediately due to the illiquidity of their exposure. This is turn transforms expectations of illiquidity of the bank into reality and more panic. Another example of self-fulfilling behaviour comes when banks on the other hand no longer trust the solidity of the businesses they lend to, this causes them not to lend or raise rates which in turn causes businesses to default. As businesses default banks begin not to trust each other either and the whole credit market freezes. Similar cases to these consist of banks not honoring contracts with each other and thereby causing chains of defaults or being forced into a fire-sale of their holdings which in turns affects other banks and their obligations to each other.

It is important to highlight that, in spite of making a distinction in terms of high and low frequency domain in the structuring of our chapters, this particular section, as partly some others, applies to both high and low frequency domains in particular with regards to fire sales, as will be discussed in the following sections.

We note that the works mentioned above mostly originate from the econometric and financial streams of this literature. Another perspective on this theme which is increasingly considered by academics, practitioners and regulators is that of Complex Systems and Econophysics approaches to the analysis of networks in financial systems. Academics in this field often have a background in Statistical or Condensed Matter Physics and the methods used to analyse connected and complex systems in this context are often more scientifically rigorous, advanced and often provide unbiased, quantifiable results.

The early work by Aharony & Swary [56] investigates the contagion effects of the three major U.S. bank failures to date. The authors find that only banks which default due to issues which implicate other banks cause contagion effects. When the cause of default is specific such as in the case of frauds there is little to no contagion. The findings support the argument that similarity and correlations between institutions do indeed cause contagion effects. The related work by Furfine [57] cites [56] as an example of the established literature on how bank failures affect the financial system, but suggests to focus on a second less popular aspect at the time: the contagion risk due to the “real” network of interbank exposures. This would indeed develop into the main stream of literature on financial systemic risk.

Early Econophysics works observe the importance of network structure in contagion and robustness as discussed in Section 1.2 and perform analyses of network structures in finance and their implications for systemic risk, shocks and stability.

Angelini et al.[58] investigate interbank clearing networks and the impact of liquidity or solvency issues of a banks and how the shock spreads in the system. They verify that is such heteregonous systems only few players cause domino effects which are systemic. This result is in line with those from general network theory robustness discussed in Section 1.2 in [37, 46, 47, 36]. The heterogeneity of interbank networks and its effects were also studied in [59] and Boss et al.[60], where the authors also observe regional clusters of banks which form community structures and highlight how these networks are different from the homogeneous ones assumed by economic theory. A series of works from Iori and co-authors investigates heterogeneity not in the structure of connections and degree distributions, but rather in the size of the institutions in the network. They highlight in [61] that homogeneous systems are less prune to shocks, yet even for heterogeneous ones the stabilising role of the interbank market remains. The interplay between the advantage of pooling risk and the risk for cascading failures is also the focus of [62] where simulations show how homogeneous systems benefit in stability from wider systems with small avalanche effects. Avalanches become more prominent with system heterogeneity and the authors achieve a critical regime of power law avalanche sizes by tuning heterogeneity and connectivity. Power law avalanche sizes are common and of great interest in real world critical systems [1].

More recent works on interbank network data contribute practical insights and monitoring techniques. Iori et al.[63] apply techniques from statistical mechanics to the Italian overnight money market to investigate the evolution of network connectivity during the maintenance period. The authors evaluate the current arrangement and its implications for system stability. They also employ null models in the context of efficiency in the absence of speculative and preferential trading relationships.

Cimini et al.[64] address the issue of partial network information where often few (if any) links are known and only the degree and other properties of individual nodes are known. They propose a network reconstruction technique and test it on the International Trade Network and the E-mid market (Electronic Market for Interbank Deposits).

In terms of monitoring techniques and useful metrics Battiston et al.[65] introduce DebtRank, a novel measure of systemic impact inspired by feedback-centrality which suggests the transition from too-big-to-fail to too-systemic-to-fail.

The method is tested on the network of financial institutions in 2008 and shown to correctly identify the ones most involved in the $1.2tn FED emergency loans program.

## 1.4 Market microstructure and High Frequency Trading players

In Sections 2.3, 2.4 we provide an overview of the topics around our study of flash crashes and contagion in high frequency markets. These topics belong to the still emerging field of market microstructure, which we define below based on [66].

Financial markets can be studied from two main perspectives: by studying prices and diving into the underlying dynamics as a result, or by analysing micro-scale actions of market players and working from the bottom up. As suggested by its name, market microstructure chooses the latter path as it focuses on trading at the micro-scale. It is important to notice though that, in a similar fashion to statistical mechanics and related fields, one cannot reduce market microstructure as the study of markets at atomic scale; rather, its study at atomic scales provides insights into emergent macro-scale phenomena which have been poorly understood for decades.

In the present section, we provide a succinct overview of the main roles of market players in modern electronic markets and continuous-time double auctions.

We first have to define continuous-time double auctions, which are the mechanism most modern exchanges rely upon. These auctions are based upon the use of a publicly observable Limit Order Book (LOB), which is made of limit orders expressing the intention of a buyer or seller to transact a given quantity at a given price. Transactions occur in these auctions whenever an incoming buyer (or seller) agrees to buy (sell) at the prices and quantities available on the LOB. LOBs are updated in real time and analysing the flow of market orders, limit orders and cancellations (i.e. order flow) allows for the bottom up perspective on liquidity and price dynamics which market microstructure aims for.

Trades mostly occur in the LOB between a limit order sitting in the book and an order of opposite sign arriving at the same price. As limit orders in the book constitute liquidity in continuous-time double auctions, the issuer of the limit order is referred to as liquidity provider while its counterparty in the trade as liquidity taker.

Liquidity providers are modern days market makers which provide liquidity on both sides of the book and profit from earning the spread on round trip trades, while keeping inventory and hence exposure to price changes to a minimum.

Liquidity takers trigger trades in the market, which are traditionally categorised as:

- **Informed trades:** triggered by sophisticated traders with information about the future price of an asset with the goal to profit from the price change.
- **Uninformed trades:** triggered by either unsophisticated traders which are unable to process information about the future price of an asset or by liquidity trades which aim to liquidate assets, reduce portfolio risk or offload an inventory imbalance. Uninformed trades are also referred to as "noise trades" since they do not correlate with future price changes and seem to occur at random.

We would like to delve deeper into the business of market makers, as they will be the target of practical insights from our results.

Market makers are mainly of two types: traditional market makers (typically investment banks) are tightly regulated and obliged to always provide quotes to buy and sell assets they deal in to their clients, usually through an API. Modern day High Frquency Trading (HFT) liquidity providers instead provide liquidity directly on the exchange by posting (and cancelling) limit orders at high frequency at different levels of the book. They are responsible for most of the liquidity on modern exchanges and play what is called the "liquidity game".

From our brief statement above it may seem that if prices behave like a martingale (i.e. they have i.i.d increments) and market makers earn the spread they are in a business of risk-free returns! Further, competition for volume is fierce between HFT market makers and this has reduced spreads from historical levels of 60bps to a few units. These two empirical facts lead to the obvious question: why has the spread not converged to zero (at least for large tick stocks)? The answer is that also market makers experience risk mostly in the form of adverse selection and price impact. Adverse selection arises from the fact that market makers must post binding quotes which can be exploited by informed trades who see an opportunity to buy or sell at an advantageous price. Further, if quotes are not adjusted properly, market makers can experience highly unbalanced flows and accumulate an inventory position which is on the wrong side of the coming price change. To mitigate this risk, market makers adapt to order flow imbalances by updating their quotes in an attempt to reduce this. Hence, as more buy (sell) orders come they move the ask and bid prices up (down) to encourage sellers (buyers). This in return causes what is called price impact of a trade, i.e. the fact that buy (sell) trades cause the price to rise (fall). These dynamics are what the non-zero spread compensates for - the risk taken by market makers.

To introduce the liquidity game we discuss the concept of skewness of price impact, i.e. the fact that the distribution of price changes after a trade is heavy tailed and skewed in the direction of the trade. This means that unusually large price changes are not unlikely and are often caused by the action of highly informed traders. This poses a great risk to market makers which can be viewed as selling insurance, where small gains need to offset large and less frequent losses. For this reason, liquidity is fragile and disappears almost instantly when liquidity providers deem the risk too high. Due to these dynamics of fleeing liquidity, ensuring the proper functioning of markets is a highly complex and everchanging task.

Another perspective on players in high frequency markets is that of individuals, companies and external providers. We report below a summary following the structure in the book by Irene Aldridge [67].

High frequency trading firms compete with other investment managers for quick access and exploitation of market inefficiencies. Competitors in this field may be trading desks at investment banks and hedge funds. Some of the largest players in HFT markets are DE Shaw, Tower Research Capital and Renaissance Technologies. Major HFT market makers instead are the likes of Citadel Securities and Two Sigma Securities.

The capital used by these firms is provided by investors such as funds of funds, hedge funds looking to diversify, or proprietary firm capital as in the case of investment banks.

There are a range of services and providers involved in this business.

Execution brokers and ECNs (electronic communication networks) route and execute trades. Execution brokers are often major investment banks while ECNs are the likes of ICAP and Thomson/Reuters.

Broker-dealers also provide custody and clearing services, i.e. safekeeping of trading capital and trade reconciliation, respectively. Transaction fees compensate these players for the risks of being responsible for the assets and insuring against counterparty defaults.

HFT operations heavily rely on advanced hardware and software, we provide a summary of components of the latter.

- Generation of trading signals is at the core of HFT strategies and translates into portfolio allocations and P&L (Profit & Loss). Signals can be obtained from complex manipulations of tick data as well as information gathering from news, satellites and alternative data providers.

- Trades need to be executed with meticulous efficiency in HF markets, as profit margins on each short term individual trade are small. This is the task of optimal execution algorithms which aim to achieve a given price in a given timeframe by balancing aggressiveness and breaking trades into optimal lots.

- The last crucial component of software in such operations is risk management, especially because at such trading speeds human oversight is ineffective at best. These components ensure that the systems stay within P&L and risk bounds as well as handle fault events.

A final key role in the HFT business and the financial industry is that of the government and regulation, which introduces the next section. In terms of regulation HFT is treated analogously to day trading and abides by common trading rules, such as the fact that the HFT execution arms of hedge funds must be completely separate in order to ensure impartiality when servicing external clients (see again the case of Citadel Securities and Two Sigma Securities). An unsuccessful attempt to introduce regulation on transaction cost surcharges was made in 2009, which leads to our broader discussion of market structures, dynamics and regulation in the next section.

## 1.5 The effect of regulation and liquidity structures

We have just seen how liquidity and price stability are at the core of electronic exchanges and HFT. Here we report an overview of current market structures, regulations and their issues [68, 69]. We also summarise proposed regulatory improvements in the literature. This ties in with our own analysis and reflections from Section 7.3 later on. Modern electronic trading presents what is commonly referred to as "liquidity fragmentation" where, for instance, the US Equity market is split between over 10 public exchanges, 30+ dark pools and hundreds of internalising broker-dealers [70].

As the number of trading venues increases so do liquidity imbalances amongst them and the complexity of guaranteeing fairness and protection on execution prices. Dark pools present the additional issue of not displaying the price process and liquidity which allows more sophisticated players to blindside less sophisticated ones in a variety of ways. We will not discuss dark pools further here though, as they do not relate to our work directly.

Liquidity fragmentation sees conflicting opinions amongst academics on its effect on modern HFT markets, in this section we report two main works and show that they just look at different aspects of what "improvement" means. O'Hara & Ye [71] find that fragmentation generally reduces transaction costs and increases execution speed. They do observe an increase in short-term volatility, but somewhat compensated for by a more efficient random walk-like price process.

The finding of increased short term volatility brings us to the opposite perspective by Golub et al.[68] which, differently from [71], is written after the crash of

May 2010.

To understand their perspective we must first provide some background on Intermarket Sweep Orders (ISOs), Auto-routing and regulatory protection such as the Order Protection Rule.

Intermarket Sweep Orders are limit orders to be quickly executed in a specific venue even when another market is publishing a better quote. As these orders do not require computation and routing to achieve the best price they can be executed faster than regular market orders. In order to comply with Regulation NMS order protection obligations, ISOs need to concurrently send orders to other markets with Protected Quotations. For a more thorough overview of ISOs we refer the interested reader to [72, 73].

Compliance with the Order Protection rule is achieved through Auto-routing, which we outline as follows. When an order is submitted to exchange $A$ not displaying the NBBO, part of that order corresponding to the "Protected Quotation" is routed to another exchange $B$. Only once the Protected Quotation is cleared on "B" the order on "A" will be executed, regardless of liquidity on that or other venues. The order potentially rips through the book causing a single-venue mini flash crash.

Hence, allowing a market participant to submit an order on exchange $A$ specifically and not at the NBBO (National Best Bid and Offer) can cause the order to consume volume at deeper levels in the book in spite of liquidity being potentially available at better levels in the book on other venues.

The Order Protection Rule is at the core of Regulation NMS which aims to provide intermarket protection in US Equities. This regulation was likely appropriate in the past, but the current HFT environment and increase in the number of mini crashes shows that market players need a better picture of liquidity and how their orders take it. A simple suggestion in [68] is that the rule be extended to the Depth of Book and not only the top which would prevent trades at deeper levels in the book on a venue when liquidity is still available at a better price in another venue. This is though difficult to uphold and fleeting liquidity deeper in the book was shown to contribute little to overall liquidity [74]. Further, Fleeting Liquidity would likely cause orders to be cancelled before the routing took place anyway.

Golub et al.[68] also advocate for the removal of ISOs as players which use them are well aware that there is not enough liquidity in the venue and the order will aggressively consume multiple levels in the book. Further, their faster execution has been used by HFTs to profit from latency arbitrage.

In a nutshell the authors of [68] do not disagree on the empirical fact that execution is faster and spreads are tighter [71], rather they say that this comes at a price of better markets on average, but worse and more frequent extreme events. This can be improved through regulation.

Linton et al. in their chapter on regulation in HFT markets [69] begin with this distinction as well and recognise that the dynamics of electronic HFT markets are different to anything before. Regulators therefore have the problem of ensuring that liquidity and price discovery are constantly provided. One of the points made which we elaborate upon in Section 7.3 is that market dynamics are too fast now for traditional ex post regulation which must transition to ex ante with the support of technology and data-driven monitoring tools. Regulators have been asking these questions and we conclude the section with their proposed solutions so far. The "flash crash" commission in the US suggested enhanced market halts, circuit breakers and changes in trading priority rules as well as more information gathering and surveillance. The European Commission also suggested changes along those lines in the form of algorithmic trading restrictions, minimum resting times, market-maker obligations and the transaction taxes already proposed and refused in 2009 [67]. The UK commissioned the Foresight project, an extensive study to evaluate these proposals.

In practice though, when it comes to passing regulation on HFTs, diverging views on their role and appropriate regulation emerged, leaving the issue unresolved to date.

We direct the interested reader to the chapter by Linton et al. for a more thorough discussion as this is beyond the scope of the present work [69]. We do point out that most regulatory proposals take a very traditional finance or econometric approach to the problem, but we aim to highlight in Section 7.3 how the academic and industry approach to HFT has moved on to complexity and advanced empirical methods.

## Chapter 2

# Introduction to the thesis

I have provided context for the overarching themes of complexity and criticaly in this thesis in Chapter 1. In this chapter I narrow down our focus by introducing the topics more closely related to our work and results in this thesis. Sections 2.1 & 2.2 continue our discussion from Sections 1.1, 1.3 and introduce market structure estimation and denoising through complex networks. Further, we present industry standard practices for portfolio construction and innovative network-based portfolio weighting methods.

The remainder of the thesis is then structured as follows. Chapter 3 introduces the methodologies used in this thesis. Simplicial persistence in particular is one of the novel concepts and methods introduced in the works that make up this thesis. I then move on to presenting the range of results for this thesis. Chapter 4 presents results on the evolution of market structures with time, investigated through simplicial persistence and null models. We conclude the chapter with a proof of concept application of temporal persistence for portfolio construction. The following two chapters focus on results in the context of high frequency markets, in particular of flash crashes. Chapter 5 analyses the volume distribution of crash and non-crash periods, showing the unboundedness of the former. A different perspective on flash crashes is then provided in Chapter 6.2, where I investigate simultaneous crashes across assets and their structure. This phenomenon ties back to our interpretation of systemic risk in Chapter 4 and has been shown to be increasingly relevant in modern markets. Following the exposition of results in these three chapters, Chapter 7 takes a step back to discuss the main takeaways emerging from our results in the context of the literature and proposals by major authorities in the field. Finally, Chapter 8 summarises the findings presented in this thesis and offers perspectives for future work.

I highlight for the reader that the material and results presented in this thesis are based upon four published works by myself and my PhD supervisors Tomaso Aste and Paolo Barucca, as well as others still unpublished which are being prepared or reviewed for journal publication at this time. More specifically, [75, 76] focus on the market structure and simplicial persistence results from Chapter 4. The results on flash crashes in Chapters 5, 6.2 instead are based upon [77, 78], respectively. [75] was awarded "Best Oral Presentation" at *Complex Networks 2019* and published in the conference proceedings. I am also the recipient of the "Entropy award" for Best PhD poster presentation at the *Econophysics Colloquium 2021*, where I presented [77]. Finally, [78] was just published in *Entropy*.

## 2.1 Complex networks and market structure

We have seen in Section 1.3 how real networks in finance are of great importance for "big picture" evaluations of systemic risk, in particular for monitoring, regulation and enforcement. Most of investment finance though focuses on portfolios, market volatility and diversification, as per the discussion in Section 1.1. Market structures are often inferred from return time series of assets in the portfolio in the form of a covariance matrix. As briefly mentioned above (Section 1.1), in order to create sparse network structures from the fully connected weighted network of empirical covariances we apply network filtering techniques. These are algorithmic techniques which allow to filter the noise and obtain the underlying network structures behind covariance matrices.

A seminal work in Econophysics and network filtering is that of Mantegna & Stanley [32]. The authors introduce the Minimum Spanning Tree (MST) filtering technique which iteratively selects the highest weight links which preserve a tree structure (no loops) of the filtered network. Mantegna shows in his first followup applied work on the method [31] that the resulting tree provides a network of stocks with meaningful taxonomy of sectors and hierarchical structure. Closely following this, another followup work with Vanderwalle & Bonanni [79] investigated the structure of stock market indices globally. The correct time horizon and choice of reference currency allow the authors to find meaningful taxonomies through the MST.

A range of authors have conducted followup works leveraging the MST method. Vanderwalle [80] considers over 6000 U.S. stocks over two decades. Their correlation matrix and MST are used to show a heavy tailed degree distribution persistent across the decades considered suggesting a highly heterogeneous structure. Further, the structure is shown to evolve slowly and locally in time and non-trivial correlations are found in the seemingly chaotic market.

Bonanno et al.[81] conduct an extensive study which aims to show how markets are complex beyond the trivial models or randomness or single market factors. To do this they compare MST networks of real covariance matrices of returns and volatilites across global indices with the more simplistic random and single factor models. This work likely elaborates on the previous one by some of the authors [82] which had shown basic differences between the structure of real markets and artificial ones in terms of degree distribution, hierarchical complexity and large scale correlation properties characteristic of complex networks. Similar results are obtained by Plerou et al.[83] with a more involved methodology based on RMT. This shows the simplicity provided by the MST and sparse networks as well as confirming that the findings can be replicated with methods from other disciplines.

Other works by the co-founder of Econophysics Eugene Stanley (with Rosario N. Mantegna) such as [83] investigate market-based networks. The work by Wang et al.[84] considers the spillover network of extreme price fluctuations amongst financial institutions with an analysis of crisis periods of interest to regulators and policy makers. Chen et al.[85] also consider correlation networks and provide insights and methodologies to leverage network topology to monitor the market.

Temporal dynamics of networks have become of interest in general network theory as they allow to investigate the evolution of systems and their structural properties. It is intuitive how this transfers well to financial correlation networks as they map the evolution of market structures. The early work by Vanderwalle discussed above [80] was followed by a range of works investigating the temporal dynamics of MST-filtered networks. The degree distribution of networks and its properties are one of the main quantitative measures of structure, the authors in [86] consider the degree distribution of MST networks of markets with time. They find as per [80] a slow evolution in the degree of each stock with a timescale in the range of years. The evolution of filtered structures is also discussed in [87], where the authors propose variations on the original MST filtering technique and show how improvements lead to more consistent and stable structures in time. This supports the idea that market structures are somewhat stable and with "long memory" characteristics, which can become more evident as the filtering procedure improves.

When discussing temporal dynamics of networks in finance the evolution in time is not the only cross-section that matters. Early Econophysics works show that the distribution of asset returns at different horizons obeys a universal scaling law [88, 89]. The corresponding phenomenon in network structures is investigated in [90] where the authors investigate the MST of return correlations in U.S. stocks at different return horizons. They find hierarchical structures in high frequency as well, but show how this structure is markedly different as the time scale lengthens, with a more in-depth investigation of how the energy sector's role within the network evolves with time scale. Ideas of scaling in financial markets are also present in the early Econophysics literature [91].

Finally, the work by Bonanno, Lillo & Mantegna [87] investigates various aspects of complexity in financial markets. The authors show in particular the complex interactions of financial time series and how hierarchical filtering/clustering allows those structures to emerge more clearly from correlation matrices. A particular further level of complexity considered in the work is that of the change in market structures during extreme events. This leads well to our introduction of high frequency markets and extreme market events, as those and their structure will be the focus of our results on high frequency markets.

## 2.2 Portfolio construction in practice

In this last section we tie together the previous ones around portfolio construction, networks, market structures and risk. We do this with a focus on works which provide practical applications to portfolio construction and allow to show the reader the need for advanced methods and their advantages.

We start with a review of traditional methods and their applications in industry. The work by Satchell & Scowcroft [92] presents applications of what is known as the Black-Litterman model [93], a Bayesian portfolio construction procedure. The authors explain portfolio construction with these models and show their value as financial management tools.

In line with the approach of this work, from having shown how these methods are usefully applied in practice [92] we now move to question their correctness and robustness and provide alternative solutions or improvements. Axioma Inc. (now part of Qontigo)[1] is one of the largest providers of portfolio construction analytics. Its CEO and Director of R&D show extensive research on the topic in [94]. They analyse the negative impact of estimation error in mean-variance portfolio optimisation. As discussed in Section 1.1, asset weights are highly sensitive to small changes in the inputs. The authors propose a robust mean variance framework designed to explicitly account for parameter uncertainty in optimisation models.

We can see from the discussion on robustness of mean-variance methods, that recent practitioner insights still show the dominance of traditional methods in portfolio construction. Specialists in the field though continue to propose advanced solutions to solve practical problems with the estimation of structures and following risk optimisation.

We therefore review a range of RMT and network-based methods with innovative approaches to portfolio construction, these lead the way to our proof of concept portfolio construction approach presented in Section 4.6.

[1] Axioma portfolio optimizer: https://qontigo.com/products/axioma-portfolio-optimizer/

Some of the main authors of the financial networks literature such as Fabrizio Lillo and Rosario N. Mantegna co-author a work [95] which tests improved covariance estimators from RMT and shows they outperform empirical estimates in terms of realised risk when $T \sim N$ and in particular with short selling allowed. With $T > N$ they cannot outperform, but provide more diversified portfolios. With $T < N$ empirical estimates are very poor. A range of methods and solutions to reduce noise in empirical estimates (even in $T < N$) is proposed in works by Tomaso Aste with parsimonious local modeling [96, 97] and similarity clustering of market states [98].

Clusters and communities are a main feature of heterogenous network structures which we have seen in Section 1.3 to characterise risk in financial networks. The authors of [99] do not explaicitly use network filtering, rather hierarchical clustering and optimise the portfolio within each cluster. They show outperformance to mean-variance optimisation out of sample with 1-4 hours rebalancing periods.

On this line, Tola et al.[100] show how clustering algorithms applied to an MST-filtered correlation network help reduce statistical uncertainty in the construction of financial portfolios

A more network native and original approach is perhaps that of Pozzi, Di Matteo & Aste [101] (supporting information in [102]). The authors explore the idea that network cores are more connected and systemic in heterogenous networks and apply it to networks filtered with the Planar Maximally Filtered Graph (PMFG) method. They consider stock centrality as the mean of betweenness and eigenvector centralities and show how peripheral stocks yield better risk-adjusted returns than stocks in the core. This provides a fully network native investment procedure which can be tailored to individual investment objectives.

## 2.3 Microstructural dynamics of flash crashes

We recall from Section 1.5 that academics, regulators and practitioners alike have noted that price efficiency has improved as a result of HFTs, but also that exchanges, assets and even markets now carry more systemic risk as a result of the new dynamics of liquidity, risk constraints and market venues. The growth in systemic risk due to electronic trading has given rise to a growing number of increasingly large flash crashes [103] of which the one on May 6th 2010 was the first notorious example.

The report on the May 6th crash by Nanex [104] begins to suggest how high frequency order placement and saturation might have worsened the extent of the crash. The authors in [105] delve deeper into this idea and investigate flash crashes between 2006 and 2011 to show a system-wide phase transition around $\sim 500ms$ to an all-electronic trading market characterised by black swan events. Further insight into the impact of HFT players in flash crashes was provided by the simulations in [106], where the authors show how lowering the number of HFT players in the simulation reduces the extent of the crash, even when the size of the large sell order which triggered it is kept constant.

The authors explain the relation between the number of HFT players and the extent of the crash as due to the "hot potato" phenomenon described in [68]. This phenomenon can be explained as follows. When an unusually large sell (buy) order hits the market it gets absorbed by liquidity providers (often HFT market makers) which, as a result, accumulate a long (short) position in their inventory. The unusually large order though impacts the price, as per our discussion of price impact in Section 1.4, and potentially triggers risk limits by those market makers holding the inventory. Meanwhile as the price drops (rises) dramatically ordinary market players withdraw from the market. As a result of the risk limits HFTs try to reduce their inventory with aggressive market orders. As everyone else has withdrawn, they trade with each other at extremely high frequency, thereby creating high trading volumes. This particular phase characterises the "hot potato" phenomenon, i.e. the inventory exposure being passed around like a "hot potato".

The positive feedback loop continues as follows. When trading volume is used as a proxy for liquidity the "hot potato" phenomenon creates apparent liquidity which triggers execution by lower frequency players which are also trying to reduce exposure as the market drops (rises).

The authors in [68] make two further points on HFTs and modern market dynamics during these extreme events. Another cause of apparent liquidity and market depth is Fleeing Liquidity by HFT market makers through the fast cancellation of limit orders which was suggested to create the illusion of market depth. This was touched upon earlier in Section 1.4.

As the number of trading venues and exchanges rises liquidity is fragmented. Sweep orders and arbitrageurs aim to move liquidity and remove arbitrage opportunities across exchanges, but can create dangerous effects. Sweep orders between markets operate at lower frequencies than these fast cancellations, thereby sweeping the book after liquidity has potentially already been removed. This worsens the systemic aspect of these events across the fragmented liquidity structure of exchanges.

Extreme events of this kind, which are characterised by non-linear reactions to shocks in the system and positive feedback loops, exist in man-made and natural systems and are instances of self-organised criticality. In this type of systemic events the system reaches a critical state where a small release in energy or imbalance triggers highly non-linear reactions in size. This is the case in avalanches, earthquakes and sandpiles [1, 107, 108].

Events with such underlying dynamics are characterised by heavy-tailed power law distributions. The tail's decay exponent in these distributions is crucial for their modeling as it indicates which moments of the distribution are defined and which diverge. This can be of great practical relevance for both regulators and practitioners, in particular traditional market makers. These have an obligation to provide liquidity at all times and cannot simply withdraw from the market under extreme conditions. This constraint brings the need to model trade volume flow in the market as well as from clients in order to provide actionable pricing also under extreme market conditions. Predictive tools for volume modeling which rely on advanced statistical methods and Artificial Intelligence are now being adopted by most market makers, but a deep understanding of the underlying data distributions are crucial to correctly fit and deploy these models. For example, unbounded distributions have major implications for the use of Gaussian-based versus distribution-agnostic (quantiles) confidence intervals and how extreme values are modeled and accounted for in terms of pricing and inventory management. Our work makes the argument that power law distributions and critical dynamics must be included in modern modeling, in particular when dealing with extreme events is the focus, as in risk modeling.

In the above discussion, we have shown the relevance and interest in these extreme events and how their dynamics were investigated mostly through simulations. In spite of their importance, to our knowledge no investigation of the statistical properties of these events in their microstructural dynamics has been conducted yet. The statistics underlying extreme events are of great relevance for the study of the dynamics of these phenomena, but mostly for market makers who aim to provide liquidity and stability during such events and should have the tools to model and best navigate them.

The results in Section 5.2 bridge this gap through analyses of flash crash events in high frequency markets from the perspective of critical systems and avalanche-like dynamics to show how the detected flash crashes constitute black swan events with unbounded trading volume distributions while the remaining order flow is sub-critical and bounded. This is of great importance for the modeling of order flow (and related assumptions) in both normal and extreme market conditions.

## 2.4 Synchronisation and co-crash structures

From the discussion in the previous section, we see that crashes of different sizes seem to involve a self-perpetuating cycle [106] with positive feedback loops.

This type of self-excited process is also investigated in [109] for the liquidity and information dependence between two sample assets, showing how liquidity shocks to an asset can propagate to related ones (and by extension to the wider market).

The frequency and size (in terms of number of securities involved) of simultaneous crashes in HFT is also investigated in the literature. For instance, the works by Lillo and co-authors [103, 110] investigate the dynamics of simultaneous flash crashes and motivate their importance by showing the growth in the number of mini crashes in recent years. They also show that the number of simultaneously crashing securities has grown over the last 10 years, thereby highlighting the increasing systemic relevance of this phenomenon.

One of the perspectives on systemic risk in this work, which will be discussed further in Sections 7.4, 7.1, is that of the risk component of an event (say a flash crash) that is given by the asset's influence over other assets and their interconnectedness. This causes the isolated event to spread in the market and affect more assets, thereby increasing its impact and relevance for all market participants. A related concept is that of "synchronisation" which is the systemic and concentration aspect that arises from the alignment and interdependence of actions between market players rather than assets. This will also be a topic of discussion later on in Section 7.4.

The systemic risk posed by HFTs has been investigated in the literature in the last decade. The work by Paulin et al.[111] simulates flash crashes through agent based modeling and highlights the importance of market structures in the systemic propagation of extreme events. The works by Abreu & Brunnermeier [112] and Bhojraj et al.[113] investigate the risks of synchronisation of arbitrageurs in financial markets and acknowledges the phenomenon. Other works investigate the systemic risk of HFT dynamics. Jain et al.[114] investigate how low-latency HFT trading can worsen extreme systemic events in financial markets and argues for the need to incorporate correlation and market structure in regulating these risks. The work by Harris [115] discusses many mechanisms, among which, systemic risks originating from order routing and self-reinforcing mechanisms causing crashes. The review by De Gruyter [116] summarises systemic aspects of HFTs and market structure such as position correlation and herd behaviour, adverse selection in orders and crowding as well as negative contribution to price discovery, at times.

In summary, co-crashes are becoming more frequent and systemic. It is therefore important to investigate their structure. In particular, it is relevant to understand which stocks are central to larger systemic events as well as the contagion structure between stocks in the market. This is a central theme in market stability for regulators as well as in risk management for market makers.

The results in Section 6.2 join the two themes of flash crashes and cross-asset systemic risk by delving deeper into the dynamics of simultaneous flash crashes of different sizes throughout 300 liquid stocks traded on the NASDAQ Exchange. We investigate the empirical distribution of crash sizes and the structure of these events in the market. We also investigate whether larger systemic events involve highly unstable stocks (which crash often) or stocks that are more stable in their price dynamics, yet more influential to trigger larger systemic events when subject to liquidity shocks. We apply tools from statistical physics to show the difference between crashes which involve a small or large number of assets. We uncover a phase transition occurring when the crash size exceeds 5 companies. Implications for systemic risk in high frequency markets are discussed from both trading and regulatory perspectives.

# Chapter 3

# Method

The literature on portfolio construction and risk reviewed in Section 1.1 focuses on correctly estimating and denoising the covariance matrix of asset returns, which represents the structure of the market. This estimation in crucial as the covariance matrix is the main input to portfolio optimisation algorithms. With this is mind we outline our methods below as follows. Section 3.1 describes our estimation of the covariance matrix with exponential smoothing. Section 3.2 the describes our "de-noising" methods which correspond to quantile thresholding and Triangulated Maximally Filtered Graph (TMFG) filtering. The most relevant simplicial substructures are then obtained from those top ranked in temporal persistance, defined in Section 3.3. These filtering methods can be questioned, as one might ask if it is not the method which yields the observed results. To remove these doubts and investigate the generating process of market structures we study and compare the null models presented in Section 3.4. We conclude by introducing a simple measure for portfolio construction in Section 3.5.

Much of the work in market microstructure deals with the complex relations between order flow activity and price returns. This is the case as most applications in this field relate to optimal order execution and market making, where price impact is key. We too investigate the relation between price returns and order flow in this work. We consider anomalous price returns (both positive and negative) and the traded volume during such return intervals. The LOB and order flow (OF) are the key representations of market state and activity in high frequency finance, we begin by introducing those. We then analyse the distribution of traded volumes. We focus on trades in particular, based on findings from the literature on market microstructure which has identified them as the most informative type of market event in the OF [117], upon which much of the market impact literature is based [118, 119, 120]. Still on the topic of flash crashes, we consider instances where multiple assets crash together and the distribution of such co-crash sizes. Finally, we use null models (introduced in Chapter 1) to show the significance of our results.

## 3.1 Market structure estimation from price data

Correlations are noisy measures of co-movement in financial asset prices, which are often non-stationary within the observation window. Longer time windows benefit the measure's accuracy, as we have more observations to estimate the $N(N-1)/2$ parameters of the matrix of $N$ assets. However, a longer observation window can come with the disadvantage of weighting more and less recent co-movements equally with the risk of averaging over a period in which the values are non-stationary. In order to compensate for this effect, we apply the exponential smoothing method for Kendall correlations [25]. As already suggested in Section 1.1, this method yields more accurate correlation estimates, as it applies an exponential weighting to the correlation window, prioritising more recently observed co-movements.

## 3.2 Network filtering: quantile thresholding and the TMFG

To obtain sparse network structures from the correlation matrices, we apply two filtering techniques with fundamental differences. The first filtering method is quantile thresholding, which corresponds to hard thresholding to generate an adjacency matrix through the binarisation of individual correlations. We consider the $N \times N$ correlation matrix $\mathbf{A}$ and its elements $\{A_{ij}\}_{i,j=0}^{N}$. The quantile $q$ of $\{A_{ij}\}_{i,j=1}^{N}$ is then $v_q$, where the adjacency matrix is defined as

$$A_{i,j} = \begin{cases} \rho_{i,j} \geq v_q, & 1 \\ \rho_{i,j} < v_q, & 0 \end{cases} \tag{3.1}$$

This filtering technique is entirely value-based with no structural or other constraints. We apply it by providing a quantile level $q$ which yields edge sparsity analogous to that of the corresponding TMFG filter.

The second filtering technique is the TMFG method [35]. This topological filtering technique embeds the matrix with topological constraints on planarity in a graph composed by simplicial triangular and tetrahedral cliques. Edges are added in a constrained fashion with priority according to their (absolute) value. The graph essentially corresponds to tiling a surface of genus 0. This technique represents a filtering method that accounts for values, but also imposes an underlying chordal structural form which might help regularising the filtered graph also for probabilistic modeling [34]. Furthermore, this technique imposes higher order structures, namely triangles and tetrahedra, which are known to be a feature of financial markets and social networks.

## 3.3 Simplicial soft persistence

We focus on temporal persistence of tetrahedral and triangular simplicial complexes (motifs) in the TMFGs and graphs filtered via quantile thesholding constructed from correlations over rolling windows as per Section 3.1. TMFG networks can be viewed as trees of tetrahedral (maximal) cliques connected by triangular faces, these are triangular cliques called separators, which different meaning in the taxonomy. As the TMFG is a tree of maximal cliques, the removal of a separator split the graph into two parts. We point out that not all triangular faces of the tetrahedral cliques are separators and we will refer to those which are not merely as triangles.

We point out though that both triangles and separators were considered for the results in Section 4.3, in order to account for all triangles in the filtered graph, as quantile thresholding does not distinguish between triangular faces and separators.

When investigating the time evolution of a graph, the literature often considers "survival", i.e. hard persistence, of an edge or a motif in time [121, 122]. This implies that to be persistent between times $t$ and $t+\tau$ the edge needs to be present in the graph at all intermediate times $\in [t, t+\tau]$. We propose the concept of "soft" persistence and argue that in highly stochastic systems where estimation is noisy this measure is more appropriate to track long term structural evolution. We define "soft" persistence as follows.

A motif corresponding to clique $\mathscr{X}_c$ is considered soft-persistent between time $t$ and $t+\tau$ if and only if the motif is present at both the initial time $t$ and at $t+\tau$. A visual intuition for motif (triangle) persistence through time is provided in Figure 3.1.

**Figure 3.1: Motif persistence visualisation.** Visual representation of a TMFG structure's motif (triangle) persistence in time. The green triangle in figure a) is persistent through figure c), while other two triangles (present in figure a) within the red triangle) do not persist due to the rewiring of an edge. Figure b) shows one of non-persistent triangles with dashed contour. The rewired edge is also dashed and colors are used for visualisation purposes to "track" the triangles through the rewiring. This visualisation aims at showing the impact of edge rewiring on motif persistence and the difference between edge and motif persistence.

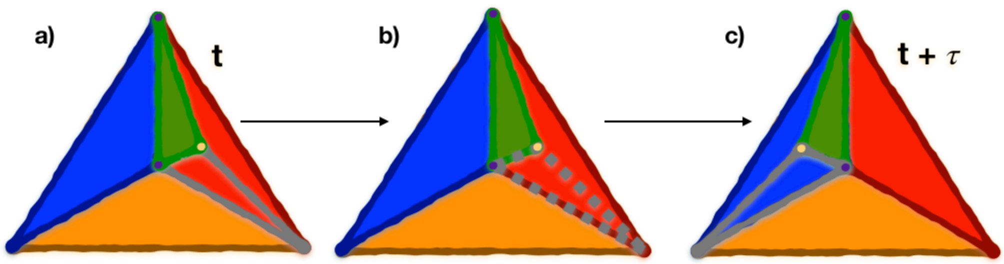


We investigate the decay in the number of persistent motifs between filtered correlation networks with observation windows progressively shifted by one trading day and we quantify how the average persistence decays with the time shift $\tau$.

As mentioned above, here we use a form of soft persistence which is different from hard persistence (survival) of motifs which is more common in the literature [121, 122].

The binary persistence value of motif $c \in C$ at time $t$ and $t+\tau$ is

$$P_m(\mathscr{X}_c^{t,t+\tau}) = (\mathscr{X}_c \in \mathscr{X}_C^t) \wedge (\mathscr{X}_c \in \mathscr{X}_C^{t+\tau}). \tag{3.2}$$

The average persistence for the entire clique set over $T$ starting points at time shift $\tau$ is

$$\langle P_m(\mathscr{X}_c^\tau) \rangle_{T,C} = \frac{1}{T} \cdot \frac{1}{|C|} \cdot \sum_{t=0}^{T} \sum_{c \in C} P_m(\mathscr{X}_c^{t,t+\tau}). \tag{3.3}$$

Where we consider the motif sets $\mathscr{X}_C^t = \{\mathscr{X}_i^t\}_{i=c_1,\dots,c_{|C|}}$ and $\mathscr{X}_C^{t+\tau} = \{\mathscr{X}_i^{t+\tau}\}_{i=c_1,\dots,c_{|C|}}$.

We observe that persistence decays as a power law with two regimes: one with a faster decay followed by one with a slower decay. The transition point between these two regimes, $\tau_{plat}$, is computed by minimising the unweighted average mean squared error (MSE) between the two power law fits over all possible transition points in time. The average motif persistence in the plateau regime is defined as

$$\langle P_m(\mathscr{X}_c) \rangle_{T,\mathscr{T}} = \frac{1}{T} \cdot \frac{1}{\mathscr{T} - \tau_{plat}} \cdot \sum_{t=0}^{T} \sum_{\tau=\tau_{plat}}^{\mathscr{T}} P_m(\mathscr{X}_c^{t,t+\tau}). \tag{3.4}$$

In order to verify that the persistence decay of motifs is not simply the consequence of the persistence decay of individual edges, we test the null hypothesis that motifs are formed by edges in the network whose existence is not mutually dependent. The hypothesis being tested is that motif persistence simply results from persistence in its component edges $(c_1, c_2, c_3)$:

$$P_m(\chi_c^{t,t+\tau}) = P_m(\chi_{c_1}^{t,t+\tau}) \cdot P_m(\chi_{c_2}^{t,t+\tau}) \cdot P_m(\chi_{c_3}^{t,t+\tau}), \tag{3.5}$$

We also compare the decay exponents for multiple random stock selections over different markets to identify whether the steepness of motif decay (edge, closed triad or tetrahedron clique) is indicative of market stability/development stage. We further investigate more liquid markets such as the NYSE from both a quantitative and qualitative point of view.We classify motifs in the plateau by their soft persistence and study the sector structure of the most persistent motifs. We also verify that these motifs are not trivially retrieved by maximum correlation edges or motifs in the correlation matrix.

## 3.4 Null models of market dynamics

We have defined null models early in Chapter 1 and discussed instances of their application in the literature. Null models are the basis of statistical testing for most results in this thesis, in line with common practice in complexity science. The null models described below are calibrated on empirical data from Section 4.1 and enforce constraints to test whether these are sufficient to reproduce the empirical observations (the persistence in this case) or not. We consider null models increasingly more constrained to show that, even imposing full pairwise structure in the generating model for time series, we are unable to replicate empirical persistence in null models. This provides a strong argument for the investigation of simplicial complexes beyond individual edges (i.e. motifs) and their long memory processes.

**Random return shuffling** Individual stock log-return ($r_t = \log Price_t - \log Price_{t-1}$) time series are randomly shuffled, i.e. a random permutation along the time dimension of each variable is applied, to obtain a null model for noisy, spurious correlations. This model maintains the overall statistics of the values of each time series but eliminates any correlation structure.

**Rolling univariate Gaussian generator** We calculate the rolling mean $\mu_{t-\delta_t,t}$ and standard deviation $\sigma_{t-\delta_t,t}$ of the log-return series over the time window $(t-\delta_t, t]$ for each node (stock) separately. We then generate ensembles by sampling the return $r_t$ at each point in time from the (rolling) univariate Gaussian distributions with sample mean and variance $r_t \sim \mathcal{N}(\mu_{t-\delta_t,t}, \sigma^2_{t-\delta_t,t})$, with $\mathcal{N}(\mu, \sigma^2)$ being a normal distribution with mean $\mu$ and variance $\sigma^2$. This intends to simulate the process as a simple moving average with uncorrelated time-varying Gaussian random noise.

**Fixed Multivariate Gaussian generator** We calculate the mean $\mu$ (for each node) and covariance matrix $\Sigma$ throughout the whole length of the log-return time series. We then generate ensembles by sampling the vector of returns $r_t$ at each point in time for all stocks from the fixed multivariate Gaussian with empirical means and covariance matrix, $r_t \sim \mathcal{N}(\mu, \Sigma)$. This aims to represent an underlying fixed market structure with sampling noise.

**Rolling multivariate Gaussian generator** After obtaining the log-return time series, we calculate the rolling mean $\mu_{t-\delta_t,t}$ (for each node) and covariance matrix $\Sigma_{t-\delta_t,t}$ between the series. We generate ensembles by sampling the return at each point in time $r_t$ for all stocks from the (rolling) multivariate Gaussian distributions with sample means and covariance matrices $r_t = \mathcal{N}(\mu_{t-\delta_t,t}, \Sigma_{t-\delta_t,t})$. This intends to detect the changing market structure and simulate the process as being generated by a multivariate Gaussian distribution with time-varying constraints on structural relations.

## 3.5 Portfolio construction methods and risk measures

In order to provide an application to systemic risk, we show that indeed the most persistent motifs represent clusters of systemic risk and high volatility. To do so, we construct a portfolio containing all stocks in the ten most persistent motifs in the plateau region, as defined in Equation 3.4 (for each market). We then compare its volatility with that of random portfolios with the same number of assets.

We conclude by defining the persistence measure $P_m(v_i)$ in Equation 3.6 to compare random portfolios weighted by $1/\sigma^2$ with those weighted by $1/P_m(v_i)$. This persistence weighting is still simple and asset-wise, but it accounts for structure in that it is based upon a non-local network measure. We do this for the four different markets, with all results showing meaningful volatility reductions.

$$P_m(v_i) = \sum_{\mathscr{X}_c \in \mathscr{X}_C | v_i \in \mathscr{X}_c} \langle P_m(\mathscr{X}_c) \rangle_{T,\mathscr{T}} \tag{3.6}$$

The measure presented in Equation 3.6 is defined for each vertex $v_i$ in the network as the sum over all $\langle P_m(\mathscr{X}_c) \rangle_{T,\mathscr{T}}$ (average pesistence of motif $\mathscr{X}_c$ in the plateau) where vertex $v_i$ belongs to clique $\mathscr{X}_c$.

## 3.6 Limit order books and order flow

In this section, we provide an introduction to limit order books, their representation in full and reduced-form, including LOB states and OF.

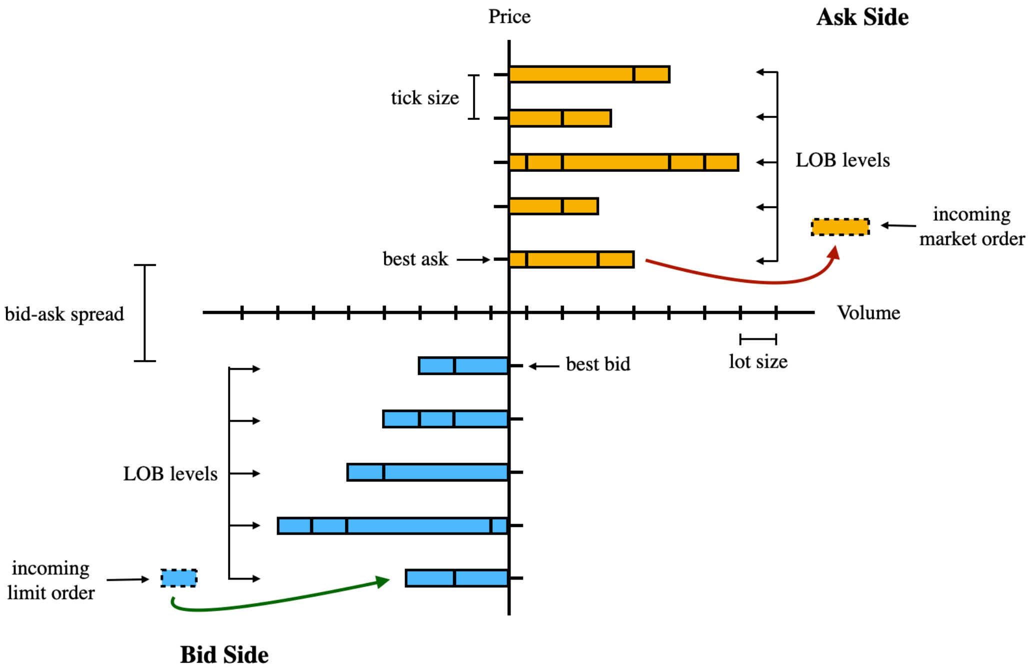


**Figure 3.2: An illustration of a limit order book.** We observe a market order being filled at the ask side and an incoming limit order being added to the bid side. The quantity of the market order then becomes traded volume.

Modern equity trading is conducted electronically. Major exchanges in the U.S. and the rest of the world facilitate this via a LOB, one per stock. The order book represents a collection of buyers and sellers, ordered by price and time, bidding and offering stock for purchase or sale. Figure 3.2 depicts an order book where buyers and sellers are shaded in blue and yellow, respectively.

At a given time $t$, the highest price buyers are prepared to buy the stock for is called the *bid price* and is denoted by $b_t$. Similarly, the lowest price sellers are prepared to sell the stock for is referred to as the *ask price* and is denoted by $a_t$. From the limit order book, we can derive the *mid-price*, $p_t := (b_t + a_t)/2$, the *bid-ask spread*, $a_t - b_t$, and the *tick size*, the smallest price increment between different price levels in the order book, as illustrated in Figure 3.2. For the Nasdaq exchange, where the data described in Sections 5.1, 6.1 is obtained, the tick size is \$0.01.

Beyond the bid and ask, there are additional *levels* in the order book.

An *order* is defined as the four-tuple (side, quantity, price, time), representing the side of the order book the order is posted to, the price at which the order is submitted, the desired amount to be traded and the time of submission. Orders may be entered, and if active, cancelled at any time. When an order is submitted the matching engine of the exchange attempts to match it with existing orders in the book. Orders which match and result in the full quantity being traded immediately are called *market orders*. Orders which do not match, or only match partially, are referred to as *limit orders*. When there are multiple limit orders with the same side and price present, they are queued chronologically at that price level according to the first-in-first-out (FIFO) principle. For an in-depth description of modern equity trading and the order book we refer the interested reader to [123, 124, 66].

## 3.7 Jump detection

We detect price jumps (up and down crashes) similarly to [103] in 1 minute non-overlapping returns, as per Sections 5.1, 6.1. We apply the basic jump detection method from [125] and detect jumps at the 1% significance level. In addition to the basic features of the method for robust volatility estimation we obtain an estimate of intraweek periodicity and adjust the return series and jump detection according to [126] as well.

### 3.7.1 Power law fitting and exponents

We analyse the volume distributions and fit power law functions to the tails of the empirical Complementary Cumulative Distribution Function (CCDF). For fitting purposes, we define the tail as the percentile interval $[90\%, 99.9\%]$, where the highest values are excluded to reduce fitting noise and discount finite size effects. We highlight that the results are robust to different choices of tail quantile intervals, as will be discussed in Section 5.2. The fit is then simply applied as a linear fit (minimum squared error) to the log-log data. In order to support the robustness of our result we also apply the method described in Clauset et al.[127] to obtain a second fit and set of exponents consistent with the first ones. To do so we use the package from Alstott et al.[128].

## 3.8 Crash size distribution

We investigate whether co-jumps which involve different numbers of stocks originate from the same dynamic process and present the same distribution. We also consider whether individual stocks are involved to the same extent across co-crashes of different sizes or if a pattern emerges.

We define the unnormalised crash frequency for stock $x$, in co-crashes with $m$ stocks as

$$f_{x,m} = \sum_{t=0}^{T} c_{x,t,m} \tag{3.7}$$

with

$$c_{x,t,m} = \begin{cases} 1, & \text{if stock } x \text{ is involved in a crash of size } m \text{ at time } t \\ 0, & \text{otherwise} \end{cases} \tag{3.8}$$

By marginalising over the ensemble of stocks $x$ we obtain the frequency distribution across co-crash sizes

$$f_m = \sum_{x} f_{x,m} \tag{3.9}$$

The changes in the composition of the crashes are investigated by computing the correlation between the involvement of firms across crashes of different sizes. Namely, for each crash size $m$ we assign to each firm $x$ a rank in decreasing order by $f_{x,m}$. We then compute the Spearman correlations between these ranks.

## 3.9 Null models and statistical testing

To support the visual intuition of our results we apply statistical testing in the form of null models.

In the first null model we apply, the Spearman correlation is used to measure rank similarity between the crash frequency distributions across stocks at different crash sizes $m$. As the frequency distributions are noisy and fat-tailed the correlation p-values from the t-test seem hard to justify as a reliable measure. Hence, we create a simple null model of correlation significance through a biased reshuffling analogous to those in Section 3.4.

To do so we sample without replacement the whole list of stocks $S_m$ according to $\propto f_m$ from Section 3.8 to obtain a biased reshuffling $G_{i,m}$ of the stocks according to their crash frequency.

For each shuffled list we calculate the Spearman correlation coefficient between the sample and the original list to form the null distribution as:

$$D_m = Spearman(G_{i,m}, S_m)_{i=1}^{10^5} \tag{3.10}$$

where the Spearman correlation can be defined as the Pearson correlation between the rank variables:

$$Spearman(X,Y) = \rho(R(X),R(Y)) = \frac{Cov(R(X),R(Y))}{\sigma_{R(X)} \cdot \sigma_{R(Y)}} \tag{3.11}$$

where $R$ is the rank operator (convert values to rank values), $\rho$ is the Pearson correlation, $Cov$ is the covariance and $\sigma$ is the standard deviation.

Given the null model distribution from Equation 3.10 we then define the significance of the correlation between sizes $m, m+\tau$ (i.e. the p-value) as the quantile of $Spearman(S_{m+\tau}, S_m)$ in $D_m$.

The second null model we apply is simpler and aims to verify that the increase in the fraction of flash crashes involving the top $N$ liquid stocks with crash size is not just an artifact of the growing crash size (ie more stocks involved in the crash and hence more likely to select one of the top $N$ most liquid). To do so we choose $N$ random stocks and compute the fraction of flash crashes with at least one of them across crash sizes. The distribution of the percentage of crashes with at least one of the $N$ sampled stocks for each crash size constitutes the null model.

## 3.10 Crash-weighted trading volume

To investigate the relationship between crash size and the involvement of highly traded stocks we define a weighted average daily Dollar Traded Volume for each crash size, where the weighting is given by the normalised crash frequency of each stock.

For crash size $m$ and crash frequency distribution $f_{x,m}$, as per Section 3.8, we define the crash-weighted Dollar Traded Volume $\mathrm{DTV}_m$ as:

$$\mathrm{DTV}_m = \frac{\sum_x f_{x,m} \mathrm{DTV}_{x,m}}{f_m} \tag{3.12}$$

This measure aims to represent how more highly traded stocks are involved at different crash sizes.

# Chapter 4

# Simplicial persistence: from long memory to portfolio construction

The main findings of this thesis on market structures and portfolio construction are described in this section. We begin with an overview of results on the long memory of edges and simplicial complexes in TMFG-filtered correlation networks. The section continues with an analysis of null models of financial market structures and a comparison with real data to gain insights about the generative process of its stochastic structure. We then suggest how soft persistence captures the underlying change in market structure by relating its decay exponent to the efficiency (a proxy for stability) or average traded volume in the market (a proxy for fluidity which yields well-defined stable structures). We analyse the taxonomy emerging from persistence by showing that the most persistent motifs correspond to stocks in the same sector with strong underlying price drivers. We conclude the section with results in systemic risk applications to financial portfolios, where we show that the most persistent motifs present strong co-movement dynamics from the portfolio of 10 most persistent motifs being highly volatile and systemic. This leads to an idea for portfolio optimisation, where we propose a non-local portfolio weighting measure based on persistence and compare it with volatility weighting to show improvements in out of sample portfolio volatility.

This section (Section 4) presents our results on network filtering and persistence in market structures and their systemic risk. In line with our approach in the introduction where Section 1.2 provides context from the physical sciences for our methods and leads to Section 1.3 which applies them to financial markets, we begin with a contextualisation and comparison of our findings in the context of statistical physics and long memory processes in its structures.

Later in this section, we make use of network filtering to investigate the temporal evolution of the market structure and we compare the persistence of certain network motifs with different levels of market efficiency. Market efficiency can be viewed as an emergent property of a system whose state is determined by the interaction of multiple agents which compete to exploit the system's inefficiencies, thereby making the system's dynamics more random as predictability in those dynamics is exploited (and thereby reduced) until no longer present. Liquidity in markets can be viewed as the level of fluidity and granularity of the system's dynamics. More fluid markets are characterised by more agents and interactions which reduce noise in the dynamic process of the emergent variable (the price).

Further, the persistence that we analyse in these systems can be compared with the autocorrelation in spin glasses. Indeed, in the physics of spin glasses, slow relaxation and correlation persistence have been studied for a very long time [129, 130]. Simulation results and experiments reveal that the dynamic correlation function decays as a power law in the proximity and below the glass transition [131] in a similar way to what we report in the present work. In such systems one also observes aging effects which show an initial exponentially fast decay of the spin-spin autocorrelation function and then a freeze into slow dynamics [129]. Intriguingly, this was observed also in simulated network systems [132] and in financial systems in the present work.

Edge persistence is also studied in Network Science in the form of survival (hard persistence) [133, 134, 135, 136] as well as in Econophysics [137, 138, 139, 140, 141]. Further, [142] defines one step node persistence analogously to the the definition for edges in the present work.

The main findings of this thesis on market structures and portfolio construction are described below. We first introduce the data used to produce the results in this chapter in Section 4.1. We then begin with an overview of results on the long memory of edges and simplicial complexes in TMFG-filtered correlation networks in Section 4.2. Section 4.3 proceeds with an analysis of null models of financial market structures and a comparison with real data to gain insights about the generative process of its stochastic structure. We then suggest how soft persistence captures the underlying change in market structure in Section 4.4, by relating its decay exponent to the efficiency (a proxy for stability) or average traded volume in the market (a proxy for fluidity which yields well-defined stable structures). In Section 4.5 we analyse the taxonomy emerging from persistence by showing that the most persistent motifs correspond to stocks in the same sector with strong underlying price drivers. We conclude with results on systemic risk applications to financial portfolios in Section 4.6, where we show that the most persistent motifs present strong co-movement dynamics from the portfolio of 10 most persistent motifs being highly volatile and systemic. This leads to an idea for portfolio optimisation, where we propose a non-local portfolio weighting measure based on persistence and compare it with volatility weighting to show improvements in out of sample portfolio volatility.

## 4.1 Data

We select the 100 most capitalised stocks from four markets: NYSE, Italy, Germany and Israel (400 stocks in total). These markets range from highly liquid and more developed ones such as the New York Stock Exchange and the Frankfurt Stock Exchange to less liquid markets such as the Italian Stock Exchange and the Tel Aviv Stock Exchange.

We investigate daily closing price data[1] for:

- New York Stock Exchange (3/01/2014 - 31/12/2018);
- Frankfurt Stock Exchange (3/01/2014 - 28/12/2018);
- Borsa Italiana (Italian Stock Exchange) (3/01/2014 - 28/12/2018);

[1] Source: Bloomberg Finance L.P.

- Tel Aviv Stock Exchange (5/01/2014 - 1/1/2019).

The data includes 1258 daily prices observations for the NYSE, 1272 for FSE and BI and 1225 for TASE.

To estimate market structures from the return series of the data we compute correlation matrices with exponential smoothing from rolling windows of $\delta = 126$ trading days with a smoothing factor of $\theta = 46$ days, as per [25] and our discussions in Sections 1.1, 3.1. This is done for all realisations of each null model ensemble (see Section 3.9) and for the real data.

## 4.2 Long memory of motif structures

Market efficiency imposes the absence of temporal memory in price returns, but the presence of long memory in higher-order moments of returns and long-term dependence (autocorrelation) of absolute and squared returns have been observed and are now considered among the important stylised facts in markets [143]. For instance, in [144, 145], later extended in [118], it was shown that order signs obey a long memory process, balanced by anti-correlated volumes which guarantee market efficiency. In financial time series analysis, through the generalised Hurst exponent analysis, it was observed that memory effects are related to the stage of maturity of the market, with more mature markets being more random [146]. However, these memory effects have been so far observed on the stochastic evolution of each single variable and not on the collective evolution of the dependency structure of the market.

With results in the present section we provide the missing piece, connecting market structure and market memory by analysing the autocorrelation of market structures [147], through persistence of its filtered correlation matrix [148].

The plot in Figure 4.1 shows the power law decay (evident from the linear trend in log-log scale) in $\langle P_m(\mathscr{X}^{\tau})\rangle_{T=200,C}$ vs. $\tau$, followed by a plateau region that also decays as a power law, but with a smaller exponent. We also observe that all motif decays have $\tau_{plat} \in [\delta t_{window}/2, \delta t_{window}]$, where $\delta t_{window}$ represents the length of the estimation window of the correlation matrix. The window used has

$\delta t_{window} = 126$ trading days and a value of $\theta = 46$ for exponential smoothing, as per [25]. The choice of $\delta t_{window}$ corresponds to roughly 6 months of trading and satisfies $N < \delta t_{window}$, with $N$ the number of assets in the correlation matrix. The correlation matrix is hence well-conditioned and invertible. On the other hand the exponential smoothing with $\theta = 46$ mainly considers recent observations from the latest few months.

There are $N-3 = 97$ tetrahedral cliques in the starting TMFG networks and $3N-8 = 292$ face triangles.

**Figure 4.1: Temporal Decay of Simplicial Soft Persistence.** Decay of triangular clique faces, separators and clique motifs persistence for 100 NYSE stocks, as a function of time interval $\tau = [0, 900]$ (average over 200 starting points). The motifs and network are obtained via TMFG filtering. The two power-law regimes are identified by the minimum MSE sum of the fits. The second region in termed "plateau" in contrast with the strong decay observed in the "decay" region although we recongnise that

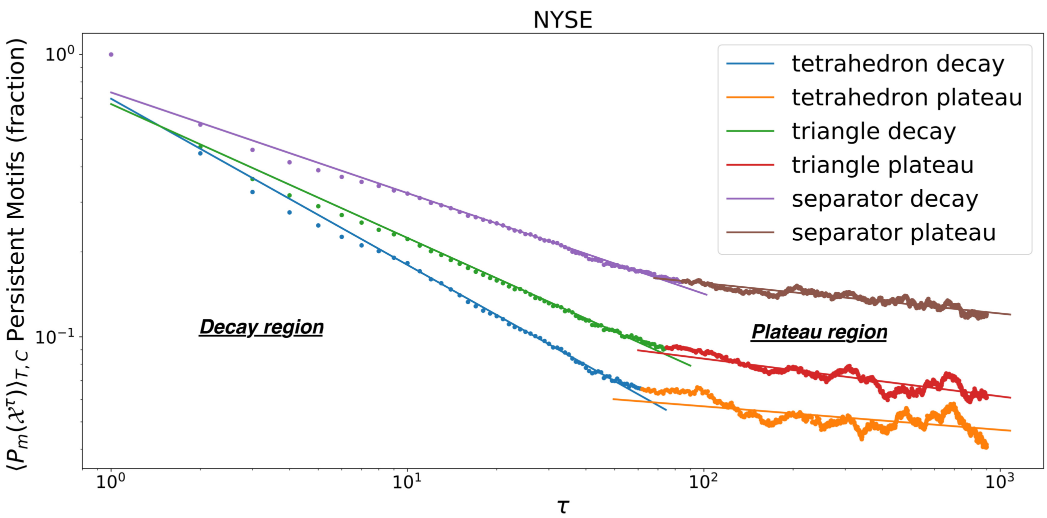


In Figure 4.1 we notice that the minimum MSE for the two linear fits is achieved at the transition point between the decay phase and the plateau. The transition point $\tau_{plat}$ can therefore be identified by minimising a standard fit measure with two phases, which strengthens the unsupervised nature of our method. The method for minimum MSE search is described Section 3.3.

## 4.3 Null models of persistence in filtered structures

The results in this section aim to demonstrate that there exists a stable sub-structure in the system which is found via the higher order relation approach (i.e. beyond pairwise simplicial relations). We observe that more efficient markets have more stable structures. We study how simplicial motifs - such as edges, triangles and tetrahedra - are persisting or decaying through the time evolution of the filtered networks computed over a rolling time window. In order to test the multivariate long memory properties of financial time series, we compare the motif persistence in networks from real data with motif persistence from a range of null models [149] - corresponding to a range of parsimonious assumptions on the underlying generative processes - for groups of time series.

Each null model preserves different aspects of the time series, allowing to validate hypotheses about the long memory of market structures by ranking persistence decays of real time series against null models.

More specifically, in this section we report results for the edge and motif (triangle) persistence for real data as well as for the null models described in Section 3.4. We compare real data with null models and TMFG filtering with quantile thresholding.

Figure 4.2 shows the decay in edge persistence for both filtering methods. We notice that the random shuffling null model lies at the bottom, as it should produce completely random structures with little residual persistence due to probabilistic combinatorics and structural filtering constraints in the TMFG. This shows that persistence is not an artifact of any of the filtering techniques used and not a mere result of return volatility of individual assets (which is preserved by return shuffling). From Fig. (4.2) we also notice that the rolling univariate Gaussian model lies just above as it does not account for structure at all and only preserves rolling means and standard deviations, this shows how persistence cannot merely be attributed to common long term trends or volatility variations. This null model carries some broad sense of structure and market direction and it shows how persistence does not merely originate from overall market trends. We then find a second cluster, of structured models, with the rolling multivariate Gaussian at the bottom. This shows how market persistence goes beyond asset means and covariances, even after spurious structures have been removed. We then find the real data, just below the stable multivariate Gaussian and with an ongoing slow, but stable decay. This shows how markets are characterised by slowly evolving long memory structures, in line with the literature reviewed in Section 1.1.

Figure 4.3 shows the decay in triangular motif persistence for both filtering methods. We notice results analogous to those in Figure 4.2 for TMFG filtered graphs. Graphs filtered through quantile thresholding instead show a high level of noise in their top cluster (where structure is present). A higher number of motifs than those of the TMFG is found, but the ranking of null models is at times inconsistent, as well as the position of the decay curve for real data. We would have expected some triangles to break when looking at motif persistence only, as well as to find that the clustering coefficient decreases in persistent graphs (as it does in TMFG graphs). The clustering coefficient for quantile persistent graphs is also found to be much higher, suggesting that the filtered structure is highly localised and clustered, while that of the TMFG is more distributed, identifying systemic groups of stocks throughout the market structure.

## 4.4 Market classification via decay exponent

We now consider how the decay exponent of TMFG graphs behaves across markets. Table (4.1) compares the decay exponents for cliques (4-cliques), triangular motifs and clique separators in the NYSE, German stock market, Italian stock market and Israeli stock market. The decay exponent $\alpha$ is obtained from the power law decay fit:

$$\langle P_m(\mathscr{X}^{\tau}) \rangle_{T,C} = \beta \cdot \tau^{\alpha} \tag{4.1}$$

We notice from the results in Table (4.1) that the NYSE, which is clearly the most developed and liquid stock market, has the lowest decay exponent (in mod-

**Figure 4.2: Decay of number of persistent edges in null models.** Decay in the number of persistent edges with $\delta_\tau$ for the time series null models of market returns and real data for the NYSE. We notice how for both TMFG filtering and quantile thresholding the real data lies between the rolling multivariate Gaussian ensemble and the stable multivariate Gaussian ensemble. This indicates that the real market structure does evolve slowly in time, but with persistence beyond what can be inferred from estimates of its covariance structure. Here we plot the number of persistent edges instead of persistence as a probability to show that both filtering methods start from the same edge sparsity.

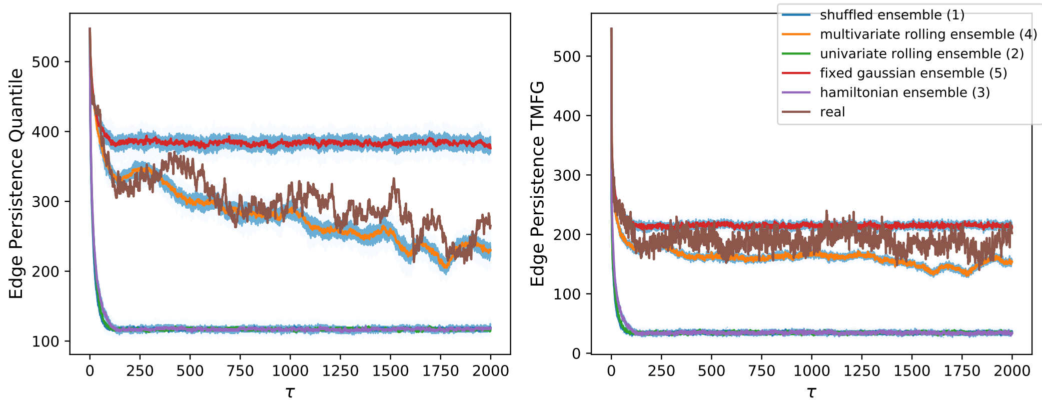


**Table 4.1: Simplicial persistence decay exponents by country.** This table reports exponents for the decay power law regime computed with MSE. The analysis refers to 100 randomly selected stocks amongst the 500 most capitalised, over time intervals $\tau = [0, 900)$ and $t = [0, ..., 200)$ different initial temporal network layers. For all motif analyses in this work, triangles and separators constitute non-overlapping sets, as these represent theoretically and taxonomically different structures and decay characteristics.

| Market | Clique | Triangular Motif | Clique Separator |
|---|---|---|---|
| NYSE | -0.392 | -0.493 | -0.245 |
| Germany | -0.792 | -0.598 | -0.381 |
| Italy | -0.785 | -0.811 | -0.174* |
| Israel | -1.024 | -0.866 | -0.728 |

* Result compromised by regimes not well identified for motif decay in large systems ($\approx$ 100 stocks).

ulus, which corresponds to the slowest decay) for both cliques and triangles. This indicates that the estimation of its correlation structure from shorter time windows is less noisy and hence more consistent in time. Germany and Italy have similar values for clique exponents, with Germany seemingly more stable in terms of tri-

**Figure 4.3: Decay of number of persistent motifs (triangle) in null models.** Decay in the number of persistent motifs with $\delta_\tau$ for the time series null models of market returns and real data for the NYSE. We notice how for TMFG filtering the real data still lies between the rolling multivariate Gaussian ensemble and the stable multivariate Gaussian ensemble (as in Figure 4.5a). We instead notice that the decay ordering is noisier for quantile thresholding, showing how the method's focus on individual connections affects it generalisation to motifs. This is despite the higher number of motifs in the quantile thresholding graph. Here we plot the number of persistent motifs instead of persistence as a probability to show that in spite of both filtering methods starting from the same edge sparsity (Fig. 4.5a) the number of initial cliques and persistent cliques is very different across methods. This is likely due to the highly clustered structure of thresholding-based networks.

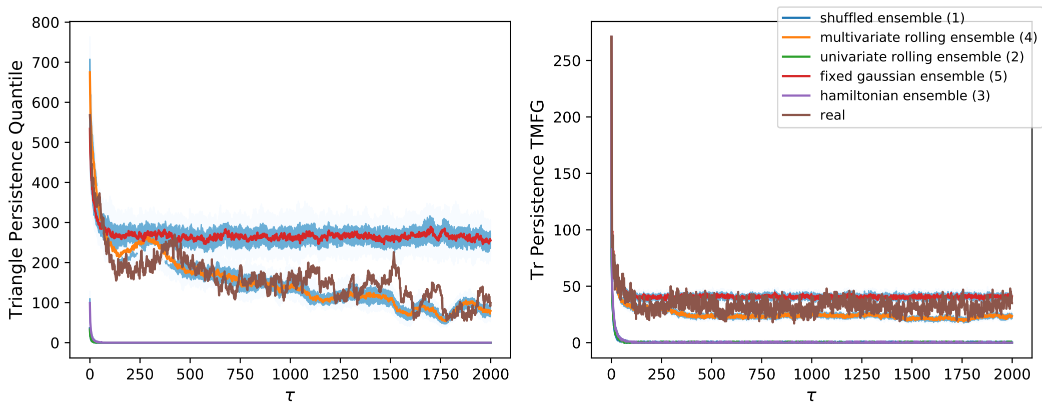


angular motifs. Israel, a younger and less liquid stock market, follows with a faster decay in both tetrahedral cliques and triangular motifs. The ordering of these markets is not clearly identifiable in clique separators as noise in the data does not allow for the two decay regimes to be correctly identified in all markets (in this case for Italy). Separators have a distinct role and meaning in the graph's taxonomy and further work should allow for a more thorough analysis of those.

We observe promising results for a monotically increasing relation between the decay exponent and the average daily volume of the market. The solidity of this result shall be investigated in future works.

We now show that simplicial persistence is indeed meaningful beyond edge persistence, which strengthens the novelty of our contribution. In Table (4.1) the decay exponent is not adjusted by the probability that all edges in the clique must be present in the temporal layer for the clique to exist. We show in Table (4.2) that, when adjusted by the probability of all its edges existing simultaneously, triangular motifs have a slower decay than individual edges. The results in Table (4.2) are obtained from a set of randomly selected stocks different to those used for Table (4.1). This adds further confidence in the results and their generality.

We stress that Table (4.2) rejects the hypothesis that motifs are formed by edges in the network whose existence is not mutually dependent (Equation 3.5). The hypothesis is rejected by the consistently lower decay exponent (in modulus) for adjusted persistence of triangular motifs. We can then conclude that motifs are more stable structures across temporal layers of the network, with significant interdependencies in their edges' existence.

**Table 4.2: Persistence of higher order motifs beyond component edges.** This table reports exponents for the decay power law regime computed with MSE. The analysis refers to 100 randomly selected stocks amongst the 500 most capitalised, over time intervals $\tau = [0, 900[$ and $t = [0, \ldots, 200[$ different initial temporal network layers.

| Market | Edge | Triangular Motif | Triangular Motif** |
|---|---|---|---|
| NYSE | -0.164 | -0.398 | -0.133 |
| Germany | -0.265 | -0.471 | -0.157 |
| Italy | -0.144* | -0.458 | -0.153 |
| Israel | -0.397 | -0.830 | -0.277 |

* Result compromised by regimes not well identified for edge decay in large systems ($\approx$ 100 stocks)
** Motif exponent adjusted by the probability of simultaneous edge persistence in the motif).

## 4.5 Sector analysis of most persistent motifs

Figure (4.4) provides a visualisation of the network components formed by the ten most persistent triangles in the NYSE. We observe that all strongly persistent triangles have elements which belong to the same industry sector. Table 4.3 shows this for the same ten triangles displayed in Figure (4.4). We notice that stock prices in the sectors in Table (4.3) are mostly driven by sector-wide fundamentals, which justify the persistent structure in the long term. Other motifs are constituted by

**Figure 4.4: Persistent NYSE motifs visualised.** Network representation of the ten most persistent triangular motifs in the TMFG layers for the 100 most capitalised stocks of the NYSE..

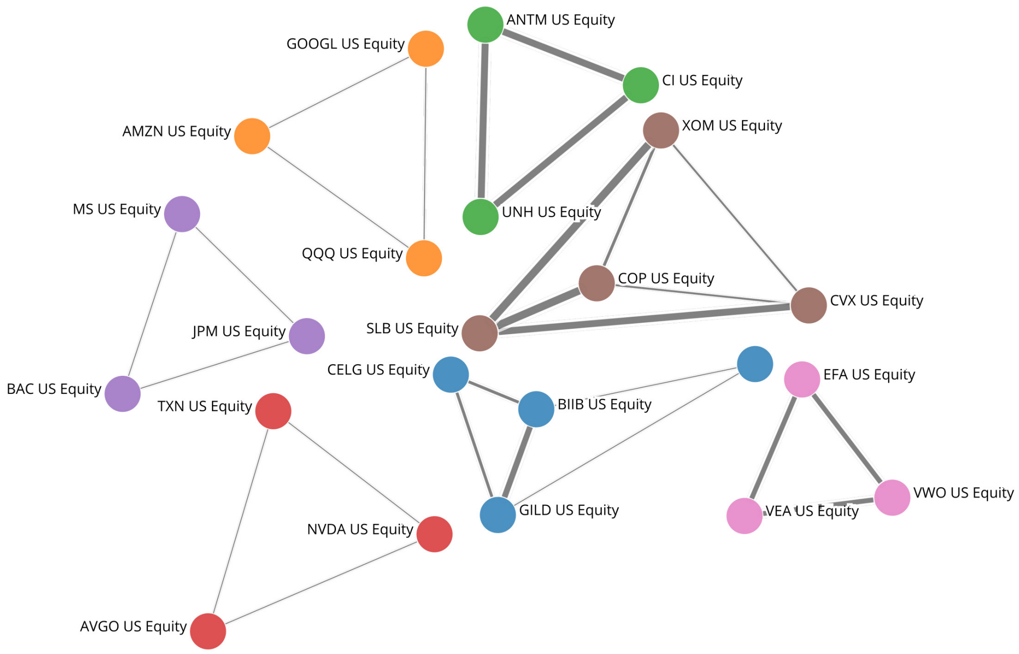


Exchange Traded Funds (ETFs) and their main holdings [2].

We also investigate whether motif persistence and motif structures can be easily retrieved from the original correlation matrix. The purpose of this is to check that our TMFG filtering method is not redundant and trivially replaceable. To test this, we consider the ten most present persistent triangles across the plateau region and check their overlap with the ten most correlated triplets in each unfiltered correlation matrix. We find that no more than one triangle lies in the intersection between the two sets, in each temporal layer. We also check the correlation between motif persistence and the average sum or product (results are equivalent for our purpose) of its individual edges' correlation for all unfiltered correlation layers. We observed through the Pearson and Kendall correlation values that the two measures are only

[2]The reason for the existence of these motifs is intuitive and does not affect our analysis, as ETF-related motifs are unlikely to be present in the network formed by a random selection of stocks or by stocks in a portfolio. These motifs are present here as we focus on the 100 most capitalised stocks in the NYSE, which include ETFs.

**Table 4.3: Motif components and Financial Times sector affiliation for the ten most persistent motifs in the NYSE's 100 most capitalised stocks.**

| Node 1 | Node 2 | Node 3 | FT Sector |
|---|---|---|---|
| Biogen Inc | Gilead Sciences Inc | Celgene Corp | Biopharmaceutical |
| UnitedHealth Group Inc | Cigna Corp | Anthem Inc | Health Care |
| Biogen Inc | Gilead Sciences Inc | Amgen Inc | Biopharma/tech |
| Bank of America Corp | JPMorgan Chase & Co | Morgan Stanley | Financials-Banks |
| Vanguard FTSE ETF** | MSCI EAFE ETF | Vanguard FTSE ETF*** | Index ETFs |
| Invesco QQQ Trust* | Amazon.com Inc | Alphabet Inc | Tech |
| ConocoPhillips | Schlumberger NV | Exxon Mobil Corp | Oil & Gas |
| NVIDIA Corp | Texas Instruments Inc | Broadcom Inc | Tech Hardware |
| Chevron Corp | Schlumberger NV | Exxon Mobil Corp | Oil & Gas |
| Chevron Corp | ConocoPhillips | Schlumberger NV | Oil & Gas |

* ETF on NASDAQ - Top Holdings include Amazon, Facebook, Apple, Alphabet
** Vanguard FTSE Developed Markets Index Fund ETF Shares
*** Vanguard FTSE Emerging Markets Index Fund ETF Shares

loosely related, as correlation explained no more than 20% of the variance in the set of variables with large persistence.

## 4.6 The volatility of motif portfolios and inverse persistence weighting

We check that a subgraph (a portfolio, i.e. group of stocks) formed by the 10 most persistent motifs in each market presents a significantly higher out of sample standard deviation $\sigma_{sub}$ (volatility - standard deviation of the mean log-return of the subgraph elements' prices) due to its stable correlations. To do this, we consider the $\sigma_{sub}$ of the motif subgraph and a distribution of $\sigma_{sub}$ for $10^5$ randomly selected subgraphs with the same number of nodes (stocks).

As expected, we observe in Figure 4.5 that the persistent motif subgraph is characterised by a $\sigma_{sub}$ over two standard deviations above the mean of the distribution as well as above its 75% quantile throughout the considered markets. We should highlight that the $\sigma_{sub}$ of subgraphs is evaluated out of sample with respect to the period the persistence was calculated on, showing that this method is not only observational, but also predictive.

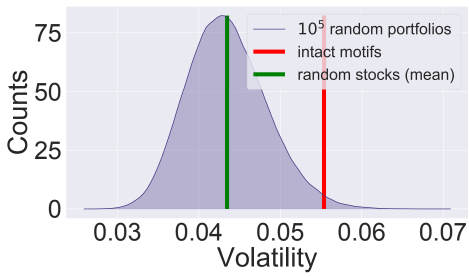


**(a)** NYSE

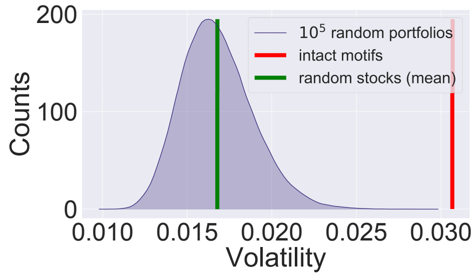


**(b)** Germany

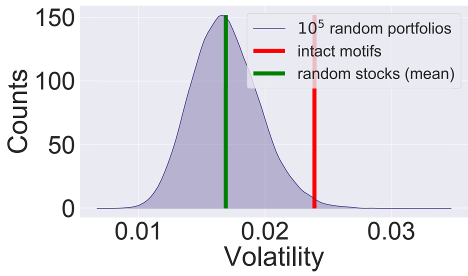


**(c)** Italy

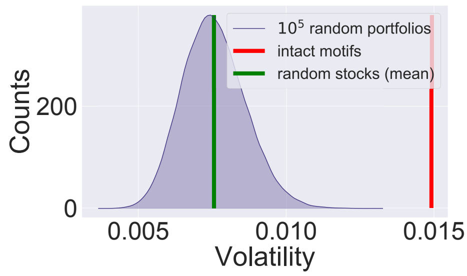


**(d)** Israel

**Figure 4.5: Persist motifs vs random portfolios volatility.** Portfolio volatility distribution for the 100 most capitalised stocks in the NYSE (a), German stock market (b), Italian stock market (c) and Israeli stock market (d). The reference portfolio (red bar) contains all stocks in the 10 most persistent triangles and distribution portfolios are formed from a random selection of stocks (mean distribution volatility represented by the green bar).

In line with the approach of this thesis, we now propose an applied proof of concept based on our findings. We wish to form a measure directly applicable to portfolio optimisation and compare it with the widespread inverse volatility ($\propto 1/\sigma$) weighting. Our findings from Figure 4.5 show that highly persistent structures have higher out of sample volatility, due to idiosyncratic risk. As we aim for low volatility, we consider the measure introduced in Equation 3.6. In line with inverse volatility weighting we weight each stock $i$ (vertex $v_i$) in our portfolio as $\propto P_m(v_i)$. In Figure 4.6 we present out of sample results where we observe a reduction in volatility throughout markets, with distributions more than one standard deviation apart. We discuss these results and their implications more in depth later, in Section 7.2.

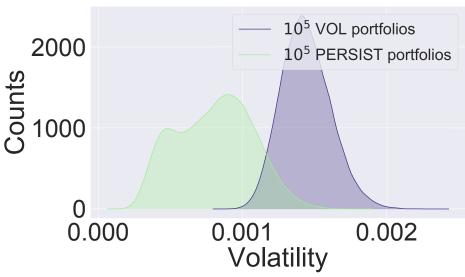


**(a)** NYSE

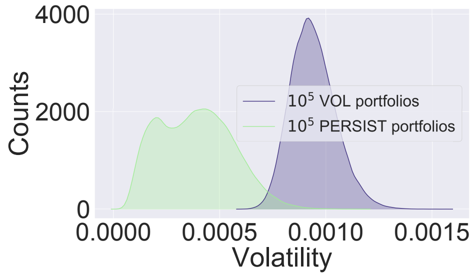


**(b)** Germany

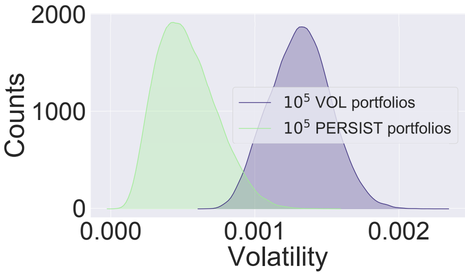


**(c)** Italy

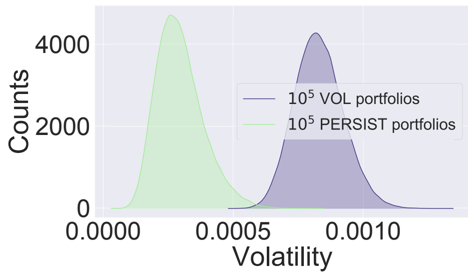


**(d)** Israel

**Figure 4.6: Volatility of persistence-weighted vs variance-weighted portfolios.** Portfolio volatility distribution for the 100 most capitalised stocks in the NYSE (a), German stock market (b), Italian stock market (c) and Israeli stock market (d). The VOL portfolios are formed by weighting a random selection of assets by $1/\sigma^2$ while the PERSIST portfolios are formed weighting assets by $1/P_m(v_i)$.

# Chapter 5

# Systemic risk in HFT: from single to multi-asset dynamics

We provide a general introduction to the ecosystem of high frequency markets in Sections 1.4, 1.5. In Section 2.3 we discussed the systemic risk aspects closer to the findings which we present in this chapter. Section 2.3 introduces the dynamics of flash crashes and the idea of positive feedback loops. This provides the context to interpret the numerical observations in this chapter.

## 5.1 Data

We consider granular order flow data for a universe of 300 liquid stocks from the NASDAQ exchange between 3/1/2017 and 25/9/2020. Section 5.2 focuses on a single asset (Apple - AAPL in our case), while co-crash results are based upon the full universe of stocks. High frequency price data is obtained from LOBSTER [150] and sampled to obtain non-overlapping one minute returns. This frequency was also adopted in [103] and other works in the literature for the detection of price jumps as it is understood that below this limit microstructural noise becomes relevant and can impact the validity of the method. The data from LOBSTER provides detailed order flow data which indicates the type (limit order, cancellation, execution of a limit order), the side (buy or sell) and volume of each event in the order flow. This allows to isolate trading volume and other quantities from the rest of the order flow and we discuss the importance of this in Chapter 3.

Return series for the jump detection method described in Section 3.7 were obtained from 1 minute non-overlapping returns of the original time series. The 1 minute intervals used for the return analysis are then used to bucket trading volumes by time interval. The minute-wise volume buckets then contain the sum of traded volumes in each non-overlapping 1 minute interval, with its corresponding jump or no jump label from the method in Section 3.7.

## 5.2 The unboundedness of crash dynamics

The relation between order flow and price returns is at the core of most studies in market microstructure, as per our more extended discussion in Section 3.6. In this section we begin our study of flash crashes in high frequency markets by comparing the order flow dynamics in "normal" and flash crash periods. The plots in Figure 5.1 show the distribution of aggregate trade volumes by minute for time intervals when a price jump was detected (Figure 5.1a) and those when it was not (Figure 5.1b).

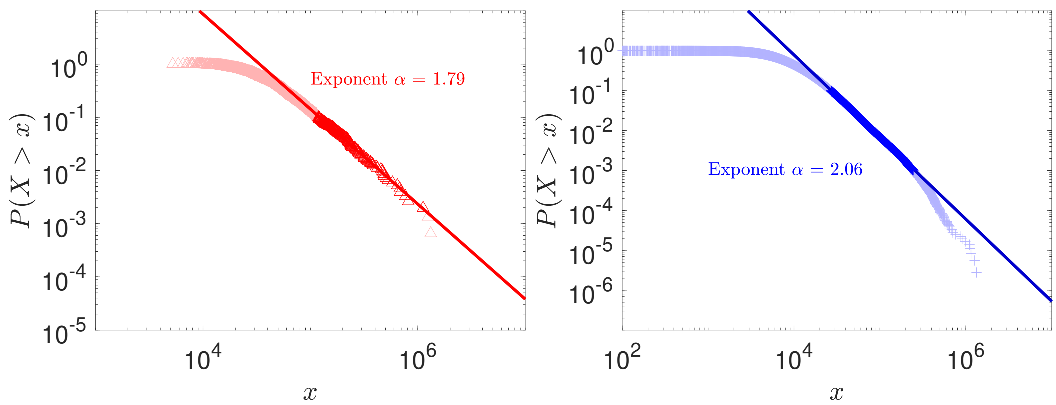


**(a)** Jump **(b)** No jump

**Figure 5.1: Trading volume distributions of jump and non-jump intervals.** Log-log plot of the volume distributions for 1 minute intervals with power. The plots report the empirical distribution and power law tail fit for intervals with jumps (Figure **(a)**) and no jumps (Figure **(b)**). In Figure **(a)** we find a decay exponent $< 2$ which indicates an unbounded critical distribution, while in Figure **(b)** we find a decay exponent $> 3$ which indicates a bounded distribution with well defined mean and variance.

We indeed observe the distribution of jump volumes to be shifted in mean and median and characterised by higher volumes, but this was expected from the relation between traded volume, price impact and volatility [117]. What is interesting is the fact that the jump distribution is not only shifted with respect to the general distribution, but also characterised by a heavier tail. At this point we highlight that our jump detection method described in Section 3.7 is only aware of price changes, thereby removing any bias on explicitly selecting anomalous volumes, apart from their potentially anomalous corresponding price jump.

**Table 5.1: Decay exponents of jump and non-jump intervals for liquid stocks.** We apply the method from Clauset et al.[127] to fit jump and non-jump volume distributions analogously to Figure 5.1 and verify the robustness of our results. We select stocks from the ones with the highest average number of book updates per day (a proxy for liquidity) within the universe of stocks and period considered. We find again a decay exponent $< 2$ for jump intervals which indicates an unbounded critical distribution, while we find a decay exponent $> 3$ for no-jump intervals which indicates a bounded distribution with well defined mean and variance.

| Stock | Jump Exponent | Non-Jump Exponent |
|---|---|---|
| MSFT | 1.90 | 2.02 |
| AAPL | 1.48 | 2.04 |
| AMD | 1.79 | 2.48 |
| ADBE | 1.88 | 2.14 |
| AVGO | 1.58 | 2.14 |
| CHTR | 1.66 | 2.25 |
| BIIB | 1.30 | 2.19 |
| CSCO | 1.52 | 2.01 |
| MS | 1.57 | 2.63 |
| AMAT | 1.76 | 2.10 |
| GOOG | 1.86 | 2.18 |
| T | 1.53 | 2.80 |
| SBUX | 1.48 | 2.02 |

We apply the methodology from Section 3.7.1 to fit the power law tails of the jump and non-jump trade volume distribution CCDFs (Complementary Cumulative Distribution Function). The decay exponents are reported in Figure 5.1. We notice how the decay exponent of jump volumes is $< 2$ in modulus, while that of non-jump volumes is $> 2$ in modulus. Decay exponents in power law distributions are of great importance [1] as they indicate which moments of the distribution are defined. Indeed a decay exponent $< 2$ indicates that the second moment of the distribution (variance) is not defined, i.e. the distribution is unbounded, while one $> 2$

characterises a bounded distribution with well-defined variance. This can also be viewed as the fact that the sample variance diverges with sample size for exponent $< 2$, while it converges to the "true" variance in the limit, for exponent $> 2$.

As the exponent values are crucial for our discussion, we wish to clarify some robustness properties of the results in relation to our description of the fitting methodology in Section 3.7.1.

We choose the percentile interval $[90\%, 99.9\%]$ to fit the power law tail, but our main results are robust to choosing a starting level of say 95% and not excluding values beyond 99.9%. The latter in particular gives a non-jump exponent $> 3$ ($\alpha = 3.21$) and a virtually unchanged jump exponent. This shows how the exponent in Figure 5.1b is a lower bound and our qualitative conclusions are robust.

In order to further show the robustness of our results we apply the fit method from Clauset et al.[127] as described in Section 3.7.1 and report consistent results with the ones from Figure 5.1 for AAPL for multiple liquid stocks in Table 5.1.

Another potential critique could arise from the fact that our fit in Figure 5.1b excludes jump intervals and this could be the cause of the observed bounded variance as extreme values are removed. We have indeed investigated the more comprehensive distribution which includes both jump and non-jump intervals and found that this has little statistical effect on the fit and resulting exponent ($\alpha = 2.02$). This also highlights that our filtering of jump intervals alone shows the unboundedness in their volume distribution and further supports our analysis and the validity of its conclusions.

A very important consequence of the above observations of the different decay exponent between jump and non-jump distributions, rather than just a shift, is the fact that indeed the two sets of intervals have different underlying distributions, generating processes and resulting characteristics. We report that the idea that the order flow distribution is not fully stationary and that it might influence the impact and subsequent order flow is already hinted by Bouchaud and co-authors in [151] on the basis of previous studies on impact [152, 153, 118].

Our finding adds the essential point to the microstructural literature than indeed

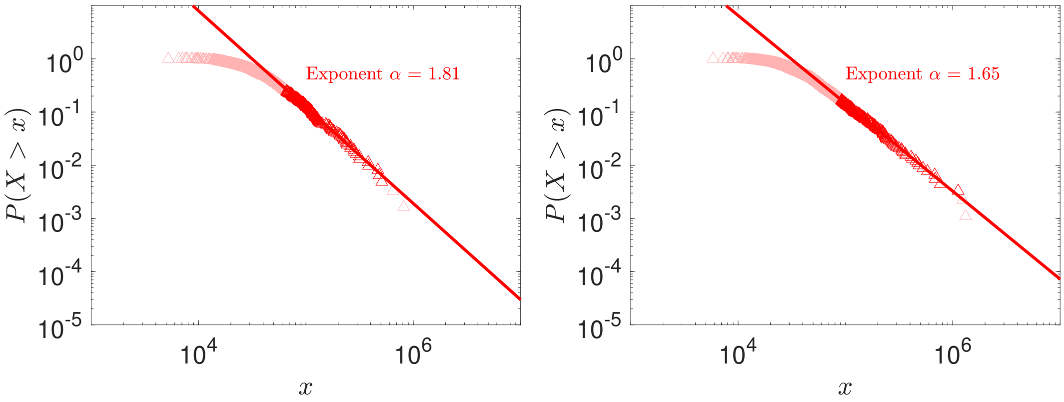


**(a)** Positive return **(b)** Negative return

**Figure 5.2: Trading volume distributions of jump intervals by return sign.** Log-log plot of the volume distributions for 1 minute intervals. The plots report the empirical distribution and power law tail fit for intervals with jumps and positive (Figure **a**) and negative (Figure **b**) corresponding price returns. In both, we find a decay exponent $< 2$ which indicates unbounded critical distributions.

volume distributions are heavy-tailed, but that under "normal" market conditions the distribution is bounded and can be well approximated by risk models which assume Gaussian properties. Gaussian assumptions in risk modeling mostly rely on the assumption of bounded variance. On the other hand, black swan events in volume bursts belong to a different distribution and process altogether. A similar idea was suggested in [105] where the authors investigate flash crashes between 2006 and 2011 to show a system-wide phase transition around $\sim 500ms$ to an all-electronic trading market characterised by black swan events. From the perspective of self-organised criticality, anomalous price jumps present an underlying volume process characterised by black swan events of onbounded magnitude where events of arbitrarily large sizes are characterised by non-zero probabilities.

The distinction is crucial in this context as black swan events characterise unbounded distributions alone and we observe how these extreme and anomalous price events are characterised by black swan events in the underlying market activity. This indeed supports the thesis from the literature described in Section 2.3 that crashes are characterised by a very large and concentrated trade activity with unbounded variance, different from the trade distribution of regular order flow.

A breakdown of the results from Figure 5.1a is presented in Figure 5.2, which analyses the volume distributions of up and down jumps separately. It is broadly accepted in the literature that the return distribution is asymmetric with a heavier tail for negative returns [66]. Based on the relation between volume imbalance and price returns [154] one could suggest that the detected jumps correspond to negative returns which present a heavier tailed volume distribution. We show that our jump detection method accounts for this asymmetry as we detect a similar number of positive and negative jumps. Further, exponents for the volume distributions of positive and negative return intervals in Figure 5.2 are both consistent with our previous findings from Figure 5.1. However, we highlight that, as expected, negative return intervals are characterised by a heavier distribution tail as events are more extreme to the downside [155, 156, 157].

# Chapter 6

# Co-crashes and their structure

We provide a general introduction to the ecosystem of high frequency markets in Section 1.4, 1.5. In Section 2.4 we discussed the systemic risk aspects closer to the findings which we present in this chapter. Section 2.4 recalls the ideas behind complex and connected systems from Section 1.2 to introduce the emerging literature on synchronisation of black swan events in high frequency markets. This aligns with the more "traditional" definition of systemic risk where extreme events spread across multiple assets due to correlated actions by market players. This introduces our basic results in Section 6.2 on the distribution of co-crash sizes. Further, in Section 6.3 we show a phase transition of the crash frequency distribution at higher crash sizes and provide statistical validation for it. We conclude the chapter with early results from future work underway on the role of liquid stocks in co-crash structures and the more general core-periphery structure of the co-crash network.

## 6.1 Data

We consider granular order flow data for a universe of 300 liquid stocks from the NASDAQ exchange between 3/1/2017 and 25/9/2020. Section 5.2 focuses on a single asset (Apple - AAPL in our case), while co-crash results are based upon the full universe of stocks. High frequency price data is obtained from LOBSTER [150] and sampled to obtain non-overlapping one minute returns. This frequency was also adopted in [103] and other works in the literature for the detection of price jumps as it is understood that below this limit microstructural noise becomes relevant and can impact the validity of the method. The data from LOBSTER provides detailed order flow data which indicates the type (limit order, cancellation, execution of a limit order), the side (buy or sell) and volume of each event in the order flow. This allows to isolate trading volume and other quantities from the rest of the order flow and we discuss the importance of this in Chapter 3.

Return series for the jump detection method described in Section 3.7 were obtained from 1 minute non-overlapping returns of the original time series. The 1 minute intervals used for the return analysis are then used to bucket trading volumes by time interval. The minute-wise volume buckets then contain the sum of traded volumes in each non-overlapping 1 minute interval, with its corresponding jump or no jump label from the method in Section 3.7.

## 6.2 The distribution of co-crashes

The work by Lillo et al.[103] discussed in Section 2.4 highlights how co-crashes have become larger and more frequent in recent years. This shows the growing importance of such events and motivates us to investigate their structure and properties further. The plot in Figure 6.1a shows the frequency distribution $f_m$ of the number of stocks involved in each flash crash. Figure 6.1b plots the cumulative frequency $f(M \geq m)$. It is evident from both figures that the distribution is heavy-tailed and there is a change in the slope around $m \approx 5$ and a finite size effect at $\approx 10^2$, i.e. when the crash involves most of the system (system size $3 \cdot 10^2$) [1]. This kind of distribution was already reported in [103], where the authors investigated and modeled flash crash sizes and frequencies as a unique Hawkes process.

The authors there suggest that each security's crash dynamics should be modeled as a self-excitation process, but they point out that this would involve tuning a large number of parameters on very noisy data. They therefore decide to model the collective self-excitation process of securities as the frequency of crashes (or co-crashes) and their size. Hence, all crash sizes are treated as instances of a multi-asset Hawkes process in [103], with no distinction between the assets involved in each crash or their structure.

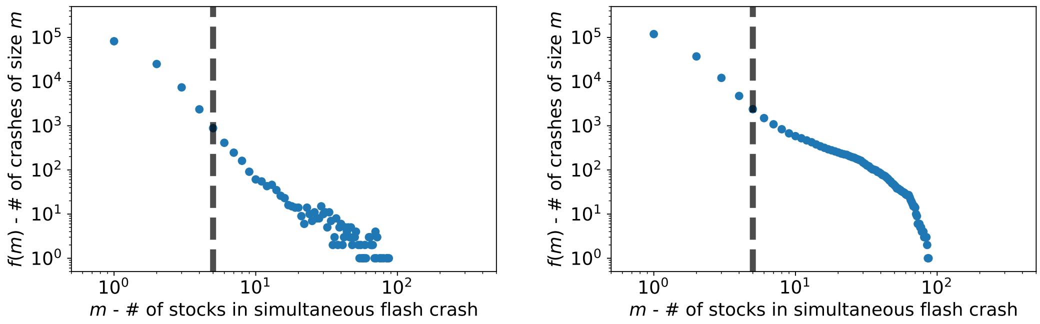


**(a)** $f(m)$ **(b)** $f(M \geq m)$

**Figure 6.1: Heterogeneous crash size distributions.** Log-log plot of the flash crash size distribution. We observe that sizes lesser than 4 diverge from the power law, with lower than expected frequency. This suggests that crashes of this size and onwards do not belong to the same self-organised process, but that this is rather a heterogeneous distribution.

## 6.3 A phase transition to systemic crashes

For the multi-asset flash crash investigation in this work we take a more granular approach and move to investigate the structure of co-crashes and the individual susceptibility of each stock.

To further investigate the difference between small and large crash sizes we report in Figure 6.2a the Spearman correlation between the ranks of crash frequency for all stocks in the dataset from Section 6.1. Specifically, each line reports the correlation between the rank of the companies in the initial crash size $m$ with all other crashes of higher size $m+\tau$. We indeed observe how crashes of smaller sizes ($m < 5$) have a substantially different composition to crashes of larger sizes. We instead observe that for sizes $m > 6$ a steady state is reached with a large component of the population represented across all crashes with similar ranks in terms of crash frequency. These steady states for $m > 5$ are significantly higher than the ones of smaller sizes, as the structure no longer evolves significantly between higher size crashes. The plot in Figure 6.2b provides a clearer visualisation of this. We highlight that already at size 5 the correlation transitions directly to the steady state, although a lower one with respect to the ones for crash size 6 and above.

To validate the visual results from Figure 6.2 we apply the null model of correlation significance between crash frequency distributions from Section 3.9.

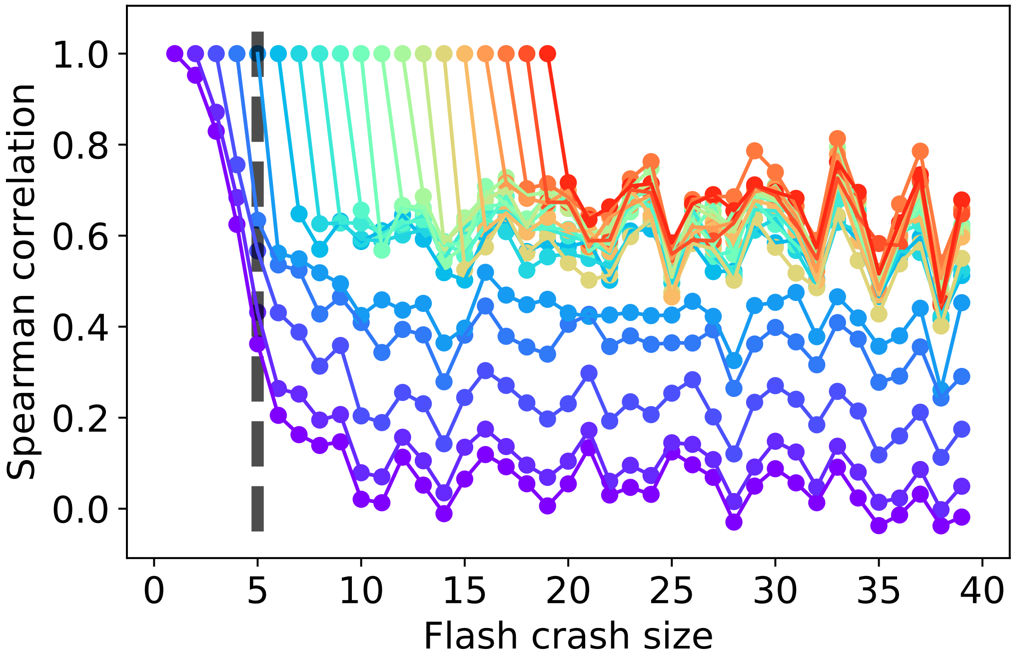


**(a)** Spearman correlation between all consecutive crash sizes.

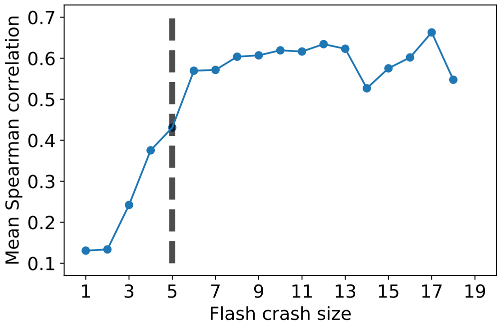


**(b)** Spearman steady state correlation mean in $[m+2, m+20]$, $\forall m$.

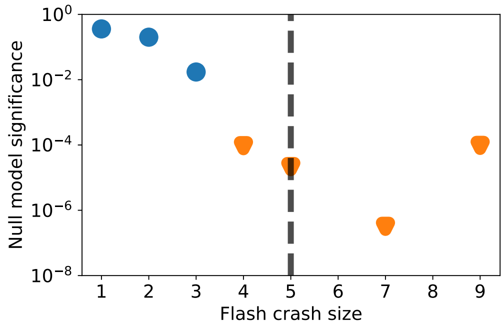


**(c)** Quantile of $Spearman(S_{m+\tau}, S_m)$ in $D_m$, as per Section 3.9. With $S_{m+\tau}$ the steady state distribution $\sim [m+2, m+10]$.

**Figure 6.2: Crash component rank correlation phase transition.** Evidence that there is a transition around $m = 5$ with crashes involving a small number of companies ($m < 5$) being substantially differently populated with respect to crashes involving a larger number of companies ($m > 5$). The plot in Figure **a** reports the Spearman correlations of ranks in frequency between each starting crash size and higher crash sizes. The color map is based upon the starting co-crash size and allows to visualise how higher co-crash sizes belong to a different distribution. The plot in Figure **b** looks at the average correlation in the range $[m+2, m+20]$ for each value of $m$ from Figure *a*, which offers better visual intuition. The plot in Figure **c** reports the steady state statistical significance of the base crash size's frequency distribution indicating a phase transition.

Figure 6.2c shows the correlation significance between the starting point $m$ on the horizontal axis and its steady state distribution $\sim [m+2, m+10]$. We observe the first significant value at 1% around $m = 4$ which confirms the intuition from Figures 6.1a, 6.1b that crash sizes up to $\approx 4$ belong to a different process than larger crashes. Indeed smaller crashes are dominated by less stable stocks and larger ones by very liquid stocks with high market capitalisation. This suggests that more influential and systemic stocks are involved in larger crashes and perhaps even trigger those. A reason for why this is not the case in small crashes can be that these stocks are systemic enough to mostly cause crashes of larger size. These are then even more relevant for systemic risk.

This is therefore further evidence of the occurrence of a transition in the process between smaller and larger crashes. The slow decay of smaller crash sizes indicates how these belong to similar distributions of non-systemic events, but as the crash size grows the steady state gets closer to the large crash level. This suggests that larger crashes have some systemic characteristics.

If we take a closer look at the top ranked stocks at each size we observe that smaller crash sizes are dominated by very volatile and illiquid stocks which are subject to large jumps perhaps due to the lack of a smooth price process in their trading. We would expect this though to make them susceptible to larger systemic events as well and hence stably ranked. Yet, we observe very low to null rank correlation between individual (and small) crash frequencies and the large crash size steady state. It seems as if not only these crashes are non-indicative, but also as indicated by the phase transition in Figure 6.2c they belong to an unrelated ranking and distribution. We highlight that we considered rankings and ranking correlation in order to avoid any sensitivity to large values or outliers at smaller or larger frequencies.

Large crash sizes involve stocks such as Microsoft (MSFT) and Apple (AAPL) as consistently high ranked. We highlight that these stocks are highly liquid and characterised by a stable price process with very few price jumps. Indeed the few times they get involved into jumps they are often part of larger simultaneous crashes, which involve more stable and systemic stocks. Further, when analysing the co-crash relations between pairs of stocks we observed a heavy-tailed distribution of centrality for these large systemic stocks which suggests a community and core-periphery like structure of the contagion network of co-crashes [158, 159, 160, 161, 162].

## 6.4 Early insights into the structure of co-crash networks

The observations from Section 6.3 above prompted us to conduct further analyses on the relation between stock liquidity (where average daily Dollar Traded Volume (DTV) is used as a proxy) and crash frequency at different crash sizes.

To validate visually and numerically our observation that highly traded stocks are more present in large crashes we present the plots in Figure 6.3. The plot in Figure 6.3b shows the average daily traded volume of a stock per crash size, weighted by its crash frequency, as per the definition in Section 3.10. This is plotted against the crash size to show a clearly increasing trend in crash-weighted traded volume with crash size. This shows how larger crashes see stocks with higher traded volumes more frequently involved.

This could though be the consequence of a subset of crashes which involve highly traded stocks. We therefore test this with results in Figure 6.3a which show how not only the average crashing stock is more "liquid" in larger crashes, but also that the fraction of crashes which involve at least one of the top 20 stocks by traded volume in our universe increases with crash size. As suggested in Section 3.9 we need to control for the increase in crashes with one of the 20 stocks due to larger crashes sampling more stocks. We use $10^4$ realisations of the null model from Section 3.9 and show qualitatively at least that the increase goes beyond this expected bias and is indeed a feature of the system.

In line with this, we test how the traded volume of each stock correlates with its crash frequency, for each crash size. We report results for the Spearman correlation coefficient in Figure 6.3c, where dots are used for correlations significant to the 5% confidence level and crosses otherwise. We see that co-crashes of size 1 and 2

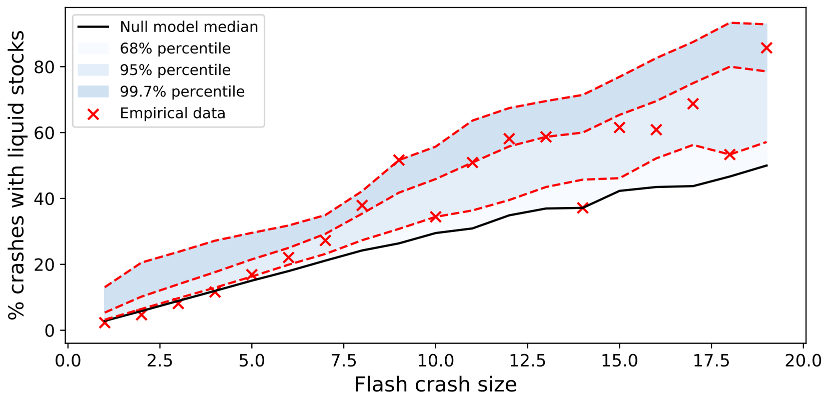


**(a)** Positive relation between fraction of crashes involving liquid stocks and crash size $m$.

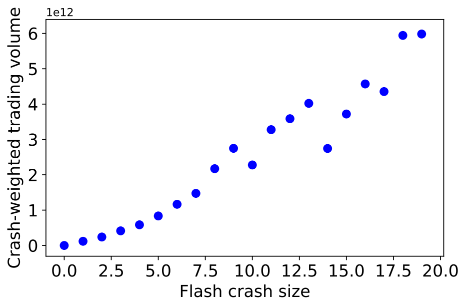


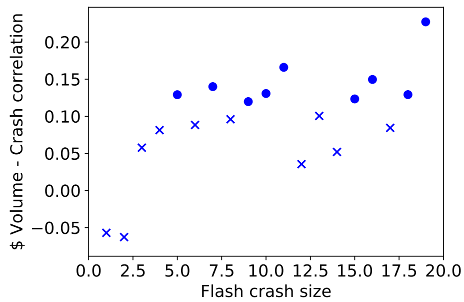


**(b)** Positive relation between crash-weighted average daily Dollar Traded Volume and crash size $m$.

**(c)** Spearman correlation between traded volume and crash frequency across crash sizes $m$.

**Figure 6.3: Relation of traded volume to crash size and liquid stock peripheries.** The figures above show evidence of a relationship between the traded volume of stocks and their involvement in crashes of different sizes. Figure **(a)** shows the general positive relation between crash size and involvement of highly traded stocks. Figures **(b), (c)** show how the relationship exists not only on average, but also how "liquid" stocks are more involved throughout crashes at higher crash sizes.

seem to have an inverse or no relation between volume traded and crash size. At our previously identified phase transition point $m \approx 5$ we see the first significant positive correlation between volume traded and crash frequency which stays somewhat stable or is slightly increasing with crash size.

This last result is less clear than the previous one, but still shows a positive correlation between volume traded (a proxy for liquidity) and crash frequency at crash sizes $m > 5$.

The presence of liquid stocks in most large crashes observed in Figure 6.3a

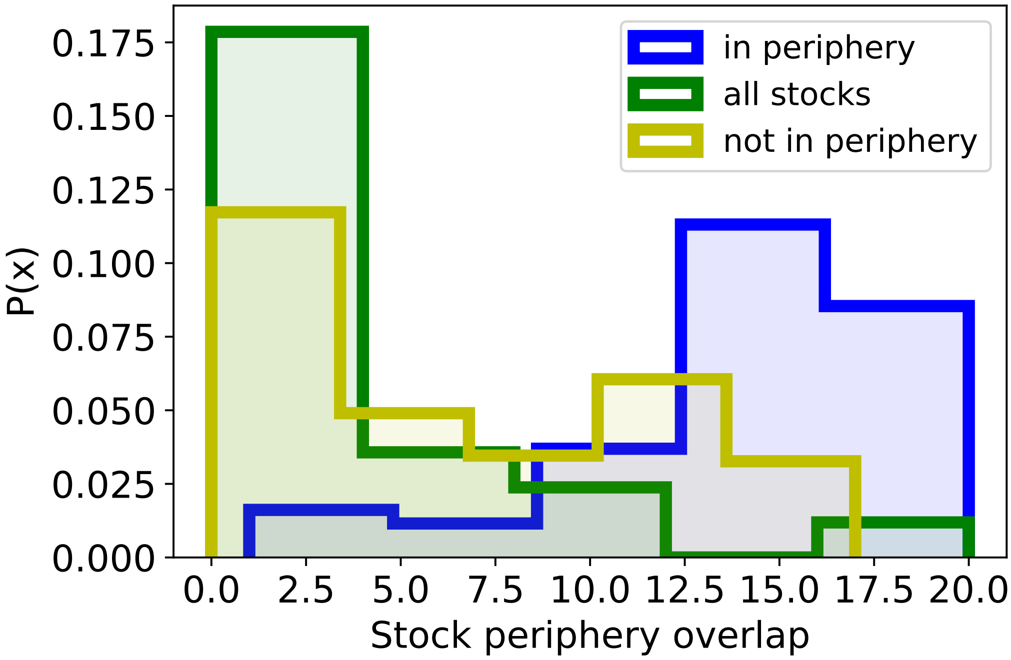


**Figure 6.4: The community structure of co-crash networks.** Overlap of top 20 periphery stocks for each pair of top (20) liquid stocks. The plot shows how, as long as the two core liquid stocks are not part of the same community, liquid stocks represent different cores with non-overlapping peripheries.

prompts questions around the periphery structure of the different liquid stocks and implications for systemic risk. Further work in this direction is already underway with promising preliminary results which we present in Figure 6.4. For the plot we select the top 20 liquid stocks by average daily Dollar Traded Volume. We then define the periphery of each as its top 20 most co-crashing stocks across crash sizes. The plot then represents the periphery overlap for each pair of liquid stocks when those belong to each other's periphery (blue) or not (yellow). We also present in green the overlap of each stock's periphery with the overall top 20 stocks by crash frequency. These results provide significant indications of a core-periphery network with multiple cores represented by one or more liquid stocks, with little overlap between different peripheries. This is strengthened by the fact that the overlap between different peripheries is similar to that with the 20 most crashing stocks (green) which can be viewed as a null model of overlap due to mere crash frequency and not structural closeness. Liquid stocks in the same community, instead, present much higher periphery overlap.

# Chapter 7

# Discussion

## 7.1 Systemic risk is real

It is now evident to the reader that the focus of this work is on systemic risk in financial markets across frequencies, market participants and assets. We should not think though that this is an issue in isolation and only relates to financial markets and risk. We have already reviewed in Section 1.2 how percolation and cascades in connected systems have been studied in telecommunication networks, electric grids, the internet, social networks and many other systems. We then saw that already in the traditional financial and economic literature there has been strong interest in systemic risk, in particular following financial and economic crises. Our discussion in Section 1.3 then moved to review applications of the concepts from Physics and Network Theory of Section 1.2 to financial markets. The growing Econophysics literature reviewed in Section 1.3 focuses complex connected systems in financial markets, at both high and low frequencies. Our review shows how different kinds of systemic risk are present in the financial system and can be analysed, monitored and prevented. Further, our discussion of regulation later in Section 7.3, reviews the literature of agent based modeling in finance and economics and the importance of modeling the complexity of systems arising from agent interactions in a bottom up fashion.

All these works analyse real data and systems of crucial importance for our daily lives and show how systemic risk is real, relevant and worth understanding.

In line with this, we discuss the results from Section 4 to gain insights into the different types of systemic risk present and how we can learn to better deal with them.

### 7.1.1 Null models and generating market structures

The power law decay of edge and simplicial motifs persistence, reported in Figure 4.1, suggests that market structures are characterised by a slow evolution with long memory. This decay type is in contrast with an exponential decay of persistence which would imply short or no memory in the system. This observation is in line with the works by Bouchaud et al. and Lillo at al. in [119, 144, 118, 146], where power law decays in autocorrelation are identified as manifestations of long-memory processes in efficient markets. However, it extends the concept to higher order structures.

The comparison between soft persistence in correlation structures of real data and artificial data generated from univariate null models (Figures 4.2 and 4.3) shows that the persistence and memory of real structures implies multivariate structures at least. We also demonstrate how real structures present higher persistence than sampling from a rolling multivariate Gaussian, hence suggesting that pairwise covariances and moving averages do not suffice to induce the long memory present in real markets. As per the analysis on motif persistence beyond that of individual edges in Section 4.4, we suggest that higher order relations in terms of structural evolution are present. As just discussed, the ordering of null models in Figures 4.2, 4.3 supports the validity of the persistence measure.

### 7.1.2 Advantages of topological filtering: the TMFG

The comparison of simplicial persistence of triangles between quantile thresholding and TMFGs, reported in Figure 4.3, reveals that quantile thresholding struggles to separate the decay of real structures from that of rolling Gaussian generated ones. This could be attributed to the "local" nature of the method, which matches the pairwise interpretation of relations generated from a rolling Gaussian. TMFGs instead, perhaps due to their non-local embedding, provide a consistent ordering of null models with relatively low noise.

The ability to correctly identify persistent motifs throughout the market sample is essential, as we have shown in Section 4.6 that the most persistent motifs are highly systemic. Persistent structures in quantile thresholded graphs present higher and more stable clustering coefficients. This suggests a very localised and compact structure. TMFGs instead present a lower clustering coefficient and a slower decay with $\tau$, as expected since some structures break. This is further evidence of the ability of the TMFG filtering method to identify meaningful persistent structures throughout the market. The issue with quantile thresholding is likely due to the method being merely value-based with no sensible structural constraint, differently from the TMFG.

The ranking of national markets based on their decay exponents in Table 4.1 can be interpreted in terms of the reduction of estimation noise in more liquid markets, as large deviations become less likely and correlations as well as prices more reflective of the underlying generative processes and structures. This suggests that more efficient and capitalised markets are characterised by structures which are more stable in time and better reflected in the data. The decay exponent ranking also leads to the conclusion that more developed markets are characterised by more meaningful underlying structures and cliques, suggesting that systemic risk may pose a greater threat in developed markets.

The results in Table 4.2 support the hypothesis that motifs constitute meaningful structures in markets, beyond their individual edges. These results apply the null model of motif edges from Section 4.4 and show solid evidence to reject it. We can then conclude that highly persistent motifs are not a mere consequence of highly persistent individual edges, but also of the correlation in those edges existing concurrently. This result supports the conclusions above on local filtering methods and generative processes.

Table 4.3 strengthens the importance of persistent motifs. Indeed, the ten most persistent motifs visualised in Figure 4.4 are representative of industry sectors in the NYSE. These sectors are not identified by the motifs with higher edge correlation, which instead are dominated by motifs often due to correlation noise in high volatility stocks. Persistence and the identification of persistent motifs are hence found to be non-trivial with respect to correlation strength of individual edges or motifs. The impact on portfolio diversification (reduction in the variability $\sigma_{sub}$) of the motifs in Figure 4.4 indicates that these structures are highly relevant for systemic risk and portfolio volatility, with high predictive power provided by the long memory property of persistence, which is an intrinsic temporal feature. As these motifs are not characterised by noticeably strong correlations, a common variance optimisation of the portfolio is unlikely to optimise the weights to sufficiently minimise the risk from these highly systemic structures.

## 7.2 Don't panic: the need for smarter portfolio risk

From the introduction to complex systems in Sections 1.2, 1.3 to results in Sections 4.2, 4.3 and discussions in Sections 7.1.1, 7.1.2 we are now able to appreciate that the world is complex. Further, financial markets are a perfect example of a system where such complexity originates from the interaction of a large number of agents, where models and their assumptions break down and only better tools for empirical modeling can help. The discussion in Sections 7.1.1, 7.1.2 highlights the importance of modern methodologies with a "practical focus", which offer a different perspective to traditional methods. What is often lacking is not anymore the awareness that such methods should be used, but rather bridging the gap between empirical insights and ways of applying them. Below we offer our contribution to turning such observational insights into practical methods for portfolio construction. Ours is just a simple proof of concept, but we hope to stimulate further work on the topic.

The systemic relevance of persistent motifs as well as their out of sample forecasting power are shown by the results in Section 4.6 and in [75], where significantly higher out of sample portfolio volatility is observed for the subgraph of persistent motifs. The motif subgraph $\sigma_{sub}$ is significantly above both the mean and median of the random subgraphs' $\sigma_{sub}$ distribution.

This is an example of how just selecting nodes (stocks) from the ten most persistent motifs forms a subgraph with higher long term variability $\sigma_{sub}$. Clearly when aiming for a reduction in systemic risk, low $\sigma_{sub}$ (the opposite) is the objective. The observations from Section 4.6 and [75] lay the ground for the construction of portfolios where asset weights aim to reduce the volatility originating from persistent correlations in motif structures.

Indeed, we provide a proof of concept for applications to portfolio diversification in Section 4.6, where we propose a simple node-specific measure for portfolio weighting and selection. We show the out of sample volatility distribution of random portfolios with weights optimised as $1/P_m(v_i)$ to be significantly lower than the distribution of portfolios optimised as $1/\sigma^2$, a widespread industry standard for portfolio weighting. This result can be explained by the persistence in time being the base of this measure, providing strong out of sample predictive power. Volatility is known to change in the medium to long term for most assets, whilst correlation is also difficult to estimate due to noise in the data and measures. This result greatly enhances the importance and applicability of this work to portfolio optimisation by providing a mapping from persistence-related observations to a direct measure for portfolio optimisation.

It is important to notice that this comparison is meaningful beyond that of an industry standard with a novel approach which outperforms it. The volatility weighting is based on individual assets taken in isolation, weakly influenced by the portfolio's composition of assets. The persistence-based weighting is instead strongly based upon the cliques (therefore the other assets) present in the portfolio. Portfolio composition also influences how the network is filtered, providing a second level of the system's influence on the asset's weighting by persistence. This shows the relevance of structure and how network analysis and complex systems can greatly enhance our understanding of real world systems beyond traditional methods.

As we have shown that both measures are relevant and weakly correlated, future works should investigate a technique to jointly optimise portfolio weighting for persistence and volatility.

## 7.3 Leading vs reactive regulation: a history of complexity and policy

Our discussion of complexity in Chapter 1 has shown how complex multi-agent systems such as financial markets present non-trivial dynamics. These dynamics are often not contemplated in classical economic and financial theories which do not attempt to model the connectedness of such global systems or use bottom-up data-driven techniques.

A major figure in applying such techniques to high frequency systems and, more recently, to the financial and economic system as a whole is J. Doyne Farmer. In his works already preceding the crisis of 2008 and the crash of May 2010 he suggests how simple models of minimal intelligence agents capture the empirical features of high frequency systems better than standard economic models [163]. These models often have a few tuning parameters which can be fit to the data.

Following the 2008 financial crisis Farmer & Foley [164] ring the alarm on the fact that, in spite of the world just having seen traditional economic theories and policies fail, policy-makers are still using common sense and anecdotal analogies to make decisions which impact the daily life of us all. The authors suggest the need to adopt Agent-based Modeling (ABM) to model economic policies and their effects in an effort to move away from the rational agent-like assumptions in economics which had just spectacularly failed.

The concept of overlapping portfolios causing systemic risk was discussed in Section 1.3 initially on the basis of anecdotal modeling and basic financial analysis of empirical data. The more advanced Econophysics works align with that of Caccioli et al.[165] which suggests that the network is "robust yet fragile" with rare, but catastrophic contagion events. The authors suggest a model and methodologies which can be calibrated on real data and used as simple tools for macro stress testing.

In spite of the empirical evidence and the growing literature the more recent work on the current state of complexity adoption in policy suggests we are still very behind [166]. We have made the superficial step forward to adopt complexity terms such as "network", "contagion", "feedback" and "resilience" yet the use of complexity methods is still at a very early stage of adoption. Recent advances in Econophysics and financial complexity provided methodologies for better monitoring and analysis of highly connected and complex economic systems. The importance of these methods could certainly not be overlooked during the COVID-19 crisis where we witnessed how advanced modeling of contagion was in high demand and the global connectedness of our supply chain systems led to massive shortages and concerns.

Rather than always realising the importance of such methodologies once the crisis has erupted and only using anecdotal evidence to prevent the next, we argue that "leading" rather than "reactive" regulation and policy can and should arise from the use of complexity in finance and economic policy at all levels.

In this spirit we advance the current literature on high frequency crashes from the perspective of connected systems and criticality in the sections below.

## 7.4 Complexity and criticality as the two cross sections of systemic risk

In the present chapter we have discussed the more "traditional" aspects of systemic risk in low frequency markets and portfolios. Section 7.3 introduced the importance of regulation in markets at both high and low frequencies. We now move to discuss the former in the general perspective of complexity and criticality of this work.

The results presented in Section 5 cover two cross-sections of systemic risk in high frequency markets: across assets and across market participants. The former is investigated in Sections 6.2, 6.3 where we discuss the structure of co-crashes and the structural characteristics of the co-crash network. The latter is somewhat less obvious and entails understanding how flash crashes happen and how interactions between market participants create positive feedback loops. Our investigation of the volume distribution of anomalous price events reveals a fat tailed distribution with unbounded variance characteristic of critical systems. These events have non-zero probability of being arbitrarily large due to their dynamics triggering a potentially endless loop of actions and reactions from agents in the system.

Whether the systemic risk discussed above is across market participants or assets we look to show clear ties to the works on percolation and phase transitions presented in Section 1.2. These phenomena constitute systemic risk in high frequency financial markets and both cross-sections can be traced back to simple rules of interaction between agents in the system, as discussed below.

### 7.4.1 Complexity: across Assets

In Sections 6.2, 6.3 we present results on co-jump structures in High Frequency markets. We investigate the distribution of co-jump sizes for 300 stocks on 1 minute returns. We highlight features of this distribution such as the finite size effect in the tail and the divergence of small crash frequencies from the distribution. We show how the ranking and structure of crash frequency throughout stocks changes drastically through a phase transition between small and large crash sizes at size $\sim 5$. We quantify this with the Spearman correlation between crash frequency ranks at different co-crash sizes. We then apply a null model of crash frequency at each crash size to test the hypothesis of a phase transition. Finally we highlight how larger crashes are dominated not by the less liquid stocks present in small crashes, but rather by highly liquid and systemic stocks which characterise communities and core-periphery like structure in co-crashes. We suggest that these systemic events can be viewed as communities centered around these most influential stocks.

We know from the literature that heterogeneous network structures can be vulnerable and highly unstable, as well as fragmented if characterised by multiple cores. One of the possible reasons for this can be inferred from the interviews with different market players following the crash of May 6th [167]. Many HFTs highlight the centralised risk constraints for volatility and P&L which cause them to withdraw from the market in case of extreme conditions or losses. As they constitute much of the liquidity in the market in particular for smaller stocks, withdrawing from those causes liquidity draughts. These are often systemic as players have central risk constraints and withdraw from the entire market (across assets) as those are triggered. Further, as systemic stocks crash, arbitrageurs come into play to level prices across the market, hence making the isolated event a systemic one. In this view well-known stocks are not systemic per se, but rather as a result of non-siloed trading by HFTs and ETFs.

In light of the present results future works shall investigate the asyncronous price changes of securities and model spreading dynamics of flash crashes and their directed structure. Lead-lag investigations of causality of these larger crashes are also suggested for future work.

Already from our results one can focus on the stocks we find most systemic in larger flash crashes for co-jumps of size 5 and higher and induce trading halts or limitations to avoid further spreading of these systemic events. This is crucial as our results combined with those of [103] suggest a systemic self-excited process in both frequency and magnitude of those crashes.

We leave the investigation of this structure for future work and highlight that this is of high importance for practitioners and regulators when dealing with market efficiency and stability, particularly as trading frequencies rise and electronic trading grows widespread across securities.

We conclude by observing that volatility and P&L-based trading breaks used by market players may worsen these events and their systemic characteristics since they cause liquidity withdrawals throughout stocks and market players. This introduces systemic synchronisation throughout the market and makes individual assets more susceptible to small trading volumes. Further, we suggest to monitor the stocks we find systemic throughout larger crashes to model the contagion of liquidity crises and halt trading before these spread and distort a larger number of assets. This should also be topic of future work aimed at smart and efficient regulation in High Frequency markets.

### 7.4.2 Criticality: across market players

In the present section we take a step back for broader perspective and provide a discussion of results from Section 5.2 in their context. The findings have crucial practical implications for financial modeling where normality assumptions are still widespread, in particular in risk models. We highlight that extreme events originating from a heavy-tailed distribution with unbounded variance are not compatible with assumptions of normality and their unbounded variance causes to overestimate significantly the variance of the "normal" underlying volume distribution while not detecting the unbounded variance of extreme risk events.

The particular features of the trade volume distribution in flash crashes support the anecdotal explanation of the "hot potato" phenomenon in the literature [68] which suggests that HFT market makers stay in the market under extreme conditions and absorb large trades which cause them to build up large inventories. Once inventory limits are hit the market makers look to reduce it more aggressively as the price moves against them, but can only trade with each other. This causes the positive feedback loop of frenetic trading activities between these market makers trying to reduce their exposure. On top of the market makers trading aggressively and driving the price further up or down on the crash, other players see an increase in traded volume which they interpret as liquidity and start adding to the imbalance in order flow, as they try to reduce their losses during the crash as well. The careful reader can indeed understand from here why these events are particularly characterised by unbounded traded volume distributions which originate from dynamics of positive feedback loops analogous to those of self-organised criticality [1]. The dynamics described above induce a positive feedback loop which is only bounded by human actions or lack of activity by lower frequency agents, but it is unbounded is size in its underlying nature. This further highlights the need for careful regulation of such events as their potential risk is unbounded.

The reader should be aware that our calibration for jump detection discounts changes in overall volatility and detects anomalously large jumps. These time intervals are not only characterised by large price changes, but by anomalous ones. These originate from avalanche-like positive feedback loops which cause black-swan trading activity events. We highlight that we have also analysed other order flow events such as limit orders and deletions, but trades are the only event we see to be bounded for non-jump intervals and unbounded for jump intervals. These events are indeed anomalous as particularly trade volumes are carefully calibrated by market participants under normal market conditions in order to minimise impact and detectability of the order flow. Further, trades are most closely related to price impact and the narrative of "hot potato" dynamics.

## 7.5 The reality of modeling: fat tails matter

As discussed in Section 2.3, event size distributions with fat tails are characteristic of system dynamics in complexity and criticality. These events can be viewed as "energy releases" in the transition of a system from one steady state (stable configuration) to another. Event size distributions in critical systems are often studied for risk assessment of catastrophic events. Some examples are the magnitudes (in their respective measures) of snow avalanches, earthquakes (and their damages) and sandpile and ricepile energy releases between stable configurations. [168, 169, 107, 108, 170, 171].

The hard sciences that deal with such events have developed tools, approaches and methodologies particularly fit for heavy tails and related risks. It is important to notice how such distributions do not provide the same kind of friendly closed form solutions of Gaussian distributions and are hence less studied in introductory undergraduate and graduate statistics courses. This leads precisely to our discussion of quantitative financial modeling below and the consequences of its early history described when introducing Section 1.

Financial modeling has evolved from traditional financial and economic theories which aim for elegant closed form solutions and are based on assumptions of Normality throughout their methods. The assumption of Normality is understandable for models which deal with the body of the distribution in financial measurables. In this restricted context, Normality can be a valid approximation and the tractability of solutions has practical benefits. Risk models instead focus on the tails of the distribution. Yet, they have evolved with the rest of financial models, in spite of Normality approximations breaking down in the tails. Below we discuss results relevant for tail modeling and financial risk and show how Normality assumptions lead to grossly incorrect risk models. This leads well onto out conclusion on the importance of complex and critical systems in financial modeling and how these more advanced methods from Physics and hard sciences should be not only discussed, but also widely implemented in practice, thereby raising the bar for quality of financial modeling in industry as well as academic settings.

In this thesis we investigate anomalous price changes in high frequency markets, also known as flash crashes. We analyse the volume distribution over 1 minute time intervals and show that intervals with no anomalous price changes have a tail exponent $> 2$ which implies a well defined variance. Intervals with anomalous price changes present a volume distribution with tail exponent $< 2$ which is characterised by black swan events and unbounded variance. Our findings reflect the analyses in the literature suggesting how such anomalous crashes originate from positive feedback loops which involve HFT market players. This is in line with findings in self-organised criticality on natural phenomena such as avalanches. Correct statistics is crucial for risk modeling, in particular when focusing on extreme events. The unboundedness of the crash volume distribution makes the argument for the use of robust heavy-tailed statistics in volume forecasting and risk modeling in high frequency markets. Further, the positive feedback loops and the risk of black swan events and market disruptions should prompt market regulators to make sure that risk constraints are being implemented in a fragmented manner across assets, exchanges and players to avoid alignments in trading halts and liquidity withdrawals.

All of the above-mentioned findings are tied together and discussed in Section 7, where we take a step back and reflect on our findings in the context of the literature as a whole. We offer direction for future works in market structure for proper persistence-based portfolio construction. Similarly for flash crash dynamics we suggest to investigate the cumulative traded volume profile during crashes vs. regular time intervals. For co-crashes we provide a preview of future work on our end to open the door for the investigation of co-crash and market structure in general in the context of High Frequency markets. We also provide a review and suggestions for regulatory investigations and potential changes in particular in Section 7.3.

# Chapter 8

# Conclusion

*In this chapter we wrap up the thesis with some overarching perspectives and directions for future work and regulatory improvements - all in the context of complexity and criticality.*

## 8.1 From empirical observations to applied methods: our work in perspective

The present work focuses on systemic risk in financial markets, analysed with the meticulous lens of statistical physics and derived disciplines. We focus on the two major fields of financial market practice: portfolio construction for long term asset allocation and market microstructure for high frequency market making and execution.

Our introduction to portfolio construction draws upon a large body of literature dating back multiple decades. We therefore adopt a chronological approach to the introduction of this topic. We begin with Section 1.1 which discusses the importance of market structure, as represented by the covariance matrix and its derived forms, from the early works of Markowitz to recent works in Econophysics and network filtering upon which our results are based. Section 1.2 follows to introduce the more recent works which focus on complex networks and extreme disruption events on these, from theory to practical observations in the physical and social sciences. Finally, Sections 1.1, 1.3 bring the original works on market structure together with the more recently developed discipline of complex networks to discuss systemic risk in financial markets and empirical works on this. Finally, portfolio construction in Section 2.2 takes such empirical findings to provide heuristics and methods for superior portfolio allocation and risk management.

Our results presented in Section 4 resemble this structure by presenting empirical evidence on the long memory of simplicial complexes in Section 4.2 and validating this through null models in Section 4.3. More empirical analysis and validation of our findings by matching our most persistent motifs with industry sector classification from the Financial Times is presented in Sections 4.4, 4.5. In line with our direction to start from empirical analysis to then propose methods with practical applications, we present in Sections 4.6 a volatility analysis of the most persistent motifs and leverage this to showcase a proof of concept for portfolio weighting based on persistence in Section 4.6.

Our introduction to HFT markets and flash crashes takes a somewhat different approach, in part because of its more recent history and because our contribution and suggestions are not purely methodological, but very much focused on the inner workings of market microstructure and its regulatory framework. With this in mind, Section 1.4 introduces market microstructure and participants in high frequency markets. Similarly, Section 1.5 presents background and current regulation in high frequency markets as well as infrastructure such as exchanges and the related issue of liquidity fragmentation.

We then focus on the topic of HFT systemic risk in the two forms investigated throughout the paper. Section 2.3 introduces flash crashes and their critical dynamics by highlighting how they arise from systemic synchronisation across market players and positive feedback loop dynamics. Analogously, Section 2.4 introduces the more traditional systemic risk across assets in the context of market microstructure with the literature on contagion dynamics and co-crashes.

Our discussion of crash dynamics from Section 2.3 is then connected to our own findings in Section 5.2 which show that crash volume dynamics are critical, with an unbounded variance distribution. These findings are evidence of different dynamics with positive feedback loops during crashes and we discuss implications for practitioners and regulators in Sections 7.5, 7.4.2.

The remaining results sections focus on co-crash dynamics and systemic risk across assets in flash crashes, with a preview of followup work underway. Section 6.2 shows general results around the heavy tailed distribution of co-crash sizes, leading to the statistically validated findings in Section 6.3 which show the phase transition from small to large crash sizes in terms of crash frequency ranking across the considered universe of stocks. Finally, Section 6.4 looks into the structure of the co-crash network to show how highly traded stocks are increasingly involved on average and throughout crashes, as crash size increases. We also show more preliminary results on the correlation significance between trading volume and crash frequency of each stock, at higher crash sizes. The final result presented is then the basis for our followup work on the topic and it shows evidence of core-periphery structures in the co-crash networks with multiple communities and one or more liquid stocks at their core. This is not an indication of causality, but at least provides the basis for major followup work by hinting that there is an interesting network structure in extreme HFT events to be investigated.

## 8.2 The solution: complexity and criticality

We have seen in Section 1 how complex systems with critical dynamics exist in the real world and in particular how critical dynamics are found in natural and anthropogenic extreme events from avalanches to defaults of electric grids.

Financial systems have been of great interest in the context of complex systems in the last decades, as per the review in Section 1.3. We also saw from Section 2.2 how finance practitioners are still dealing with the issues of traditional methods. A similar phenomenon emerges from Section 7.3, where we report that regulators still use traditional finance and economic methods and little to no advanced live monitoring and analysis. The need for ex ante regulation was also highlighted by major authorities in the field [69]. The chapter by Linton et al. [69] suggests that, analogously to how markets have changed fundamentally, so must regulation. The speed of markets no longer allows for reactive regulation, rather regulators must establish a set of automated responses to market disruptions and systemic risk events. To do so, high frequency technology must be used by regulators to gather and process data for live monitoring and intervention, as suggested by other authorities in the field as well [166].

Terminologies from complex and connected systems are now being adopted, but the implementation of methods is still at the very early stages.

After showing and discussing the importance of advanced findings in this work across frequencies and systemic cross-sections we hope to have contributed motivation and practical methodologies to promote the adoption of complex systems approaches to finance from Econophysics and networks to market microstructure. These disciplines and tools are of particular importance in the context of monitoring and controlling risk for both institutions and regulators. More specifically, we highlight how our observations on long memory in persistent structures and the proposed persistence weighting lay the foundation for future work on portfolio weighting based on simplicial complexes and validated filtered structures. Further, our findings on flash crashes should prompt regulators to reason about how these phenomena can be controlled and allow market makers to prepare for this emerging high frequency systemic risk.

Beyond the monitoring aspect of these advanced methods we have understood that dynamics are complex and emerging phenomena are non-trivial throughout financial markets. In line with our review in Section 7.3 we advocate for smart regulation which uses ABMs and other advanced modeling methods to simulate the complex outcomes of new regulation and induced behaviour.

In analogy with how traditional modeling and regression methods are evolving into the use of Artificial Intelligence we suggest that the transition to electronic markets should progress to “smart electronification”, where regulation and market structures are modified to reduce fleeing liquidity and positive feedback loops which result in the observed black swan events. Trading halts are an example of a traditional “temporary” fix, but we believe that more advanced methods and analyses can lead to the smoother and more orderly functioning of markets and avoid such abrupt measures.

In essence, we have shown that complexity and criticality are inherent in financial markets and institutions across frequencies, scales and cross sections. We have also highlighted how the appropriate methods are not widely adopted yet amongst practitioners and regulators. From our historical reviews and more we can see the potentially catastrophic consequences of this. This prompts us to propose the methods presented in this work alongside practical applications in order to promote adoption amongst practitioners and regulars and raise the bar for the quality and soundness of production modeling.

This thesis has summarised a large body of literature on financial systemic risk and we have produced highly relevant and impactful results for systemic risk in both portfolios and high frequency markets. Still, we believe to have just scratched the surface of themes such as temporal persistence of higher order structures in financial networks. Future work should analyse motifs of different sizes and refine a portfolio weighting method which jointly optimises individual volatilities and correlation structures in portfolios. Our investigation of flash crash dynamics and co-crash structures in high frequency markets is perhaps of even greater interest at this time and certainly less documented. Work underway shall investigate the co-crash network at different points in time and any spreading dynamic on it. Further, the cumulative traded volume profile during a flash crash should be obtained and investigated against null models to better understand empirical volume dynamics in jump processes.

# Bibliography


[1] Kim Christensen and Nicholas R Moloney. *Complexity and criticality*, volume 1. World Scientific Publishing Company, 2005.

[2] Shlomo Havlin and Armin Bunde. Percolation ii. In *Fractals and disordered systems*, pages 97–150. Springer, 1991.

[3] Chen Ning Yang. The spontaneous magnetization of a two-dimensional ising model. *Physical Review*, 85(5):808, 1952.

[4] Barry M McCoy and Tai Tsun Wu. *The two-dimensional Ising model*. Harvard University Press, 2013.

[5] Benoit B Mandelbrot. The variation of certain speculative prices. In *Fractals and scaling in finance*, pages 371–418. Springer, 1997.

[6] Benoit B Mandelbrot. The variation of the prices of cotton, wheat, and railroad stocks, and of some financial rates. In *Fractals and scaling in finance*, pages 419–443. Springer, 1997.

[7] Benoit Mandelbrot and Howard M Taylor. On the distribution of stock price differences. *Operations research*, 15(6):1057–1062, 1967.

[8] Laurent E Calvet, Adlai J Fisher, and Benoit B Mandelbrot. Large deviations and the distribution of price changes. 1997.

[9] John C Bogle. Black monday and black swans. *Financial Analysts Journal*, 64(2):30–40, 2008.

[10] M Mitchell Waldrop. Computers amplify black monday, 1987.

[11] George Soros. After black monday. *Foreign Policy*, (70):65–82, 1988.

[12] Benoit B Mandelbrot. *Fractals and scaling in finance: Discontinuity, concentration, risk. Selecta volume E.* Springer Science & Business Media, 2013.

[13] Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

[14] Mark Newman. *Networks*. Oxford university press, 2018.

[15] Ginestra Bianconi. *Multilayer networks: structure and function*. Oxford university press, 2018.

[16] Mark Ed Newman, Albert-László Ed Barabási, and Duncan J Watts. *The structure and dynamics of networks.* Princeton university press, 2006.

[17] Giulio Cimini, Tiziano Squartini, Fabio Saracco, Diego Garlaschelli, Andrea Gabrielli, and Guido Caldarelli. The statistical physics of real-world networks. *Nature Reviews Physics*, 1(1):58–71, 2019.

[18] Laurent Laloux, Pierre Cizeau, Jean-Philippe Bouchaud, and Marc Potters. Noise dressing of financial correlation matrices. *Physical review letters*, 83(7):1467, 1999.

[19] Jerome Garnier-Brun, Michael Benzaquen, Stefano Ciliberti, and Jean-Philippe Bouchaud. A new spin on optimal portfolios and ecological equilibria. *Available at SSRN 3817718*, 2021.

[20] Vasiliki Plerou, Parameswaran Gopikrishnan, Bernd Rosenow, Luís A Nunes Amaral, and H Eugene Stanley. Universal and nonuniversal properties of cross correlations in financial time series. *Physical review letters*, 83(7):1471, 1999.

[21] Moshe Levy, Haim Levy, and Sorin Solomon. A microscopic model of the stock market: cycles, booms, and crashes. *Economics Letters*, 45(1):103–111, 1994.

[22] Laurent Laloux, Pierre Cizeau, Marc Potters, and Jean-Philippe Bouchaud. Random matrix theory and financial correlations. *International Journal of Theoretical and Applied Finance*, 3(03):391–397, 2000.

[23] Bernd Rosenow, Vasiliki Plerou, Parameswaran Gopikrishnan, and H Eugene Stanley. Portfolio optimization and the random magnet problem. *EPL (Europhysics Letters)*, 59(4):500, 2002.

[24] Szilárd Pafka, Marc Potters, and Imre Kondor. Exponential weighting and random-matrix-theory-based filtering of financial covariance matrices for portfolio optimization. *arXiv preprint cond-mat/0402573*, 2004.

[25] Francesco Pozzi, Tiziana Di Matteo, and Tomaso Aste. Exponential smoothing weighted correlations. *The European Physical Journal B*, 85(6):1–21, 2012.

[26] Joël Bun, Jean-Philippe Bouchaud, and Marc Potters. Cleaning large correlation matrices: tools from random matrix theory. *Physics Reports*, 666:1–109, 2017.

[27] Federico Musciotto, Luca Marotta, Salvatore Miccichè, and Rosario N Mantegna. Bootstrap validation of links of a minimum spanning tree. *Physica A: Statistical Mechanics and its Applications*, 512:1032–1043, 2018.

[28] Michele Tumminello, Salvatore Micciche, Fabrizio Lillo, Jyrki Piilo, and Rosario N Mantegna. Statistically validated networks in bipartite complex systems. *PloS one*, 6(3), 2011.

[29] Salvatore Miccichè and Rosario Nunzio Mantegna. A primer on statistically validated networks. *arXiv preprint arXiv:1902.07074*, 2019.

[30] Riccardo Marcaccioli and Giacomo Livan. A pólya urn approach to information filtering in complex networks. *Nature communications*, 10(1):745, 2019.

[31] Rosario N Mantegna. Hierarchical structure in financial markets. *The European Physical Journal B-Condensed Matter and Complex Systems*, 11(1):193–197, 1999.

[32] Rosario N Mantegna and H Eugene Stanley. *Introduction to econophysics: correlations and complexity in finance*. Cambridge university press, 1999.

[33] Michele Tumminello, Tomaso Aste, Tiziana Di Matteo, and Rosario N Mantegna. A tool for filtering information in complex systems. *Proceedings of the National Academy of Sciences*, 102(30):10421–10426, 2005.

[34] Wolfram Barfuss, Guido Previde Massara, T. Di Matteo, and Tomaso Aste. Parsimonious modeling with information filtering networks. *Phys. Rev. E*, 94:062306, Dec 2016.

[35] Guido Previde Massara, Tiziana Di Matteo, and Tomaso Aste. Network filtering for big data: Triangulated maximally filtered graph. *Journal of complex Networks*, 5(2):161–178, 2016.

[36] Alain Barrat, Marc Barthelemy, and Alessandro Vespignani. *Dynamical processes on complex networks*. Cambridge university press, 2008.

[37] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *nature*, 406(6794):378–382, 2000.

[38] Paul Erdos, Alfréd Rényi, et al. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

[39] Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’networks. *nature*, 393(6684):440–442, 1998.

[40] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

[41] Albert-László Barabási, Réka Albert, and Hawoong Jeong. Mean-field theory for scale-free random networks. *Physica A: Statistical Mechanics and its Applications*, 272(1-2):173–187, 1999.

[42] Sidney Redner. How popular is your paper? an empirical study of the citation distribution. *The European Physical Journal B-Condensed Matter and Complex Systems*, 4(2):131–134, 1998.

[43] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *The Structure and Dynamics of Networks*, pages 195–206. Princeton University Press, 2011.

[44] Reuven Cohen, Keren Erez, Daniel Ben-Avraham, and Shlomo Havlin. Breakdown of the internet under intentional attack. *Physical review letters*, 86(16):3682, 2001.

[45] D Stauffer and A Aharony. Introduction to percolation theory, 1992.

[46] Duncan S Callaway, Mark EJ Newman, Steven H Strogatz, and Duncan J Watts. Network robustness and fragility: Percolation on random graphs. *Physical review letters*, 85(25):5468, 2000.

[47] Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.

[48] Franklin Allen and Douglas Gale. Financial contagion. *Journal of political economy*, 108(1):1–33, 2000.

[49] Franklin Allen and Douglas Gale. Bubbles and crises. *The economic journal*, 110(460):236–255, 2000.

[50] Larry Eisenberg and Thomas H Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.

[51] Matthew Elliott, Benjamin Golub, and Matthew O Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–53, 2014.

[52] Franklin Allen, Ana Babus, and Elena Carletti. Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, 104(3):519–534, 2012.

[53] Francis X Diebold and Kamil Yılmaz. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of econometrics*, 182(1):119–134, 2014.

[54] Antonio Cabrales, Piero Gottardi, and Fernando Vega-Redondo. Risk sharing and contagion in networks. *The Review of Financial Studies*, 30(9):3086–3127, 2017.

[55] Matthew O Jackson and Agathe Pernoud. Systemic risk in financial networks: A survey. *Annual Review of Economics*, 13:171–202, 2021.

[56] Joseph Aharony and Itzhak Swary. Contagion effects of bank failures: Evidence from capital markets. *Journal of Business*, pages 305–322, 1983.

[57] Craig H Furfine. Interbank exposures: Quantifying the risk of contagion. *Journal of money, credit and banking*, pages 111–128, 2003.

[58] Paolo Angelini, Giuseppe Maresca, and Daniela Russo. Systemic risk in the netting system. *Journal of Banking & Finance*, 20(5):853–868, 1996.

[59] Paolo Angelini et al. Liquidity and announcement effects in the euro area. Technical report, Bank of Italy, Economic Research and International Relations Area, 2002.

[60] Michael Boss, Helmut Elsinger, Martin Summer, and Stefan Thurner 4. Network topology of the interbank market. *Quantitative finance*, 4(6):677–684, 2004.

[61] Giulia Iori, Saqib Jafarey, and Francisco G Padilla. Systemic risk on the interbank market. *Journal of Economic Behavior & Organization*, 61(4):525–542, 2006.

[62] Giulia Iori and Saqib Jafarey. Criticality in a model of banking crises. *Physica A: Statistical Mechanics and its Applications*, 299(1-2):205–212, 2001.

[63] Giulia Iori, Giulia De Masi, Ovidiu Vasile Precup, Giampaolo Gabbi, and Guido Caldarelli. A network analysis of the italian overnight money market. *Journal of Economic Dynamics and Control*, 32(1):259–278, 2008.

[64] Giulio Cimini, Tiziano Squartini, Diego Garlaschelli, and Andrea Gabrielli. Systemic risk analysis on reconstructed economic and financial networks. *Scientific reports*, 5(1):1–12, 2015.

[65] Stefano Battiston, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli. Debtrank: Too central to fail? financial networks, the fed and systemic risk. *Scientific reports*, 2(1):1–6, 2012.

[66] Jean-Philippe Bouchaud, Julius Bonart, Jonathan Donier, and Martin Gould. *Trades, quotes and prices: financial markets under the microscope*. Cambridge University Press, 2018.

[67] Irene Aldridge. *High-frequency trading: a practical guide to algorithmic strategies and trading systems*, volume 604. John Wiley & Sons, 2013.

[68] Anton Golub, John Keane, and Ser-Huang Poon. High frequency trading and mini flash crashes. *Available at SSRN 2182097*, 2012.

[69] Oliver Linton, Maureen O'Hara, and Jean-Pierre Zigrand. The regulatory challenge of highfrequency markets. *High-Frequency Trading*, page 227, 2013.

[70] Mary L Schapiro. Strengthening our equity market structure. *US Securities and Exchange Commission Economic Club of New York, New York*, 2010.

[71] Maureen O'Hara and Mao Ye. Is market fragmentation harming market quality? *Journal of Financial Economics*, 100(3):459–474, 2011.

[72] Sugato Chakravarty, Pankaj Jain, James Upson, and Robert Wood. Clean sweep: Informed trading through intermarket sweep orders. *Journal of Financial and Quantitative Analysis*, 47(2):415–435, 2012.

[73] Tom McInish, James Upson, and Robert A Wood. The flash crash: trading aggressiveness, liquidity supply, and the impact of intermarket sweep orders. *Financial Review*, 49(3):481–509, 2014.

[74] Mao Ye, Chen Yao, and Jiading Gai. The externalities of high frequency trading. *WBS Finance Group Research Paper*, (180), 2013.

[75] Jeremy D Turiel and Tomaso Aste. Sector neutral portfolios: Long memory motifs persistence in market structure dynamics. In *International Conference on Complex Networks and Their Applications*, pages 573–585. Springer, 2019.

[76] Jeremy D Turiel, Paolo Barucca, and Tomaso Aste. Simplicial persistence of financial markets: filtering, generative processes and portfolio risk. *arXiv preprint arXiv:2009.08794*, 2020.

[77] Jeremy D Turiel and Tomaso Aste. Self-organised criticality in high frequency finance: the case of flash crashes. *arXiv preprint arXiv:2110.13718*, 2021.

[78] Jeremy D. Turiel and Tomaso Aste. Heterogeneous criticality in high frequency finance: A phase transition in flash crashes. *Entropy*, 24(2), 2022.

[79] Giovanni Bonanno, Nicolas Vandewalle, and Rosario N Mantegna. Taxonomy of stock market indices. *Physical Review E*, 62(6):R7615, 2000.

[80] N Vandewalle, F Brisbois, and X Tordoir. Non-random topology of stock markets. *Quantitative finance*, 1(3):372, 2001.

[81] Giovanni Bonanno, Guido Caldarelli, Fabrizio Lillo, Salvatore Micciche, Nicolas Vandewalle, and Rosario Nunzio Mantegna. Networks of equities in financial markets. *The European Physical Journal B*, 38(2):363–371, 2004.

[82] Giovanni Bonanno, Guido Caldarelli, Fabrizio Lillo, and Rosario N Mantegna. Topology of correlation-based minimal spanning trees in real and model markets. *Physical Review E*, 68(4):046130, 2003.

[83] Vasiliki Plerou, Parameswaran Gopikrishnan, Bernd Rosenow, Luis A Nunes Amaral, Thomas Guhr, and H Eugene Stanley. Random matrix approach to cross correlations in financial data. *Physical Review E*, 65(6):066126, 2002.

[84] Gang-Jin Wang, Chi Xie, Kaijian He, and H Eugene Stanley. Extreme risk spillover network: application to financial institutions. *Quantitative Finance*, 17(9):1417–1433, 2017.

[85] Lin Chen, Qian Han, Zhilin Qiao, and H Eugene Stanley. Correlation analysis and systemic risk measurement of regional, financial and global stock indices. *Physica A: Statistical Mechanics and its Applications*, 542:122653, 2020.

[86] Salvatore Miccichè, Giovanni Bonanno, Fabrizio Lillo, and Rosario N Mantegna. Degree stability of a minimum spanning tree of price return and volatility. *Physica A: Statistical Mechanics and its Applications*, 324(1-2):66–73, 2003.

[87] Michele Tumminello, Fabrizio Lillo, and Rosario N Mantegna. Correlation, hierarchies, and networks in financial markets. *Journal of economic behavior & organization*, 75(1):40–58, 2010.

[88] Rosario N Mantegna and H Eugene Stanley. Scaling behaviour in the dynamics of an economic index. *Nature*, 376(6535):46–49, 1995.

[89] Miquel Montero, Josep Perelló, Jaume Masoliver, Fabrizio Lillo, Salvatore Miccichè, and Rosario N Mantegna. Scaling and data collapse for the mean exit time of asset prices. *Physical Review E*, 72(5):056101, 2005.

[90] Giovanni Bonanno, Fabrizio Lillo, and Rosario N Mantegna. High-frequency cross-correlation in a set of stocks. 2001.

[91] Rama Cont, Marc Potters, and Jean-Philippe Bouchaud. Scaling in stock market data: stable laws and beyond. In *Scale invariance and beyond*, pages 75–85. Springer, 1997.

[92] Stephen Satchell and Alan Scowcroft. A demystification of the black–litterman model: Managing quantitative and traditional portfolio construction. *Journal of Asset Management*, 1(2):138–150, 2000.

[93] Fischer Black and Robert Litterman. Global portfolio optimization. *Financial analysts journal*, 48(5):28–43, 1992.

[94] Sebastián Ceria and Robert A Stubbs. Incorporating estimation errors into portfolio selection: Robust portfolio construction. *Journal of Asset Management*, 7(2):109–127, 2006.

[95] Ester Pantaleo, Michele Tumminello, Fabrizio Lillo, and Rosario N Mantegna. When do improved covariance matrix estimators enhance portfolio optimization? an empirical comparative study of nine estimators. *Quantitative Finance*, 11(7):1067–1080, 2011.

[96] Pier Francesco Procacci and Tomaso Aste. Portfolio optimization with sparse multivariate modelling. *arXiv preprint arXiv:2103.15232*, 2021.

[97] Wolfram Barfuss, Guido Previde Massara, Tiziana Di Matteo, and Tomaso Aste. Parsimonious modeling with information filtering networks. *Physical Review E*, 94(6):062306, 2016.

[98] Yuanrong Wang and Tomaso Aste. Dynamic portfolio optimization with inverse covariance clustering. *arXiv preprint arXiv:2112.15499*, 2021.

[99] Diego León, Arbey Aragón, Javier Sandoval, Germán Hernández, Andrés Arévalo, and Jaime Niño. Clustering algorithms for risk-adjusted portfolio construction. *Procedia Computer Science*, 108:1334–1343, 2017.

[100] Vincenzo Tola, Fabrizio Lillo, Mauro Gallegati, and Rosario N Mantegna. Cluster analysis for portfolio optimization. *Journal of Economic Dynamics and Control*, 32(1):235–258, 2008.

[101] Francesco Pozzi, Tiziana Di Matteo, and Tomaso Aste. Spread of risk across financial markets: better to invest in the peripheries. *Scientific reports*, 3(1):1–7, 2013.

[102] F Pozzi, T Di Matteo, and T Aste. Supporting information for the paper:"spread of risk across financial markets: better to invest in the peripheries". 2013.

[103] Lucio Maria Calcagnile, Giacomo Bormetti, Michele Treccani, Stefano Marmi, and Fabrizio Lillo. Collective synchronization and high frequency systemic instabilities in financial markets. *Quantitative Finance*, 18(2):237–247, 2018.

[104] LLC Nanex. Nanex flash crash summary report, 2010.

[105] Neil Johnson, Guannan Zhao, Eric Hunsader, Jing Meng, Amith Ravindar, Spencer Carran, and Brian Tivnan. Financial black swans driven by ultrafast machine ecology. *arXiv preprint arXiv:1202.1448*, 2012.

[106] Mark Paddrik, Roy Hayes, Andrew Todd, Steve Yang, Peter Beling, and William Scherer. An agent based model of the e-mini s&p 500 applied to flash crash analysis. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr)*, pages 1–8. IEEE, 2012.

[107] Zeev Olami, Hans Jacob S Feder, and Kim Christensen. Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes. *Physical review letters*, 68(8):1244, 1992.

[108] Vidar Frette, Kim Christensen, Anders Malthe-Sørenssen, Jens Feder, Torstein Jøssang, and Paul Meakin. Avalanche dynamics in a pile of rice. *Nature*, 379(6560):49–52, 1996.

[109] Giovanni Cespa and Thierry Foucault. Illiquidity contagion and liquidity crashes. *The Review of Financial Studies*, 27(6):1615–1660, 2014.

[110] Giacomo Bormetti, Lucio Maria Calcagnile, Michele Treccani, Fulvio Corsi, Stefano Marmi, and Fabrizio Lillo. Modelling systemic price cojumps with hawkes factor models. *Quantitative Finance*, 15(7):1137–1156, 2015.

[111] James Paulin, Anisoara Calinescu, and Michael Wooldridge. Understanding flash crash contagion and systemic risk: A micro–macro agent-based approach. *Journal of Economic Dynamics and Control*, 100:200–229, 2019.

[112] Dilip Abreu and Markus K Brunnermeier. Synchronization risk and delayed arbitrage. *Journal of Financial Economics*, 66(2-3):341–360, 2002.

[113] Sanjeev Bhojraj, Robert J Bloomfield, and William B Tayler. Margin trading, overpricing, and synchronization risk. *The Review of Financial Studies*, 22(5):2059–2085, 2009.

[114] Pankaj K Jain, Pawan Jain, and Thomas H McInish. Does high-frequency trading increase systemic risk? *Journal of Financial Markets*, 31:1–24, 2016.

[115] Larry Harris. What to do about high-frequency trading, 2013.

[116] Antonio Sánchez Serrano. High-frequency trading and systemic risk: A structured review of findings and policies. *Review of Economics*, 71(3):169–195, 2020.

[117] Marc Potters and Jean-Philippe Bouchaud. More statistical properties of order books and price impact. *Physica A: Statistical Mechanics and its Applications*, 324(1-2):133–140, 2003.

[118] Jean-Philippe Bouchaud, J Doyne Farmer, and Fabrizio Lillo. How markets slowly digest changes in supply and demand. In *Handbook of financial markets: dynamics and evolution*, pages 57–160. Elsevier, 2009.

[119] Jean-Philippe Bouchaud, Yuval Gefen, Marc Potters, and Matthieu Wyart. Fluctuations and response in financial markets: the subtle nature of 'random'price changes. *Quantitative finance*, 4(2):176–190, 2004.

[120] Jean-Philippe Bouchaud, Julien Kockelkoren, and Marc Potters. Random walks, liquidity molasses and critical response in financial markets. *Quantitative finance*, 6(02):115–123, 2006.

[121] Danilo Dessì, Jacopo Cirrone, Diego Reforgiato Recupero, and Dennis Shasha. Supernoder: a tool to discover over-represented modular structures in networks. *BMC bioinformatics*, 19(1):318, 2018.

[122] Nicoló Musmeci, Tomaso Aste, and T Di Matteo. Risk diversification: a study of persistence with a filtered correlation-network approach. *Network Theory in Finance*, 1(1):77–98, 2015.

[123] Martin D. Gould, Mason A. Porter, Stacy Williams, Mark McDonald, Daniel J. Fenn, and Sam D. Howison. Limit order books. *Quantitative Finance*, 13(11):1709–1742, 2013.

[124] Frédéric Abergel, Marouane Anane, Anirban Chakraborti, Aymen Jedidi, and Ioane Muni Toke. *Limit Order Books*. Cambridge University Press, 2016.

[125] Suzanne S Lee and Per A Mykland. Jumps in financial markets: A new nonparametric test and jump dynamics. *The Review of Financial Studies*, 21(6):2535–2563, 2008.

[126] Kris Boudt, Christophe Croux, and Sébastien Laurent. Robust estimation of intraweek periodicity in volatility and jump detection. *Journal of Empirical Finance*, 18(2):353–367, 2011.

[127] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[128] Jeff Alstott, Ed Bullmore, and Dietmar Plenz. powerlaw: a python package for analysis of heavy-tailed distributions. *PloS one*, 9(1):e85777, 2014.

[129] Malte Henkel, Michel Pleimling, and Roland Sanctuary. *Ageing and the glass transition*, volume 716. Springer, 2007.

[130] Kurt Binder and A Peter Young. Spin glasses: Experimental facts, theoretical concepts, and open questions. *Reviews of Modern physics*, 58(4):801, 1986.

[131] E Donth. *The glass transition: relaxation dynamics in liquids and disordered materials*, volume 48. Springer Science & Business Media, 2013.

[132] Tomaso Aste and David Sherrington. Glass transition in self-organizing cellular patterns. *Journal of Physics A: Mathematical and General*, 32(41):7049, 1999.

[133] James R Lincoln and Michael L Gerlach. *Japan's network economy: Structure, persistence, and change*. Cambridge University Press, 2004.

[134] Aaron Clauset and Nathan Eagle. Persistence and periodicity in a dynamic proximity network. *arXiv preprint arXiv:1211.7343*, 2012.

[135] Natacha P Chacoff, Julian Resasco, and Diego P Vázquez. Interaction frequency, network position, and the temporal persistence of interactions in a plant–pollinator network. *Ecology*, 99(1):21–28, 2018.

[136] Leonhard Horstmeyer, Christian Kuehn, and Stefan Thurner. Network topology near criticality in adaptive epidemics. *Physical Review E*, 98(4):042313, 2018.

[137] Michael C Münnix, Takashi Shimada, Rudi Schäfer, Francois Leyvraz, Thomas H Seligman, Thomas Guhr, and H Eugene Stanley. Identifying states of a financial market. *Scientific reports*, 2:644, 2012.

[138] Piero Mazzarisi, Paolo Barucca, Fabrizio Lillo, and Daniele Tantari. A dynamic network model with persistent links and node-specific latent variables, with an application to the interbank market. *European Journal of Operational Research*, 281(1):50–65, 2020.

[139] Tim Hellmann and Mathias Staudigl. Evolution of social networks. *European Journal of Operational Research*, 234(3):583–596, 2014.

[140] Fragkiskos Papadopoulos and Kaj-Kolja Kleineberg. Link persistence and conditional distances in multiplex networks. *Physical Review E*, 99(1):012322, 2019.

[141] Paolo Barucca, Fabrizio Lillo, Piero Mazzarisi, and Daniele Tantari. Disentangling group and link persistence in dynamic stochastic block models. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(12):123407, 2018.

[142] Vincenzo Nicosia, John Tang, Cecilia Mascolo, Mirco Musolesi, Giovanni Russo, and Vito Latora. Graph metrics for temporal networks. In *Temporal networks*, pages 15–40. Springer, 2013.

[143] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

[144] Fabrizio Lillo and J Doyne Farmer. The long memory of the efficient market. *Studies in nonlinear dynamics & econometrics*, 8(3), 2004.

[145] Fabrizio Lillo, Szabolcs Mike, and J. Doyne Farmer. Theory for long memory in supply and demand. *Phys. Rev. E*, 71:066122, Jun 2005.

[146] Tiziana Di Matteo, Tomaso Aste, and Michel M Dacorogna. Long-term memories of developed and emerging markets: Using the scaling analysis

to characterize their stage of development. *Journal of Banking & Finance*, 29(4):827–851, 2005.

[147] Giuseppe Buccheri, Stefano Marmi, and Rosario N. Mantegna. Evolution of correlation structure of industrial indices of u.s. equity markets. *Phys. Rev. E*, 88:012806, Jul 2013.

[148] Sunil Kumar and Nivedita Deo. Correlation and network analysis of global financial indices. *Phys. Rev. E*, 86:026101, Aug 2012.

[149] Giovanni Bonanno, Guido Caldarelli, Fabrizio Lillo, and Rosario N. Mantegna. Topology of correlation-based minimal spanning trees in real and model markets. *Phys. Rev. E*, 68:046130, Oct 2003.

[150] Ruihong Huang and Tomas Polak. Lobster: Limit order book reconstruction system. *SSRN Electronic Journal*, 12 2011.

[151] Zoltan Eisler, Jean-Philippe Bouchaud, and Julien Kockelkoren. The price impact of order book events: market orders, limit orders and cancellations. *Quantitative Finance*, 12(9):1395–1419, 2012.

[152] J Doyne Farmer, Austin Gerig, Fabrizio Lillo, and Szabolcs Mike. Market efficiency and the long-memory of supply and demand: is price impact variable and permanent or fixed and temporary? *Quantitative finance*, 6(02):107–112, 2006.

[153] Austin Gerig. A theory for market impact: How order flow affects stock price. *arXiv preprint arXiv:0804.3818*, 2008.

[154] Rama Cont, Arseniy Kukanov, and Sasha Stoikov. The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88, 2014.

[155] Błażej Mazur and Mateusz Pipień. Time-varying asymmetry and tail thickness in long series of daily financial returns. *Studies in Nonlinear Dynamics & Econometrics*, 22(5), 2018.

[156] Peter Verhoeven and Michael McAleer. Fat tails and asymmetry in financial volatility models. *Mathematics and Computers in Simulation*, 64(3-4):351–361, 2004.

[157] William L Beedles. On the asymmetry of market returns. *Journal of Financial and Quantitative Analysis*, 14(3):653–660, 1979.

[158] Vito Latora, Vincenzo Nicosia, and Giovanni Russo. *Complex networks: principles, methods and applications*. Cambridge University Press, 2017.

[159] M Puck Rombach, Mason A Porter, James H Fowler, and Peter J Mucha. Core-periphery structure in networks. *SIAM Journal on Applied mathematics*, 74(1):167–190, 2014.

[160] Martin G Everett and Stephen P Borgatti. Extending centrality. *Models and methods in social network analysis*, 35(1):57–76, 2005.

[161] Paolo Barucca, Daniele Tantari, and Fabrizio Lillo. Centrality metrics and localization in core-periphery networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(2):023401, 2016.

[162] Marcio Rosa Da Silva, Hongwu Ma, and An-Ping Zeng. Centrality, network capacity, and modularity as parameters to analyze the core-periphery structure in metabolic networks. *Proceedings of the IEEE*, 96(8):1411–1420, 2008.

[163] J Doyne Farmer, Paolo Patelli, and Ilija I Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102(6):2254–2259, 2005.

[164] J Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460(7256):685–686, 2009.

[165] Fabio Caccioli, Munik Shrestha, Cristopher Moore, and J Doyne Farmer. Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245, 2014.

[166] Stefano Battiston, J Doyne Farmer, Andreas Flache, Diego Garlaschelli, Andrew G Haldane, Hans Heesterbeek, Cars Hommes, Carlo Jaeger, Robert May, and Marten Scheffer. Complexity theory and financial regulation. *Science*, 351(6275):818–819, 2016.

[167] SEC CFTC and US SEC. Findings regarding the market events of may 6, 2010. *Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*, 104, 2010.

[168] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Physical review letters*, 59(4):381, 1987.

[169] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Physical review A*, 38(1):364, 1988.

[170] Kim Christensen, Álvaro Corral, Vidar Frette, Jens Feder, and Torstein Jøssang. Tracer dispersion in a self-organized critical system. *Physical review letters*, 77(1):107, 1996.

[171] Kim Christensen and Zeev Olami. Sandpile models with and without an underlying spatial structure. *Physical Review E*, 48(5):3361, 1993.