University of Nottingham
UK | CHINA | MALAYSIA

Describing financial crisis propagation through epidemic modelling on multiplex networks

Thesis submitted to the University of Nottingham for the degree of **Doctor of Philosophy, September 2024.**

Malvina Bozhidarova

20302453

Supervised by

Reuben O'Dea
Frank Ball
Gilles Stupfler
Yves van Gennip

The project is in collaboration with

Russell Group

Abstract

In this thesis we employ various methods from network science, together with epidemic modelling and extreme value theory, to build and analyse financial crisis propagation models. We use stock price, geographical location, and economic sector data for a set of 398 companies to construct multiplex networks and propose a novel framework for modelling financial contagion using an SIR (Susceptible–Infected–Recovered) epidemic model. We compare different shock transmission models and explore their effectiveness in predicting the spread of financial shock during the 2008 financial crisis and the 2020 financial crisis. To enhance the accuracy of our models, we introduce a change point detection method to detect significant changes in historical crisis data and integrate them into our models accordingly, improving their adaptability to major market events. Additionally, we develop a model that prioritizes recent observations under the assumption that they provide a more accurate reflection of current market conditions and trends, assigning greater weight to recent data while reducing the influence of older data. Our findings highlight the importance of the multiplex network structure, differentiating between various transmission pathways, and demonstrate the value of incorporating change points and weighted observations for more accurate predictions of affected companies, sectors and continents. In addition, there is no single model that performs best in all scenarios. Hence, different predictions tasks, whether forecasting the number of infected companies or making company-specific predictions, may require distinct approaches to achieve more accurate results.

Acknowledgements

First and foremost, I would like to thank my quartet of supervisors for the invaluable support throughout these challenging four years. I am grateful for the opportunity to work on such a fascinating project and for your patience and guidance over this period. Together, my supervisors have kept me sane and on track, despite starting my PhD during the crazy old times of COVID-19 pandemic. I would also like to thank Reuben O'Dea for handling all the tedious administrative tasks and Yves van Gennip for welcoming me to Delft and being an exceptional host.

I want to express my deep gratitude to my family and friends for their unwavering support and encouragement. A special thanks goes to my parents for supporting me not only during my PhD studies but also in every decision I have made throughout my life. I am also deeply thankful to my partner, Marin, for helping me maintain a work-life balance and standing by me through all the ups and downs of my PhD journey.

The work presented in this thesis has been done in collaboration with Russell Group Limited. I would like to acknowledge their support for providing the data and the funding necessary for this project, as part of the EPSRC CASE doctoral studentship.

Contents

Chapter 1

Introduction

1.1 Research outline

The increasing frequency, magnitude and international scope of financial crises has made global financial stability one of the key concerns of economic policy-makers and decision-makers [71]. The global financial system is characterised by extensive interconnectedness, where companies can be connected in multiple ways, such as via mutual claims and obligations [92], or through transactions between them [25], forming a network structure. Understanding how financial crises experienced by certain companies or sectors can spread, potentially leading to wider crises, is self-evidently of interest to policy-makers, investors and business owners, as even a small disruption in a single company may cause long-term problems and significant losses, as well as a global financial crisis [155]. Given the impact of financial links on global economies, it is no surprise that financial system analysis is of great importance. Insights into how financial crises propagate across a network of companies would be beneficial not only to current large organisations, but also to the development of smaller businesses and the formation of new financial ties.

The network structure is important in determining how an initial shock propagates across the system of companies. Hence, rather than being viewed as a standalone entity, a company and the risk it faces should be evaluated in conjunction with the network of companies with which it interacts and its broader financial environment. Allen and Gale [11] study how network topology influences the propagation of risk in financial systems. They emphasize that the existence of network connections can create channels for contagion spread, increasing the probability of risk transmission within the network. Since the publication of this pivotal paper, network models have increasingly been used in both theoretical and empirical studies of financial contagion.

The goal of this thesis is to build and analyse financial shock propagation models. Our objective is to present novel methodologies and approaches for modelling shock propagation in multiplex financial networks, integrating techniques from network science, epidemiology and extreme value theory. In particular, we explore how network science can be used to model the spread of financial contagion within a system of companies. Our main focus is on predicting what will happen in the future of a financial crisis, given events in the preceding days of the crisis. This research is crucial for investors, businesses and policymakers, as it offers insights into predicting and mitigating future risks. We demonstrate the effectiveness of our approach by considering the two most recent crises: the 2008 Global financial crisis and the 2020 COVID-19 crisis.

1.2 Motivation

1.2.1 Why study financial crises?

Financial crises are the primary economic concern of our time [59]. As the 2008 Global financial crisis developed, the world received a reminder of the significance of these events. Studying financial crises is crucial due to their historical propensity to recur over time, their ability to affect both developed and developing countries, and the severe and long-lasting harm they can do to economies and societies.

The 2008 financial crisis is the most severe crisis in the global economy since the Great depression (1929–1939). The crisis started in the United States in 2007 when the US housing market began to collapse and prices plummeted. The subprime mortgage crisis quickly evolved into a financial catastrophe as in 2008 two giant mortgage-lending companies, Fannie Mae and Freddy Mac, filed for bankruptcy. Soon the banking sector began to struggle too as the well-known bank Lehmann Brothers declared bankruptcy. The instability surrounding the credit and equity markets further compounded the issues already affecting the banking sector as a whole.

As the financial downturn in the US continued, the recession also started spreading globally. In the last decades there has been a significant growth in the interactions between banks, investors and companies, so the crisis was accelerated by the increased global economic interdependence. As a result, the entire stock market had a sharp decline in response to the crisis towards the end of 2008. Numerous banks and subsidiaries were forced to close throughout Europe; banks in Germany, for instance, experienced significant losses as they had made huge investments in American real estate securities [111]. Furthermore, Europe faced its own housing bubble, with home prices in Ireland,

Iceland, Spain and Denmark plummeting by up to 40% by the end of 2011. In addition, the global trade declined by more than 15% from 2008 to 2009 [89]. The decline in exports in the US and Europe had an impact on suppliers from China. Chinese manufacturers, who were heavily dependent on the US market, experienced difficulties when their primary markets withdrew their investment and decreased demand. As a consequence, more than 10 million workers lost their jobs in China [89]. In Europe and the US the unemployment rates also increased significantly, reaching 10% at the end of the crisis, the highest value since the 70's [213]. Overall, as a result of the crisis the US alone lost more than 8.7 million jobs, while worldwide more than 30 million jobs were lost.

In addition, the 2008 financial crisis caused protests, riots and unrest worldwide. The median family income in the US dropped from $\$126,000$ in 2007 to $\$77,400$ in 2010, leading to more than 46 million people approaching the US government for food stamps and costing the treasury $75.7 billion during 2011 [19]. Falling incomes caused unrest not only in the United States but also in Europe. Numerous protests were organised around the world: on November 16, 2011 thousands of angry US citizens protested against social inequality and poverty at Succoth (a park in New York); on November 4, 2011, the Italian Premier, Mario Monti, announced radical measures to cut public spending, increase tax rates and the retirement age limit, leading to a series of protests continuing throughout the whole year 2012; on July 20, 2012 in Spain a record number of people protested against cuts in wages, increases in tax rates and reductions in the education and healthcare budgets [19]; in October 2009, Greece declared a financial emergency and required a €45 billion bailout from the EU and the International Monetary Fund (IMF), leading to cuts in wages and increases in taxes, which caused series of protests.

The above discussion has demonstrated that the 2008 crisis had a profound

impact on lives, economies and politics worldwide. As financial crises become increasingly common, it is crucial to understand their underlying mechanisms and develop strategies to prevent and mitigate their effects.

1.2.2 Why use network science?

Over the past two decades, advancements in data collection and computing power have increased the availability of data across many natural and man-made systems [164]. This surge in data availability has led to the emergence of network science, a multidisciplinary field in which researchers explore the structural properties and characteristics of complex networks [163]. Within this context, network science has become an increasingly popular framework for modelling financial systems. As a result, many different financial network models have been developed in recent years [83, 154, 221], all demonstrating the effectiveness of the use of network theory approaches to model financial systems.

The 2008 financial crisis illustrated how the global financial interconnectedness can lead to significant losses and global financial instability. The crisis has emphasized the importance of studying network linkages and interactions between financial institutions for understanding systemic risk [13]. In general, a network is a collection of nodes and the links between them. In the financial context the nodes usually represent financial institutions and the links between them can represent an interaction between them, such as common investors or similar portfolios.

Network science offers many advantages for studying and modelling financial crises. For instance, traditional financial models usually focus on individual institutions or markets, not considering their connections. However, by

taking into account the wider financial system and the complex network of relationships among its individual components, network science offers a comprehensive perspective, which is crucial for identifying systemic risks that may not be apparent when examining institutions individually. For example, Qiu et al. [186] and Cai et al. [49] argue that using network science can enrich our understanding of financial systems. They demonstrate the importance of the network approach for evaluating financial stability and understanding how the risk associated with a single institution can impact the entire system. In addition, network analysis provides tools to identify key institutions (nodes) and connections (edges) that play a significant role in the stability of the financial system. Regulators and policymakers can then develop strategies to protect the most influential parts of the network, thereby enhancing overall resilience [10]. In addition, financial crises often spread through contagion, where the failure of one institution leads to a cascade of failures across the whole financial system. Network models are particularly useful for simulating these so-called spillover effects, allowing for better understanding and mitigation of potential chain reactions in financial markets [178]. Allen and Babus [10] argue that regulations that focus on individual institutions, while also considering vulnerabilities from network interconnections, can help prevent local crises from spreading globally. All the advantages mentioned above, combined with the growing body of research demonstrating its effectiveness in financial modelling and the considerable potential for additional contributions to this field, motivates the use of network science as a foundation for the research presented in this thesis.

1.3 Thesis structure

The remainder of the thesis is organised into seven chapters, whose content is described in the following paragraphs.

Chapter 2 provides a review of the existing literature in the fields of network science and epidemiology, focusing on their application to financial modelling. Specifically, we review how the fundamental methods and concepts developed in these fields can be used to model financial contagion. We explore some of the key studies related to these topics and discuss their significant findings and conclusions.

Then, in Chapter 3, we introduce an innovative framework for modelling the spread of financial crises. Utilizing stock price, geographical location, and economic sector data from 398 companies, we build a financial network and employ a Susceptible–Infected–Recovered (SIR) transmission model to simulate the spread of financial shocks among the network of interconnected companies. Our findings highlight the importance of differentiating between various transmission pathways. Additionally, the framework is versatile in nature and can be applied to any financial crisis, whether within networks of companies, countries or economic sectors. Most of the work presented in Chapter 3 is published in [45].

In Chapter 4, we examine the use of change point detection techniques to enhance the model's responsiveness to major market events. External shocks, like geopolitical tensions and natural disasters, can significantly impact stock prices and market dynamics. Hence, it is crucial to build a model that can quickly adapt to such changes. By using change point detection, we identify and integrate these events into the model, enabling its predictions to adjust to the changing market. The method illustrates the importance of considering

significant past events for achieving more accurate forecasts.

Chapter 5 presents a model that prioritizes recent observations under the assumption that they provide a more accurate reflection of current market conditions and trends. This approach assigns greater weight to recent data while reducing the influence of older data. The results indicate that incorporating this weighting of past observations improves accuracy compared to models that treat all past data equally important. In addition, we combine the observations weighting approach with the model introduced in Chapter 4. Specifically, we use both change point detection techniques and weighting of past observations in order to build a more accurate model that can adapt to significant market changes quickly and prioritizes recent observations simultaneously.

Chapter 6 offers a comprehensive comparison of the various models discussed throughout the thesis. It evaluates each model based on its effectiveness in addressing different types of predictions, including predicting the number of infected companies, the continents and economic sectors that will be most affected, and the specific companies that will be infected during future crisis days. The chapter then identifies the most suitable model for each specific prediction scenario.

In Chapter 7 we explore how different network and node characteristics can influence risk propagation. In particular, we study the correlation between different centrality measures and the probability of infection. We also analyse if a company's credit score can be used as a predictor how long it would stay infected in a crisis. Finally, we conduct a simulation study of various counterfactual crisis scenarios, such as initiating a crisis with the most central nodes or varying the number of initially infected companies, demonstrating the practical value of such studies.

Finally, Chapter 8 summarises the thesis, identifies existing research gaps and outlines potential directions for future research.

Chapter 2

Literature review

In this chapter we conduct a comprehensive review of the relevant literature on network science and epidemic modelling on networks, which are the two main research areas of our thesis. Our aim is to outline the main methods and concepts that will be used in this thesis. We review the key studies that have been conducted on these topics and highlight the significant findings and conclusions.

Specifically, in Section 2.1, we provide an overview of network science, highlighting its fundamental concepts and theories and its main contributions to economics and finance. Then, Section 2.2 discusses the relevant literature on modelling epidemics on networks and examine some of the significant findings in the field of finance. Finally, in Section 2.3 we discuss some of the research gaps in the existing literature and explain how our thesis addresses these issues.

2.1 Network science

Technology and computation have improved dramatically in the twenty-first century, resulting in a huge amount of data for a wide range of natural and synthetic systems, including biological systems such as the human nervous system, technological systems like the Internet, and social systems such as friendship groups. The increased data availability has sparked the formation of a new field of study known as network science (also called network theory), in which researchers have developed various methods to model and understand the behaviour of complex systems [164]. The goal of network science is to create models that reflect the structural properties of real-world systems and investigate how these structures are linked to network characteristics. In this section, we discuss some of the key ideas, applications and recent developments in network theory.

2.1.1 Background and history

Network theory, which apply concepts from graph theory, is a mathematical approach to the study of complex systems and their relationships. Specifically, graph theory is a branch of mathematics that studies graphs as abstract structures representing relationships between objects, while network theory applies these concepts to real-world systems, focusing on their structure and dynamics. The beginning of graph theory is considered to be in the 18^{th} century, when Leonard Euler published a paper providing the solution of the Königsberg bridge problem [85]. The city of Königsberg includes four different islands, connected by bridges (see left plot in Figure 2.1). The question was if it is possible to walk through all the islands, without crossing any bridge twice, and starting and finishing the walk on the same island. Euler repre-

Figure 2.1: The 'Seven bridges of Königsberg' problem illustrated in graphs. A graph, as seen in the most-right plot, is represented by a set of points called vertices or nodes, connected by a set of lines called edges. The image is taken from https://en.wikipedia.org/wiki/Seven_Bridges_of_K%C3%B6nigsberg .

sented the city as a graph (see right plot in Figure 2.1), where the islands are represented as vertices (blue circles) and the bridges are represented as edges (black lines) and proved that it would be possible only if all islands (vertices) had an even number of bridges connected to them (even degree). He also generalised this statement for any number of bridges and islands. His sketch is one of the first examples of graphs.

Since then, graph theory has been constantly developed and extensively studied. In the last century, graph theory has evolved significantly, leading to the introduction of a variety of definitions and methods for characterising the structural aspects of graphs, which led to the development of network theory. The first widely recognised examples of the application of graph theory methods to the study of real-world networked systems occurred in the field of social networks analysis in the early part of the 20^{th} century [164]. This early research was primarily limited to small networks that could be drawn by hand and the investigation of the attributes of individual vertices, due to constraints in the availability of data and computing capacity at the time. Jacob Moreno's study on the number of friendships held by individuals within friendship groups is an example of such research [152]. Since the late 1990s and early 2000s, network science has grown into a well-established and productive field. Initial research relied on empirical networks and sought to

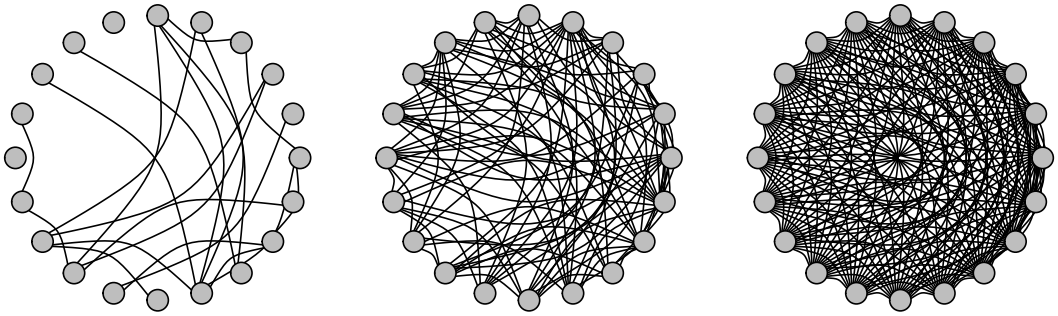

Figure 2.2: Erdős–Rényi network models with 20 nodes and probability of an edge existence $p = 0.1,\ p = 0.5,\ p = 1$ from left to right.

better understand how networks form and what qualities they have. Degree heterogeneity (variation in the number of neighbours of each node) has been discovered to be a common feature of many real-life networks, ranging from biological to political and social networks [21]. As a result of better knowledge of how real-world networks arise and evolve, a variety of theoretical and synthetic network models have emerged. One of the most popular synthetic network models was created by Erdős and Rényi in 1959 [97]. In this model, a graph is created with a given number of nodes n, where each possible (undirected) link between each pair of nodes is present with an equal probability p. Figure 2.2 shows three different Erdős–Rényi network models, each with 20 nodes and each edge having a probability of being present $p = 0.1$, $p = 0.5$, $p = 1$, respectively, independently of the other edges. In 1998 Watts and Strogatz published a paper in which they proposed another very popular random model used in network theory, called a small-world network model [225]. The small-world model is a network model constructed to 'mimic' properties, which are found to be present in many real-life systems, such as social networks [184] and brain networks [28]. Such properties include small distances between the nodes and a high level of clustering. Small-world networks are constructed in the following way. The first step is creating a ring with N nodes and connecting each node with its k nearest neighbours on each side. Then each edge in the graph is randomly rewired to another node with a prob-

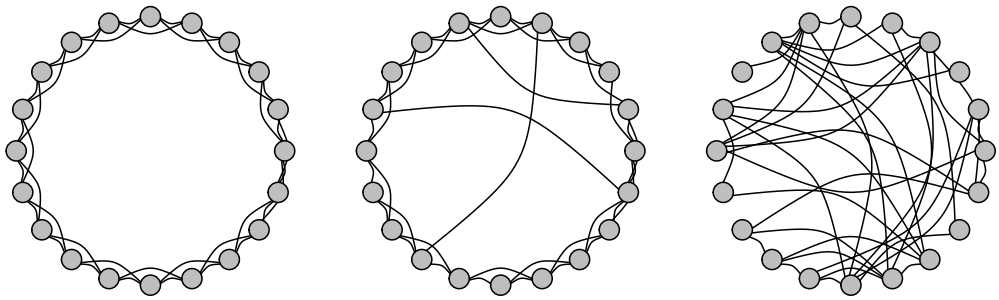

Figure 2.3: Illustration of the small-world network construction. Starting with a ring lattice where each node is connected to its $k = 2$ nearest neighbours on each side (left plot), the edges are rewired with probability p. The graph in the middle shows a realisation with $p = 0.1$ and the right plot illustrates an example where $p = 1$.

ability p. An example is shown in Figure 2.3. In 1999, Barabási and Albert introduced their preferential attachment model [22]. The model begins with a small graph of just a few nodes. New nodes are added sequentially, each connecting to existing nodes with a probability proportional to their degrees ('rich-get-richer' principle, based on the idea that highly connected nodes have a greater ability to form new connections). This process continues until the network reaches the desired size n. These models are typically used to represent social networks, where popular individuals can continue to make friends via their pre-existing connections, while the less popular members of the network remain isolated. An example is shown in Figure 2.4. It can be seen that in comparison to the model where new nodes are attached to the existing ones randomly (with equal probability of attachment), the preferential attachment model produces a network, where most nodes have a low degree, but there are a few highly connected nodes. This feature of the node degree distribution is typically displayed by power-law distributions. All these theoretical models of networked systems are valuable as they help uncover the fundamental mechanisms driving the formation and development of networks found in real-world scenarios.

In recent years network theory has become a common tool for analysing the

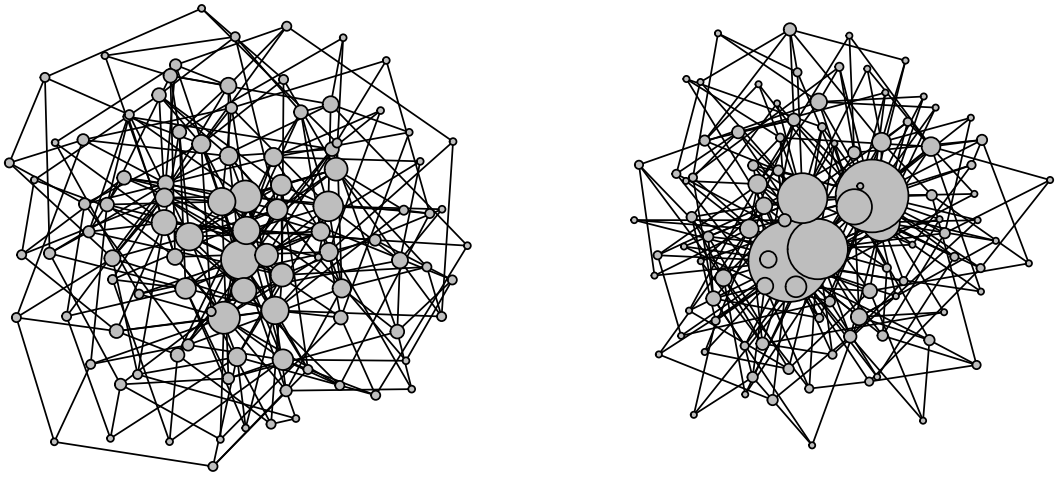

Figure 2.4: The figure illustrates two models of network attachment. The left plot represents an equal probability of attachment (each new node is connected to an existing node randomly), while the right plot represents a preferential attachment model (each new node is connected to an existing node with probability proportional to node's degree). The initial network consists of 3 nodes. The size of each node is proportional to its degree.

connections in social [8, 224], biological [29, 84, 144, 205] and financial systems [11, 72, 93]. Previous work on social network analysis [224] has shown that analysing community structures has important applications in marketing studies [75], homeland security [230] and epidemic modelling [84, 144]. In addition, network theory is successfully applied in modelling interactions of neurobiological systems [29], describing the protein connections in human cells [205], flood risk assessment in rivers' basins [82] and studying the topology of World Wide Web networks [8].

Many complex networks are dynamic in nature. Usually, we say that a network is 'dynamic' if its set of edges, and sometimes its set of vertices, change with time. Alternatively, we can say that a network is 'dynamic' if the vertex or edge attributes are functions of time. If both cases are present, then we say that the network is 'coevolving' [131]. Edge-set dynamics is usually seen in social networks. For example, in a social network friendship connections form and break over time. Newcomb [161] used temporal data to study the friendship creation processes in a group of males living in the same house.

Also, Moody et al. [151] studied the romantic networks between the students in a school, trying to model how sexually transmitted diseases spread. For such networks it is important to capture how often a student changes a partner. In addition, node and edge attribute dynamics is commonly seen in neuroscience [199]. For example, during development or learning, an adjustment of myelin thickness or internode length can influence conduction speed of myelinated axons [14]. Other fields where the usage of dynamic networks is very popular include biology and medicine [7, 40, 173]. Applications include modelling protein interactions [6], cancer research [141, 210] and epidemiology [20]. Dynamic networks are also used to model complex networks in covert operations. One of the most famous examples where dynamic network was used is the capture of Saddam Hussein [44] where cellular phone data was used to study his location and contacts [98].

However, in practice most network analyses assume that the network of interest is static. One reason is that the methodologies for analysing dynamic networks are not as well developed as those for static networks. Also, in many cases working with temporal networks is computationally expensive and it is sometimes difficult to collect data over time. It is important to note, however, that it is not always necessary to use dynamic networks, as a static network can sometimes provide a sufficiently good approximation to the temporal system [127]. Overall, modelling dynamic networks is not straightforward and each different problem requires different techniques. Nowadays, there are still not many models and methods for dynamic networks. However, the field of dynamic networks is progressing fast, improving our understanding of interconnected systems and leading to significant advancements in many areas.

Network theory has become an increasingly important tool for analyzing the world around us. The history of network theory reflects a long and rich tradition of using mathematical and computational tools to understand the struc-

ture and function of complex systems. This overview has highlighted some of the methods developed to construct and analyse network structures, theoretical models to represent networks and empirical studies on real networks.

2.1.2 Multilayered networks

Real-life systems are usually too complicated to be described accurately by a single-layered network. Frequently, the system of interest will contain different types of relations between the nodes and/or different types of nodes for which a single-layered approximation may be inadequate. In such cases, a multilayered network is required. In such networks, there are two types of edges: intra-layer edges, that connect nodes within the same layer, and inter-layer edges, which connect nodes in different layers (see Figure 2.5a). For example, sometimes in a social network we might have different types of relationships, such as friendship, coworker, or family connections. Each of these relationships can be defined in a distinct network layer. Networks in which each edge is categorized by its type are called 'multiplex' or 'multirelational' networks [129]. In multiplex networks, each node is present in each layer and the inter-layer links only connect instances of the same node in different layers. In other cases we might have the same connection types, but different types of nodes, such as for example males and females in a social network. Such network is called a 'node-coloured' network or a 'network of networks' [31]. In these networks each layer represents a set of nodes of the same type (see Figure 2.5b).

Multilayered networks are employed in a variety of fields of study, such as neuroscience [64], ecology [179], biomedicine [106], epidemiology [118] and cancer research [51]. Although their usage has become very popular in recent years, there are relatively few models and analysis tools for multilayer networks. In

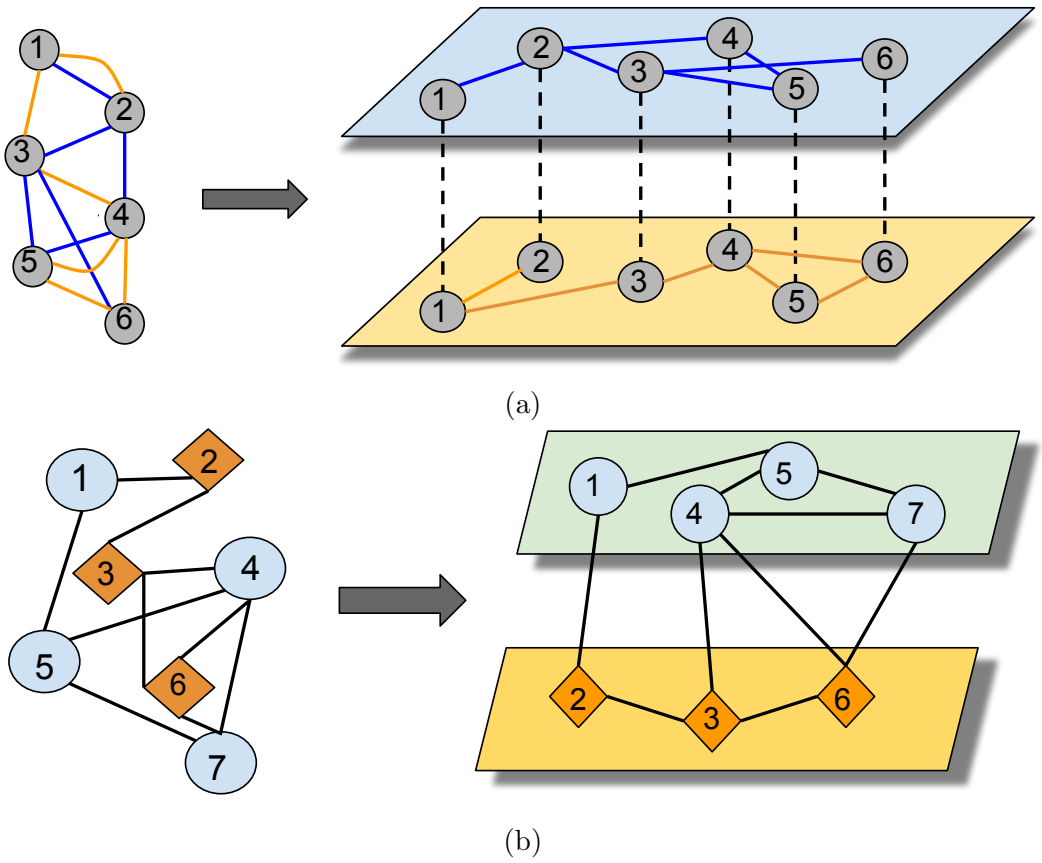

Figure 2.5: The figure illustrates a) a multiplex (or edge-coloured) network, where each layer consists of the same nodes, but represents a different type of connection; b) a node-coloured network, where nodes in each layer are of the same type, and nodes in different layers represent different types.

fact, many of the methods used in multilayer network analysis have been developed from the methods used in the analysis of single-layered (monoplex) networks. For example, the notions of node degrees [30], centrality measures [69], clustering coefficients [62] and community structures [110] have been generalised to multilayer networks. In many cases, analysis of a multilayer network is undertaken by 'overlapping' it into a monoplex network (simplifying it into a single-layer network by combining or merging the different layers into one) and then using well-known methods for analysing single-layered networks. The negative consequence is that usually some important data is lost. Nevertheless, the field of multilayered networks is developing quickly, as the research interest in the area is constantly growing.

Overall, multilayered networks are a powerful tool for modelling complex systems that exhibit multiple types of interactions and dependencies. The ability to represent and analyze multiple layers of interactions allows for a more accurate and nuanced understanding of how different components of a system are connected and how they interact. As the complexity of systems continues to increase, multilayered networks will become increasingly important for understanding and managing complex systems. Further research in this area is needed to develop more sophisticated methods and tools for analysing multilayered networks and to apply them to new domains and applications. Ultimately, the continued development and application of multilayered networks will contribute to a better understanding of complex systems and inform the development of more effective strategies for their management and control.

2.1.3 Networks in finance

In his 1758 work, 'Tableau Economique' [187], François Quesnay conceptualised the circular flow of financial funds in an economy as a network. This

is considered to be the first publication on the topic of financial networks. His fundamental idea has been employed in creating so-called financial flow of funds accounts, which provide a statistical depiction of the movement of money and credit within an economy [158]. Since then the application of networks in economy and finance has evolved and expanded. Many theoretical models that focus on trade linkages [95], financial linkages [70], as well as models on information asymmetries and investor behaviour [76] are among those used in the economics and finance literature. The types of networks in finance can be divided into two main categories: similarity-based networks and direct interaction networks, as explained in [50]. In the first type a link (or edge) between two nodes represents a similarity of the agents (companies, industries, stocks, markets). In the second, the links represent a direct dependence between agents (ownership, customer-supplier relationship, partnership, money flow, transaction between agents). In Section 2.1.3.1 we discuss popular methods for constructing similarity-based networks, and in particular their applications in finance. In Section 2.1.3.2 we discuss some of the existing literature related to direct interaction networks in finance. Finally, in Section 2.1.3.3 we describe the applications and importance of multiplex networks in finance.

2.1.3.1 Similarity-based networks

A similarity-based network, or a similarity graph, is a network where a link between two nodes (agents) indicates how similar they are. A typical similarity measure used in literature is the correlation coefficient between two time series. Different correlation measures have been used such as Pearson correlation, Kendall's tau and the Spearman rank coefficient [197]. Other popular methods to construct similarity graphs include, but are not limited to, using Euclidean distance [120], cosine similarity [231] and Hamming distance [169] between

nodes. These distances are computed based on feature vectors associated with each node, which quantify the characteristics or attributes of the nodes.

Typically, in similarity graphs a link between two nodes exists if there is a strong similarity between them. In such cases one needs to construct a criterion by which to decide if this similarity is strong enough, such that a link exists between the agents [50]. In general, consider a system of N nodes, each represented by a set of variables. These variables might be different properties of the nodes, or represent the values of the variables at different times (in which cases each node will be represented as time series). A popular way to create a similarity graph is to consider an $N \times N$ similarity matrix S, with entries $s_{i,j}$ representing the similarity between node i and node j. Then, if $s_{i,j}$ is larger than some threshold, there is a link between nodes i and j, typically with weight associated with the value of $s_{i,j}$. However, choosing the threshold is not straightforward. Therefore, an important topic in the study of similarity networks is filtering the similarity matrix (or graph) such that only the most 'important' entries (or edges) are kept. Some of the popular methods for filtering the networks include the minimum spanning tree (MST) and planar maximally filtered graph (PMFG) approaches. The latter is the filtering method chosen in this thesis. More information and detailed descriptions of each of these methods are found in [50].

Similarity-based networks are popular in finance. One of the most popular ways to construct similarity graphs in finance is using stock prices and measuring different types of correlation between them. For example, Bonanno et al. [42] investigated how the cross-correlation between pairs of 100 US stocks depended on the time horizon used to compute the stock returns. The authors found that clusters of economic sectors emerge in the network structure as the time horizon gets larger. The correlation based similarity graph ap-

proach has also been applied to stock indices and cryptocurrencies data. For instance, the authors of [43] studied the correlation between 51 different countries' stock indices and found that there are regional clusters of European and Asian countries. In addition, [101] built a correlation-based similarity network to study the transmission of price information among different bitcoin market exchanges.

Correlation measures can be useful in modelling financial risks. For instance, modern portfolio theory proposes that diversification can lessen the volatility of a portfolio's returns, thereby reducing risk [133]. By examining the correlation coefficient between historical returns, investors can determine if incorporating a new investment into a portfolio will enhance its diversification. Even though the traditional correlation measures mentioned above have been extensively used in finance, they typically consider the entire distribution of the time series. However, the dependence in the tail of the distribution can differ significantly from that in its mid-range. When extreme realizations are important, such as estimating the dependence between risky returns, extreme value theory techniques are commonly used. More details are given in the following section.

2.1.3.1.1 Tail dependence to model financial networks

Extreme value theory (EVT) is a statistical approach that is specifically designed to analyse rare and extreme events. It has found a wide range of applications in fields such as climatology [107], hydrology [218], biomedical data processing [192], sports [94], public health [214] and medicine [38]. While there are many theoretical and statistical tools for univariate extremes (see [32] for more information), techniques for high-dimensional and complex data sets are still limited. Only recently have appropriate ideas of sparsity been

created, as well as linkages to other fields such as machine learning, graphical models and high-dimensional statistics [82]. Despite the fact that both EVT and network science have been successful separately in modelling financial and economic data, there has been limited research merging the two. Using the concept of conditional dependence, recent work by Engelke and Hitz [81] illustrates how to integrate the two by building graphical models on extremes. Another recently developed class of extreme networks is the directed graphical model for extremes. Although directed graphical models for extreme values (also known as max-linear Bayesian networks) were only introduced in 2018 by Gissibl and Klüppelberg [100], they have already found use in hydrology [219] and risk modelling [78].

In the last few decades, EVT gained popularity in the financial industry, where it is used to model and manage financial risks associated with extreme events such as stock market crashes, credit defaults and natural disasters [80]. Such events can have significant economic and social impacts. Hence, the application of EVT in finance has become an important tool for financial institutions in risk management and investment decision-making. In fact, EVT has recently become very important in credit risk management [80]. Numerous studies have examined the dramatic fluctuations to which financial markets are vulnerable, primarily as a result of currency crises, stock market crashes and loan defaults [99]. More information about the opportunities that EVT offers for risk management in finance, including the ability to estimate tail risks more accurately, to model dependencies among tail events and to incorporate time-varying risk parameters can be found in Diebold et al. [73].

Nowadays, crises continue to arise at various financial levels, both national and international [132]. Scientists' interest in this topic is growing and the significant contributions to the study of the global economy have increased considerably over the last decade. The study of extreme occurrences, such as

the 2008 financial crisis, is a hot topic among investors [35]. EVT provides methods and distributions that are able to fit such extreme events very accurately [16]. One effective approach to studying financial crises, particularly focusing on how they are transmitted, involves the use of tail dependence. This method examines the extreme co-movements between financial assets or institutions, providing insights into how distress in one entity can significantly impact others during periods of severe market stress.

Tail dependence is a statistical concept that measures the strength of the relationship between extreme events (tail events) in two or more variables. In finance, tail dependence is used to model the interconnectedness and systemic risk of financial networks, such as banks and other financial institutions. By analyzing the tail dependence of different financial variables, analysts can identify potential sources of systemic risk and quantify the potential impact of extreme events on the overall financial system. The notion of tail dependence was first introduced by Sibuya [200] and so far many parametric and nonparametric measures of multivariate tail dependence have been introduced [3, 138, 139, 177, 188].

The use of tail dependence to model financial networks is a relatively new approach in the field of financial economics. It is based on the idea that traditional correlation measures are not sufficient to capture the complex interdependencies between financial institutions in times of stress, and that tail dependence can provide a more accurate picture of the risks involved. Portfolio theory, credit spread analysis and risk management, all rely on estimating the dependency between risky asset returns [183], for which the quantification of high losses in volatile market conditions is crucial. However, measuring the dependency between risky asset returns is not a straightforward task. When extreme realisations matter, the traditional correlation measures can be poor indicators of dependency as the dependence in the tail of the distribution

might not be the same as the dependence in its mid-range. As an alternative, tail dependence coefficients can be used to measure the dependence between two or more random variables in their tails [138, 196].

One of the earliest studies on tail dependence to model financial networks was conducted by Embrechts et al. [80], who proposed the use of copula functions to model the dependence structure between different financial variables. Copulas are statistical tools used for modelling the joint distribution of two or more variables while capturing their dependence structure. The authors showed that copulas could be used to model the tail dependence between financial variables and demonstrated their use in analysing the risk of portfolio losses. Since then, there has been a growing body of literature on tail dependence and systemic risk in financial networks, with a focus on developing new models and techniques to capture the complex and dynamic nature of these networks.

For example, Danielsson et al. [67] studied the tail dependence of the joint distribution of stock returns in the Swedish stock market and found evidence that extreme events were more likely to occur simultaneously across different stocks, suggesting the presence of systemic risk in the market. Gai et al. [92] introduced a method to model tail dependence in financial networks by incorporating 'systemic importance' as a measure that considers the potential impact of institutions on the overall system during a shock. Their findings suggested that systemic importance enhances predictions of contagion spread in a network compared to models that only consider connectivity. Singh et al. [201] analysed the tail dependency structure between the Australian and five other international stock markets. They found that the Australian stock market is more dependent on other major international markets for extreme losses than for extreme profits. Moreover, they discovered that in the case of extremely high gains, the Australian stock market is significantly more

influenced by the stock markets of Japan and Hong Kong than the US market. Similarly, Wang et al. [223] studied the tail dependence structure of the foreign exchange market and discovered that the upper-tail dependence network and the lower-tail dependence network each contain distinctive currency clusters, cliques, and communities. They concluded that market participants should consider various topological elements while deciding on investment or hedging methods, depending on the market environment.

Tail dependence has also been used to analyse propagation of shocks across financial networks. For example, Acemoglu et al. [2] proposed a model to study the propagation of financial contagion in a network of interdependent financial institutions. The authors focused on the role of tail dependence in the transmission of shocks across the network. They found that the presence of tail dependence can significantly amplify the impact of shocks on the network, particularly when the network is highly interconnected. Also, Wen et al. [226] examined the tail dependence relationships between 47 global stock markets. They estimated tail dependence coefficients using the empirical copula method and constructed a financial network, finding that it exhibits a small-world structure with strong tail dependence relationships between some stock markets. They also applied the tail dependence network to portfolio optimization and showed that incorporating tail dependence information can significantly improve the performance of a global portfolio.

Overall, tail dependence has proven to be a valuable tool in modelling financial networks and assessing contagion risk in the financial sector. It allows us to analyse the extreme events and dependencies that occur when multiple assets or institutions experience significant losses or gains simultaneously. Given its success in modelling financial networks, tail dependence is becoming a popular way to assess the interconnectedness of financial systems and identify potential sources of systemic risk. In this thesis, we utilize tail dependence to build our

similarity-based financial network model, focusing on extreme co-movements of stock prices.

2.1.3.2 Direct interaction networks

Direct interaction networks are graphs, in which the links between the nodes represent an interaction between them. In a financial context such interactions may be a transaction between companies, percentage of ownership of one company in another or the amount of debt of one company to another. For example, the authors of [160] studied the corporate ownership network in the automotive industry, constructed using shareholders data on the 30 largest automakers (ranked by production in 2016) and their largest holders (companies, banks, passive investment funds, mutual and pension funds and governments). By studying the network properties they found that most central financial institutions in the corporate ownership network are passive investment funds, highlighting their influential role in the financial system. In addition, Elliott et al. [79] built a financial network using equity claims between banks, while Jackson and Pernoud [117] included both debt and equity. Other examples include customer-supplier networks [105, 191] and transaction networks [115, 135]. However, the lack of publicly available financial data makes it difficult to construct and analyse such interaction networks. This challenge often requires the usage of incomplete datasets, which can limit the comprehensiveness and accuracy of the studies. Nevertheless, the potential for using direct interaction networks in financial modelling is growing with the increasing data availability.

2.1.3.3 The multiplex nature of finance

In finance, multiplex networks have been used to model the interconnectedness of financial institutions and the risks associated with their interactions. The applications of multiplex networks in finance include portfolio optimization, risk management and contagion modelling. For example, the authors of [209] presented a model for assessing the degree of diversification and systemic risk in interbank networks. The model incorporates both the size and concentration of a bank's lending portfolio, as well as the connectivity of the interbank network. The authors demonstrated how the model can be used to analyse the effects of different shock scenarios on the stability of the financial system. Also, Del et al. [70] constructed a multiplex network where the nodes represent countries, the edges represent cross-country financial assets of different types, and layers represent asset types. They then simulated different scenarios of financial shocks in the network and measured the extent and speed of contagion across the network. The authors found that the multiplex network structure can both amplify and dampen the effects of financial shocks. In particular, they showed that the presence of different layers of relationships can create pathways for contagion to spread quickly, but can also provide alternative routes for the flow of funds, which can mitigate the impact of shocks. Also, the authors of [204] constructed a multiplex network that combines the global trade network and the global investment network, and used network analysis tools to investigate how the failure of a country or a company in one network affects the other network and the global economy as a whole.

Numerous studies use various methods and approaches to compare multiplex and monoplex (single-layered) network models in the context of financial contagion, including simulations, network metrics, and empirical analysis of real-world financial networks. In [92] the authors investigated the dynamics

of financial contagion in interbank lending networks using a model that incorporates both random and targeted shocks. They compared the impact of these shocks in monoplex and multiplex network models and found that the multiplex model can better capture the complexity of the interbank lending structure and the potential for contagion to spread across different financial markets. Also, Poledna et al. [182] examined the nature of systemic risk in financial systems. The authors argued that financial systems are multilayered networks, with layers representing different types of financial institutions and interconnections between them. They showed that systemic risk is substantially underestimated when computed on single layers only, as is the typical approach. In addition, Bargigli et al. [24] used data of supervisory reports on Italian banks to build multiplex networks. They investigated the structure and properties of interbank networks, which represent the connections between banks through lending and borrowing relationships. The authors argued that interbank networks are not simple, one-dimensional structures, but rather multiplex networks that consist of multiple layers of interbank lending and borrowing relationships. They found that monoplex approaches poorly depict interlinkages and may cause biased systemic risk estimation, with the multiplex structure having important implications for stability. Bardoscia et al. [23] compared the use of monoplex networks (which consist of a single layer of interbank lending and borrowing relationships) to the use of multiplex networks (which consist of multiple layers of interbank lending and borrowing relationships) in modelling financial instability and systemic risk. According to the authors, employing multiplex networks offers a more accurate and detailed understanding of the intricate connections and interrelationships among financial institutions and markets. The authors of [125] also compared multiplex and monoplex networks by studying the dependency relations among international stock market indices. They concluded that the multiplex network approach is better suited for studying contagion and systemic risk in

stock markets due to their ability to capture complex interactions.

All these studies demonstrate that multiplex network models offer a more comprehensive and accurate depiction of systemic risk and contagion dynamics compared to traditional monoplex (single-layered) models. Hence, it is crucial to consider the different connections between financial institutions, as each connection represents a potential pathway for contagion transmission. By modelling financial assets as nodes and representing various types of relationships, such as correlations or co-movements, through distinct layers, multiplex networks provide a more comprehensive and realistic representation of financial markets than traditional approaches. In addition, analysis of the interconnectivity between different financial institutions can be used to identify potential systemic risks and vulnerabilities in the financial system, enabling regulators and policymakers to implement more effective risk management strategies and prevent financial crises. In this thesis, we build a multiplex network model to leverage these advantages. Specifically, we study the spread of financial contagion within a multiplex network, allowing for a more accurate assessment of systemic risk.

2.2 Epidemics on networks

The increasing use of methods from network science has resulted in a new modelling paradigm that combines multiple fields of study, including physics, mathematics, biology, and social sciences. Infectious illness transmission between nodes in a network has become an important area of research in recent years, with many studies exploring the dynamics of disease transmission and the impact of network structure on epidemic control. In this section we explore epidemics on networks and demonstrate their usefulness for modelling

financial contagion on networks.

2.2.1 Background and history

Infectious diseases create a significant social burden. For example, globally, the COVID-19 pandemic has, as of July 2024, claimed more than 7 million lives. In addition to the human cost, the financial impact of the COVID-19 pandemic has been huge [156]. For instance, the Congressional Budget Office predicts a total of $7.6 trillion in lost output in the US during the next decade [66]. New diseases emerge on a regular basis, while old ailments remain. Mathematical models have been used to guide policymakers in the response to the emergence of diseases such as SARS [166], influenza [170] and Ebola [36]. Epidemic propagation has been mathematically modelled for more than a century. The majority of epidemic models are based on a compartmentalisation of individuals or hosts dependent on their infection status [124]. The two most widely used compartmental models for the spread of infectious diseases in populations, the SIR (Susceptible - Infected - Recovered) and the SIS (Susceptible - Infected - Susceptible) epidemic models, were introduced by Kermack and McKendrick in the 1920s [126]. The SIR model assumes that once an individual has recovered it cannot get reinfected again, while the SIS model assumes that once an individual recovers it is susceptible to the disease again. Hence, the SIR model is commonly used to describe infectious diseases that lead to a longlife immunity, such as measles [145]. The SIS model, on the other hand, is mostly used for sexually transmitted diseases (STDs), as individuals who recover can often become susceptible to reinfection [91].

The work of Kermack and McKendrick resulted in two standard sets of differential equations, which serve as a foundation of almost all of mathematical

(a)

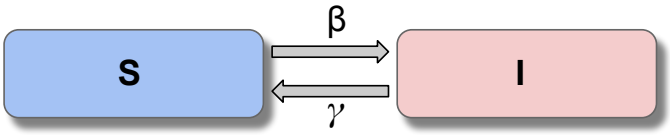

(b)

Figure 2.6: Schematic representation of the a) SIR and b) SIS epidemic models. Boxes represent different compartments and the arrows indicate transitions between compartments, based on their respective rates.

epidemiology [124]: the SIR model

$$\begin{aligned} \frac{\mathrm{d}S}{\mathrm{d}t} &= -\beta\frac{SI}{N}, \\ \frac{\mathrm{d}I}{\mathrm{d}t} &= \beta\frac{SI}{N} - \gamma I, \\ \frac{\mathrm{d}R}{\mathrm{d}t} &= \gamma I, \end{aligned} \tag{2.1}$$

and the SIS model

$$\begin{aligned} \frac{\mathrm{d}S}{\mathrm{d}t} &= -\beta\frac{SI}{N} + \gamma I, \\ \frac{\mathrm{d}I}{\mathrm{d}t} &= \beta\frac{SI}{N} - \gamma I. \end{aligned} \tag{2.2}$$

In the above equations $S = S(t)$, $I = I(t)$ and $R = R(t)$ are the number of susceptible, infected and recovered individuals at time t, respectively, and N is the population size, which is assumed to be constant over time, so no new individuals are 'born' and no individuals 'die'. Here, β is the rate at which susceptible individuals become infected and γ is the rate at which each infected individual recovers (in the SIR model) or becomes susceptible again (in the SIS model) (see Figure 2.6). These models are based on the assumption that each individual has an equal chance of coming into contact with any other individual, also called a homogeneous-mixing population assumption. In

practice individuals have different numbers of contacts, usually considerably smaller than the population size, making the homogeneous mixing population assumption unrealistic. Models which incorporate network structure avoid the homogeneous mixing assumption and enable the calculation of the spread of an epidemic across the entire population based on the infection patterns at the individual level [124]. The concept of social networks and their impact on disease transmission began to emerge in the late 20th century, with studies such as the work of Granovetter [104] on the strength of weak ties in social networks. Nonetheless, it was not until the late 1990s that the study of epidemics on networks began to gain momentum with the advent of modern computing. In the following section, we provide a comprehensive review of various applications of epidemic modelling on networks.

2.2.2 Epidemics on networks: Applications

Directly transmitted infectious disease epidemiology and networks are fundamentally linked. The use of networks to model epidemics allows for realistic infection contact patterns to be incorporated. For example, real-world data can be used to parametrise network models that include high variability in the number of contacts that an individual has, the higher probability that individuals connect with individuals with similar properties or the tendency of already connected individuals to share common contacts. The comprehensive and objective examination of the transmission mechanism, as well as real contact network data, support the use of networks in disease modelling [128].

One influential study in this field is the work of Watts and Strogatz [225], who demonstrated that diseases can spread quickly through small-world networks, even if the network has only a small number of nodes with large degree.

Another important development in the study of epidemics on networks was the introduction of the network epidemic model (NE model) by Pastor-Satorras and Vespignani [175]. This model extends the classic SIR model to include the effect of network structure on disease transmission. In the NE model, individuals are represented as nodes in a network, and the edges between nodes represent social connections. The probability of an individual becoming infected depends on the infection rate of the disease and the number of infectious neighbors they have. The probability of an infected individual recovering depends on the recovery rate of the disease. The NE model also allows for the study of different types of network structures. Through simulations, Pastor-Satorras and Vespignani found that the NE model can be used to predict the spread of epidemics in a variety of network structures, and can help identify the most effective strategies for controlling the spread of infectious diseases.

Many other studies have also explored the impact of network structure on disease transmission and the development of new epidemic models and control strategies. For example, Newman [162] showed that scale-free networks, which have a few highly connected nodes and many weakly connected nodes, are particularly vulnerable to the spread of epidemics. In addition, Berger et al. [37] developed a novel epidemiology-based modelling framework to demonstrate how network structure influences the spread of quality issues in supply networks. By applying the model to 21 real-world networks, they found that the magnitude of quality disruptions is strongly influenced by the origin node and network topology. Also, Hiram et al. [109] discussed graph-based epidemiological models and showed how their use may significantly improve the disease spreading control. The importance of studying epidemics on networks has been underscored by the recent global health crisis brought about by the COVID-19 pandemic, which highlighted the complex interplay of human interactions and mobility patterns in the spread of pathogens. Researchers have

used network models to study the impact of different interventions, such as social distancing and vaccination, on disease transmission and to identify the most effective strategies for controlling the spread of the virus [142, 180, 235].

Epidemic models are commonly used to study the spread of infectious diseases in populations, but as indicated in [46], they offer a wide range of applications, such as modelling social networks [55, 140], computer networks [143, 189] and transportation networks [140]. In the next section we discuss the applications of epidemic modelling on networks in the field of finance.

2.2.3 Financial epidemics on networks

The term 'financial contagion' first appeared in 1997 during the Asian crisis but the first instances of the phenomenon must have occurred much earlier in financial history [60]. Financial contagion refers to a shock (an unexpected event that significantly disrupts financial markets, leading to widespread economic impact, such as a group of large companies suddenly unable to pay back their loans) that begins with a small number of financial institutions and spreads to the entire financial system, often impacting economies across multiple countries. The Asian crisis of 1997, the Russian default of 1998, and the global financial crisis of 2008 are among the recent events that are thought to be results of contagion spread [1].

Epidemic modelling has become an increasingly popular tool to study the spread of financial contagion, particularly since the global financial crisis of 2008. The spread of financial contagion can be seen as similar to the spread of an infectious disease, with the nodes of the network representing financial institutions and the edges representing their interconnections [149]. In this literature review, we examine some of the key studies that have explored this

topic, highlighting the main findings and contributions of each.

One of the early studies in this area was conducted by Kleindorfer et al. [130] who examined the interconnectedness of modern businesses and economies and its impact on strategy, risk management and profitability. The book discusses the use of epidemic models to study financial risk contagion in interconnected networks, exploring network structure, contagion mechanisms and policy interventions to mitigate contagion. Past financial crises are analysed, including the 2008 global financial crisis, for which the authors compared the spread of the SARS virus to the risk spread caused by the Lehman Brothers' bankruptcy and concluded that the infection conditions of both were compatible with virus transmission characteristics in their respective networks. In another study, Garas et al. [93] used an SIR epidemic model to study the spread of financial crisis among different countries, where susceptible, infected and recovered individuals represented the healthy, financially distressed (lacking the necessary income or revenue to fulfill their financial commitments and obligations) and recovered economies, respectively. The authors validated their model using data from the 2008 global financial crisis and demonstrated that it can accurately capture the propagation of financial distress across countries. Similarly, Huaihu et al. [113] proposed a complex network model to study banking crisis contagion. The authors used an edge-dynamic weighted network to represent interbank lending relationships and applied an SIR epidemic model to simulate contagion. The model is validated using data from the Chinese banking system. The study suggests that early intervention and coordinated policy responses are important for preventing the spread of banking crises. Toivanen [216] also proposed an SIR epidemic model to study financial contagion in the interbank network. The model is validated using Finnish interbank network data and allows for policy intervention analysis. Also, Yuanyuan et al. [232] studied crisis spreading in shareholding networks of businesses and

their primary holders using data on ownership structure in Chinese stock markets. They created a correlation network based on stock price movements and simulate crisis contagion using an SIR epidemic model. The characteristics of shareholding networks are studied and the results show that if there are some failures or attacks on highly connected vertices, the crisis will spread quickly and create catastrophic harm. Demiris et al. [71] used an SIR model to study the dynamics of financial crises in a network of countries, adapting the epidemiological framework to model the spread of financial distress. The method is illustrated on a number of currency crises. The empirical findings point to a rising trend in global currency crisis transmission over time. Also, the authors of [88] proposed a time delayed SIR model to study the credit risk contagion in stock markets, where the time delay is the amount of time that a bank has before it becomes defaultable. The model has been used to find the steady states according to different bank support policies.

These studies highlight the usefulness of epidemic modelling for studying the dynamics of various types of financial contagion. In general, epidemic modelling is a popular tool for describing financial contagion because both phenomena share some common characteristics, such as the spread of a disturbance through a network of interconnected units. While in epidemiology the disturbance is a disease, in finance it usually is a financial crisis. In both cases, the spread can be influenced by the network topology, the properties of the units, as well as the interactions between them. The SIR model, in particular, is chosen in this thesis for its effectiveness in capturing the essential dynamics of financial contagion. This model allows us to simulate how financial distress propagates through interconnected financial institutions and to assess the impact of various network structures and interaction patterns.

2.3 Discussion and conclusion

This literature review provides a critical analysis of existing research, focusing on the two primary research topics used in this thesis: network science and epidemic modelling on networks. Insights from the topics discussed will be utilized throughout the thesis to support the primary goal: modelling financial crisis propagation using multiplex networks. Network science will form the foundation of much of the research, in which we employ multiplex networks and epidemic models on networks to study contagion in complex financial systems.

In summary, modelling financial contagion on networks is a complex problem that requires interdisciplinary expertise and advanced analytical techniques. The application of network theory to the analysis of financial systems has proven highly effective and holds great potential for the future. While significant progress has been made in understanding how financial shocks propagate through networks, the increasing availability and resolution of data suggest that this research field is still in its early stages. Moreover, the network approach offers valuable insights into practical issues like portfolio optimization, asset allocation, risk diversification and systemic risk [50]. Ultimately, improving our understanding of financial contagion on networks can help us to better prevent and mitigate the effects of future financial crises.

In this thesis we address several research gaps that are evident from the literature. Firstly, there has been limited research on studying the combination of multiplex network structure and EVT methods. While multiplex networks provide a richer representation of financial systems and EVT methods focus on modelling the impact of extreme events, the intersection of these two approaches has remained limited. We address this gap by using the notion of tail dependence for building one out of four layers in our proposed multiplex

financial network. Secondly, the majority of studies do not consider the 'pure contagion', which arises from global effects, rather than from specific, identifiable channels such as direct financial relationships [60]. In our network we introduce a 'global layer', incorporating the 'pure contagion' assumption. Thirdly, most of the literature fails to address the significance of exogenous shocks (external events like natural disasters or geopolitical crises) in financial contagion spread. This thesis demonstrates the importance of detecting such shocks and incorporating them into the model for more accurate predictions. In addition, unlike most static models in the literature, our model is designed to be updated in real time. Moreover, the framework is universal in nature and could be applied to analyse any contagious financial crisis, whether it involves networks of companies, countries or economic sectors. Finally, the majority of studies on econometric and empirical finance models rely primarily on maximum likelihood methods for estimation [206]. Traditional maximum likelihood approaches assume that all sample observations are weighted equally. However, we show that improvements can be made by introducing a weighting scheme, which accounts for the importance of each observation.

Chapter 3

Modelling financial crisis propagation as an epidemic on multiplex networks

This chapter proposes a novel framework for modelling the spread of financial crises in complex networks, combining financial data, extreme value theory and an epidemiological transmission model. We accommodate two key aspects of contagion modelling: fundamentals-based contagion, where the transmission is due to direct financial linkages, and pure contagion, where a crisis might trigger additional crises due to global effects. We use stock price, geographical location and economic sector data for a set of 398 companies to construct multiplex networks of four layers, on which a SIR transmission model is defined, in order to model the spread of financial shocks between companies by accounting for their interconnected nature. By utilizing stock price data for the 2008 and 2020 financial crises, we investigate and assess the effectiveness of our model in forecasting the propagation of financial shocks through the network, where a shock is detected by measuring stock price volatility.

The results suggest that the proposed framework is effective in predicting the spread of financial crises. Our findings demonstrate the significance of each layer of the multiplex network structure, which differentiates between various transmission pathways, for predicting the number of affected companies, as well as for company-, sector- or location-specific predictions.

3.1 Introduction

In this chapter we propose a novel framework for modelling financial contagion that is based on an SIR epidemic model defined on a multiplex network constructed from financial data. We employ a stochastic epidemic transmission mechanism in which financial crises can spread locally (to network neighbours) as well as globally (to any company). Then, by considering their local and global connectivity, we simulate how a financial shock spreads from the original infected companies to the others. To demonstrate our approach we construct two multiplex networks, representing the financial dependence of 398 companies in the 2008 and the 2020 financial crises, where each node represents a company and each layer represents a different type of connection between the companies. Both networks consist of four layers: a tail dependence network layer, a continents layer, a sectors layer and a global layer. The tail dependence layer measures the strength of dependence between two companies using tail dependence coefficients, which are calculated using daily stock price data. This weighted (complete) network is filtered via the planar maximally filtered graph (PMFG) method [220], to remove weak and potentially spurious links. The continents and sectors layers, respectively, connect companies under the assumption that companies in the same continent or sector are more likely to be affected by a financial crisis simultaneously. Finally, the global layer is a complete network, in which each company is connected to

every other company. This layer corresponds to the 'pure contagion' assumption that a crisis in any company may trigger a crisis in any other company. In addition, in our model, a company may experience a financial shock not just as a result of direct linkages to the initially infected company, but also as a result of indirect connections within the network of companies, amplifying the spread and impact of the financial shock. As a result, we allow for the so-called 'cascading effect', a phenomenon where the impacts of a financial crisis spread and intensify through interconnected channels, resulting in a broader and more severe contagion than initially anticipated [102], which is commonly overlooked in the literature.

We apply the model to two recent financial crises, the 2008 financial crisis and the 2020 financial crisis, and investigate and evaluate its utility in predicting the spread of financial shock across the network. We first identify which companies have been 'infected' in each of the two crises using stock price volatility. We then study how using the previous n crisis days in each of the two crises can be used to predict the infections in the future k days for different combinations of n and k. The results suggest that for each crisis a different combination of n and k give the most accurate predictions. For both crises the proposed model outperforms the homogeneous mixing population approach in predicting the number of infected companies, the continents and economic sectors that will be most affected, and the sets of specific companies that will be infected during the future crisis days.

The rest of the chapter is organized as follows. Section 3.2 describes the data set. Then, in Section 3.3 we describe the two parts of the modelling framework: the multiplex network construction procedure and the transmission mechanism. In Section 3.4 we apply the model to the 2008 and the 2020 financial crises. We first define the concept of infection in a financial context and then we study how the model can be used to predict future infections in each

of the two crises, using past data. Finally, we assess the significance of each layer within our multiplex network by conducting a comparative analysis of its predictive accuracy on omission of various subsets of its layers. Section 3.5 concludes and discusses our findings, the limitations of our approach as well as avenues for further research.

3.2 Data

The analysis in this paper is based on the closing daily stock price of 398 companies from 17/01/2002 to 18/07/2022 (inclusive), representing $n = 5229$ trading days. The data are collected from https://finance.yahoo.com/ and the companies are selected such that for each company there are consistent data going back as far as 17/01/2002, covering a sufficient time period before the 2008 financial crisis. We separate the companies into groups, based on the Bureau van Dijk[1] company database. Firstly, we group the companies according to the geographical location of their headquarters, resulting in six groups: Africa (2), Asia (77), Europe (115), North America (194), Oceania (9) and South America (1). The numbers in parentheses indicate the total number of companies in each respective continent. The disparity in geographical representation arises from the distribution of available data meeting our date span criteria. North America, for example, has a highly developed and mature financial market and hosts numerous publicly traded companies, many of which have extensive historical data available. This makes it easier to find companies with consistent data spanning back to 2002. In contrast, some regions, especially emerging markets in Africa or South America, may have fewer publicly traded companies or less robust historical financial data, making it more challenging to include a comparable number of companies from

[1] A significant business information publisher, Bureau van Dijk specialises in private corporate data together with software for searching and analysing businesses.

those regions in the dataset. Secondly, we separate the companies into 13 groups based on their primary economic sector, as defined by the Bureau van Dijk dataset: Finance (47), Oil and gas industry (36), Pharmaceutical industry (36), Automotive industry (35), Airline industry (17), Food industry (23), Mining activities (20), Electricity (17), Software industry (38), Electronics (58), Telecommunications (10), Chemicals (8) and Others (53), with the numbers representing the number of companies in each sector.

The stock price returns for each company i at day t for $2 \leq t \leq 5229$ are calculated by taking the logarithmic difference of successive closing prices as follows:

$$x_{i,t} = \ln(p_{i,t}) - \ln(p_{i,t-1}), \tag{3.1}$$

where $p_{i,t}$ denotes the closing stock price of company i at day t for $1 \leq t \leq 5229$.

Tables 3.1 and 3.2 show the characteristics of the studied data for the periods prior to, and after, the 2008 financial crisis. We calculate the minimum, maximum, mean and standard deviation of each company's stock returns and average them by continent and sector to highlight the characteristics of the studied data. We can see that in all the cases the log-returns have mean close to 0, with potentially asymmetric distributions. The asymmetry becomes particularly noticeable when examining the minimum and maximum values within each dataset. For instance, in certain sectors or geographic regions, the minimum and maximum log-return values differ notably in absolute terms, indicating potential skewness or asymmetrical behaviour in the return distribution. Positive skewness indicates that the distribution has a longer tail on the positive side, implying that extreme positive values occur more frequently than extreme negative values, which in financial contexts is generally associated with the potential for large gains. Negative skewness, on the

other hand, refers to a distribution skewed towards the left and has a longer tail on the negative side, indicating a higher frequency of extreme negative values, highlighting the possibility of more frequent and severe losses.

	Before the 2008 crisis				After the 2008 crisis			
	Mean	Min	Max	Sd	Mean	Min	Max	Sd
Finance	1.97e-04	-0.2555	0.1836	0.0231	-2.79e-05	-0.3288	0.3219	0.0256
Oil and gas industry	3.71e-04	-0.4108	0.3361	0.0336	2.67e-05	-0.3476	0.3705	0.0293
Pharmaceutical industry	-3.12e-05	-0.3289	0.3053	0.0298	1.98e-04	-0.3336	0.3831	0.0274
Automotive industry	2.26e-04	-0.2357	0.1864	0.0271	2.16e-04	-0.2454	0.2878	0.0269
Airline industry	1.98e-04	-0.4155	0.2829	0.0286	7.27e-05	-0.3342	0.2418	0.0261
Food industry	1.74e-05	-0.2281	0.1281	0.0212	1.85e-04	-0.2472	0.1994	0.0216
Mining activities	4.25e-05	-0.5125	0.1565	0.0309	-1.11e-04	-0.3303	0.2526	0.0310
Electricity	3.43e-04	-0.2966	0.2475	0.0241	4.94e-05	-0.2076	0.1563	0.0191
Software industry	-3.55e-04	-0.4669	0.2713	0.0402	4.85e-04	-0.2231	0.2037	0.0229
Electronics industry	-4.64e-04	-0.4841	0.2382	0.0368	2.41e-04	-0.3531	0.4159	0.0311
Telecommunications industry	1.85e-04	-0.1610	0.1599	0.0265	1.85e-04	-0.3495	0.3160	0.0232
Chemicals	-6.83e-05	-0.1939	0.1310	0.0205	2.01e-04	-0.2450	0.2262	0.0225
Others	5.61e-05	-0.2859	0.1625	0.0267	3.23e-04	-0.3682	0.4089	0.0257

Table 3.1: Descriptive statistics of stock returns, with the minimum, maximum, mean and standard deviation calculated for each company's stock returns. The columns display these values, averaged by sector. The statistics are provided for two distinct periods: before and after the 2008 financial crisis.

	Before the 2008 crisis				After the 2008 crisis			
	Mean	Min	Max	Sd	Mean	Min	Max	Sd
Asia	1.11e-04	-0.3523	0.3134	0.0353	9.77e-05	-0.2815	0.3356	0.0273
Europe	1.53e-05	-0.2607	0.2052	0.0208	-3.94e-05	-0.3299	0.3203	0.0262
North America	1.19e-04	-0.4107	0.2044	0.0307	3.69e-04	-0.3245	0.3291	0.0265
Others	3.96e-04	-0.3439	0.2459	0.0254	-1.58e-04	-0.2466	0.2080	0.0270

Table 3.2: Descriptive statistics of stock returns, with the minimum, maximum, mean and standard deviation calculated for each company's stock returns. The columns display these values, averaged by continent. The statistics are provided for two distinct periods: before and after the 2008 financial crisis.

Our objective is to model financial contagion using two different data sets: daily stock prices from the period 17/01/2002 to 30/06/2007, to model contagion during the 2008 financial crisis, and from 17/01/2002 to 29/02/2020 to do so in the 2020 crisis following the onset of the COVID-19 pandemic. We emphasise that here, and in all subsequent occurrences, datasets defined over stated date ranges are understood to be inclusive of start and end dates. The following section gives a detailed explanation of the network construction approach used in our analyses.

3.3 Model formulation

Our primary focus is on modelling the process of financial contagion that can spread to a specific population of companies. Specifically, we aim to predict which companies, financial sectors, and continents are most likely to be affected in the future, using information based on the past days of the crisis. Our modelling framework comprises two main parts. We first build a multiplex financial network, where the nodes correspond to companies and the edges in each layer represent different types of connections between companies. By incorporating multiple network layers, we can capture the various ways in which financial contagion may spread between companies. We then employ an SIR epidemic model on each network layer. The model's key parameters are the transmission probabilities (i.e., the probability of an infected node transmitting the infection to a susceptible node on a given day) and the recovery probabilities (i.e., the probability that an infected node becomes recovered on a given day), which we estimate using a maximum likelihood approach by fitting the model to past crisis data. Then, by simulating the spread of financial contagion using the SIR model with the estimated parameters, we can identify the companies, financial sectors, and continents that are predicted to be most vulnerable to future contagion events.

3.3.1 Network model

We construct a multiplex network (see Figure 3.1), where each node represents a company and each layer represents a different type of connection between the companies. We construct four layers: a tail dependence network layer, a continents layer, a sectors layer and a global layer. The motivation for, and method for construction of, these networks are detailed in the following

subsections.

3.3.1.1 Tail dependence network layer

The relationship between tail dependence and the propagation of financial crisis risk is highlighted by a number of studies [1, 5, 215]. Tail dependence is used to study the likelihood of joint tail events, where the occurrence of extreme movements in one asset's return is associated with a higher likelihood of extreme movements in another. This phenomenon reflects the interconnectedness of financial markets, whereby shocks or disruptions in one asset class or market segment can trigger correlated movements in other assets. The tail dependence coefficient is a common measure of financial dependence between two companies. For example, the concept of marginal expected shortfall (MES), a widely recognized risk measure that evaluates the potential losses of a company given that another experiences an extreme loss, is intricately linked to tail dependence coefficients, thereby underscoring the relevance of tail dependence in capturing the tail behavior of financial assets [48]. To study how likely it is that two companies experience extreme losses together we construct complex financial networks, via the following two-step process. Firstly, we calculate the tail dependence strength between each pair of companies' stock returns. Secondly, we filter the edge information required for network building using the PMFG approach.

3.3.1.1.1 Tail dependence estimation

Let $\{(-x_{i,t}, -x_{j,t}) : t = 1, 2, \ldots N\}$ be the realisations of the bivariate negative stock return (X_i, X_j), where $x_{i,t}$ is as defined in (3.1). We assume throughout that X_i and X_j have continuous distribution functions. For each pair of

negative stock returns (X_i, X_j) of companies i and j, the marginal aspects of the joint distribution can be removed by transforming the bivariate negative returns into unit Fréchet marginals (S_i, S_j) by using the following transformation:

$$S_i = -1/\ln F_i(X_i) \text{ and } S_j = -1/\ln F_j(X_j), \tag{3.2}$$

where F_i and F_j are the marginal distribution functions of X_i and X_j, respectively. In practice, the functions F_i and F_j used in (3.2) are estimated by the empirical marginal distribution functions of the two random variables. This transformation does not affect the dependence structure of the bivariate joint distribution, so (S_i, S_j) possesses the same dependence structure as (X_i, X_j).

Since we are interested in the probability that one company experiences an extreme financial loss, given an extreme loss in another (the likelihood of crisis transmission), for each pair (S_i, S_j) we estimate the upper tail dependence coefficient (upper TDC) $\chi^U_{i,j}$, defined as

$$\chi^U_{i,j} = \lim_{q \to 1^-} \mathrm{P}(F_j(S_j) > q \mid F_i(S_i) > q).$$

Hence the upper TDC corresponds to the likelihood that one margin will surpass a high threshold if the other margin also exceeds this threshold. The coefficient $\chi^U_{i,j}$ takes values in the range $[0, 1]$, describing the strength of the tail dependence between S_i and S_j: $\chi^U_{i,j} = 0$ means that the two variables S_i and S_j are upper tail independent and $\chi^U_{i,j} > 0$ indicates upper tail dependence.

The TDC can also be defined using the concept of a copula, introduced by Sklar [202]. A fundamental result shown in [202] states that $F_{i,j}$, the joint distribution function of (X_i, X_j), can be represented as $F_{i,j}(s_i, s_j) = C_{i,j}(F_i(s_i), F_j(s_j))$, where $C_{i,j}$ is a copula function (a bivariate distribution

function with uniform margins). Then, as shown in [61],

$$\chi^U_{i,j} = \lim_{q \to 1^-} \frac{1 - 2q + C_{i,j}(q,q)}{1-q}.$$

In practice we estimate the strength of tail dependence for each pair (S_i, S_j) and threshold $q \in (0,1)$ as follows:

$$\hat{\chi}^U_{i,j} = \hat{\chi}^U_{i,j}(q) = \frac{1 - 2q + \hat{C}_{i,j}(q,q)}{1-q}, \tag{3.3}$$

where $\hat{C}_{i,j}$, the empirical counterpart of $C_{i,j}$, is computed via

$$\hat{C}_{i,j}(u,v) = \frac{1}{N} \sum_{n=1}^{N} \mathbb{1}\left(r^n_i \leq N - \lfloor N(1-u) \rfloor, r^n_j \leq N - \lfloor N(1-v) \rfloor\right).$$

Here, r^n_i and r^n_j are the ranks (the index of the element in an ascending list) of the n^{th} observations of S_i and S_j, respectively. Note that the transformation in (3.2) is monotonically increasing, so that the rank of an observation from S_i is the same as that for the corresponding X_i.

The analysis in the remainder of the chapter is based on the estimated upper tail dependence coefficients $\hat{\chi}^U_{i,j}(0.95)$ which use a threshold of $q = 0.95$. This choice is consistent with the existing literature using tail dependence to build financial networks [54, 137, 201]. Moreover, in this chapter we construct separate networks employing all the data in our set prior to the 2008 and to the 2020 crises, respectively. We use the values of the upper TDC between each pair of companies i and j to measure the strength of dependence between the companies in our dataset, and are key to our construction of complex financial networks and our SIR model for financial contagion: the higher the upper TDC between two companies, the higher the probability of crisis transmission.

3.3.1.1.2 Planar maximally filtered graph

The PMFG method was first introduced in [220]. The primary goal is to filter complex networks by retaining only the most important links that do not break planarity [168]. By doing so PMFGs can assist with eliminating spurious (weak) connections, thereby emphasizing topological properties such as communities and easing computational burden. Moreover, planarity ensures easier network visualisation. PMFGs are maximally filtered, which means that they are constructed in such a way that the number of connections between nodes is maximized while still maintaining planarity. This property ensures that the PMFG captures the most important relationships between the nodes, making it a useful tool for analyzing complex networks [148].

PMFGs constructed from financial datasets have been used to detect fundamental market changes and community structures [157], to study the spread of financial risk [185] and to analyse financial networks describing correlations (or other dependencies) between financial assets [18, 90, 220]. In addition, the authors of [203] show that PMFGs can be used to reduce the complexity and dimensionality of financial networks, while keeping the clustering structure. Prior to the study of [203], the two most popular tools for filtering the edge information in complex financial networks were the minimum spanning tree (MST) algorithm [18, 146] and the correlation coefficient threshold method [41]. However, the latter is extremely dependent on the threshold decision [229]; for the former, the key advantage of the PMFG algorithm is that it preserves more information: the MST has $n-1$ edges, while the PMFG has $3(n-2)$ edges (compared to $n(n-1)/2$ of the complete network with n nodes). Furthermore, the PMFG always contains the MST, so it is a connected network.

To construct the two PMFG financial networks we employ the procedure

described in Algorithm 1. For clarity, we specify that this method is applied only to the tail dependence layer.

Algorithm 1 Construct PMFG layer

1. Start with the graph $G = (V, E)$, where V is the set of all companies, represented by the nodes and $E = \emptyset$.
2. For each pair S_i and S_j $(i < j)$ calculate $\hat{\chi}^U_{i,j}(0.95)$ via equation (3.3) and sort those values in a list L from greatest to smallest.
3. Add edge (i, j) to E if $\hat{\chi}^U_{i,j}$ is the first element of L and the network G is still planar after adding the edge.
4. Remove $\hat{\chi}^U_{i,j}$ from L and go back to Step 3.
5. Continue the procedure until the graph has $3 \times (n - 2)$ edges.

3.3.1.2 Additional layers

In addition to the tail dependence network (hereafter denoted the PMFG for brevity) layer, we include layers to incorporate other known relations between the companies and describe other possible crisis transmission channels.

3.3.1.2.1 Sector and continents layers

The 2008 financial crisis demonstrated the importance of interconnectedness as it quickly spread from the subprime mortgage market in the United States to the wider financial sector, causing significant losses for institutions, leading to a decline in consumer spending and demand for goods and services [211]. On the other hand, sectors such as healthcare and technology did relatively well [227]. In addition, the 2008 global financial crisis provides an example of how different continents can be affected by a crisis at different time periods [136]. The crisis originated in the United States with the collapse of the subprime mortgage market, which led to a wave of defaults and foreclosures.

After that, the crisis expanded to other regions of the world, such as Europe, Asia and South America, although the timing and intensity of its effects differed between continents. Ireland and Spain, two countries in Europe that had recently had huge property market booms and had significant exposure to the US subprime mortgage sector, were among the first to be affected by the crisis. As the crisis expanded, countries with close ties to the financial industry, such as the United Kingdom and Switzerland, were also affected [212]. Some Asian countries also experienced a slowdown in economic development and a fall in exports but Asia was generally protected from the worst of the crisis because of its comparatively strong banking sector [39]. In Latin America the crisis led to a decline in commodity prices and a decrease in demand for exports, which affected countries such as Brazil and Mexico [114]. The financial crisis of 2008 hence showed how various economic sectors and continents might experience a crisis at various times, based on their susceptibility to the underlying causes of the crisis and their resilience to external shocks.

Furthermore, in the early stages of the 2020 financial crisis, initiated by the COVID-19 pandemic, the healthcare sector was the most directly affected, as it had to respond rapidly to the surge in demand for medical services and supplies. At the same time, the travel and tourism industry was severely impacted, as travel restrictions and lockdown measures were put in place, which caused a sharp decline in demand for air travel, hotels, and other related services. As the pandemic continued to spread and the global economy went into recession, other sectors also began to feel the impact [171]. For example, the closure of non-essential businesses and social distance initiatives had a substantial negative impact on the retail and restaurant industries, resulting in significant decreases in employment and sales. Also, disruptions to global supply networks and a decline in consumer demand have had an impact on manufacturing [15]. The pandemic struck Asia early and hard; in an effort

to stop the virus's spread, China experienced massive lockdowns and travel restrictions. Several Asian nations, including South Korea and Japan, were also affected early and took severe precautions to limit the epidemic [58]. In Europe, the pandemic hit in late February and early March of 2020, with an especially bad outbreak in Italy that overwhelmed its healthcare system. Other European countries, including Spain, France, and the United Kingdom, also experienced high levels of transmission and implemented strict lockdown measures to control the spread of the virus [181]. The epidemic hit the Americas later than it did Asia and Europe, with a spike in cases and fatalities in the United States in the spring and summer of 2020.

To account for these features, we add undirected 'sectors' and 'continents' layers in which companies are connected if they are in the same sector or continent, respectively. Hence, each connected component in the sector and continents layers is a complete network.

3.3.1.2.2 Global layer

In addition to the 'fundamentals-based contagion' embedded in the above network layers, we additionally allow for 'pure contagion', whereby crises may spread due to global effects that are not explicitly accounted for within the connectivity so far defined. Hence, we add a complete network, denoted the 'global layer'. Figure 3.1 presents a schematic representation of our multiplex network.

3.3.2 Contagion Model

We employ a discrete-time SIR epidemic model defined on the network of n companies to simulate financial crisis propagation. At each time step a

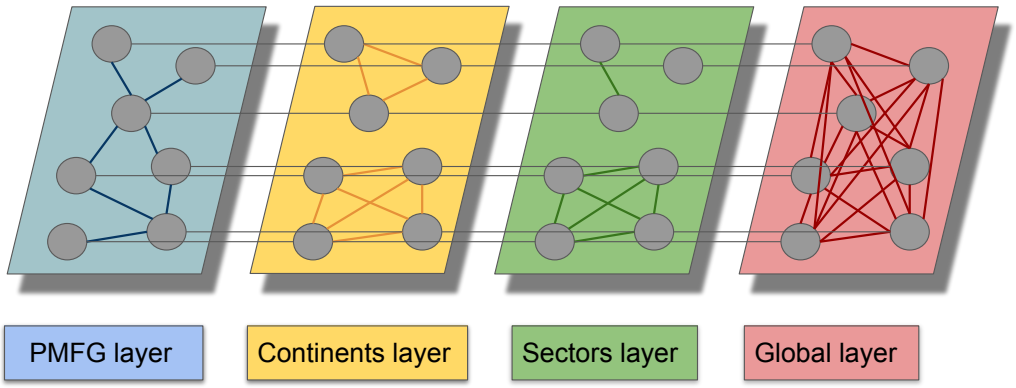

Figure 3.1: Schematic diagram of the 4-layer multiplex network described in Section 3.3.1, comprising a weighted PMFG layer and binary continents, sectors and global layers. The inter-layer links are only between the same node, as well as between each pair of layers.

company is either susceptible (S), infected (I) or recovered (R). Let the integer-valued functions S, I and R represent the number of companies that are in the state S, I and R, respectively, at time t.

The process starts at day $t = 0$, with $m \geq 1$ initially infected companies ($I(0) = m$) and the remainder being susceptible ($S(0) = n - m$, $R(0) = 0$). Then at each day $t = 1, 2, 3, ...$, an infected company i infects each susceptible neighbour j on layer α independently with probability $w_{i,j}^{[\alpha]}$, after which each infected company i recovers independently with probability p. Once recovered, a company cannot be reinfected again. Infection or recovery of a node occurs simultaneously on all layers. The process continues until there are no more infected companies.

We model the transmission probabilities per edge (i, j) in each layer $\alpha \in \{1, 2, 3, 4\}$, where the values $\alpha = 1, 2, 3, 4$ correspond to the PMFG, continents, sectors and global layers respectively, as

$$w_{i,j}^{[\alpha]} = \begin{cases} \hat{\chi}_{i,j}^{U} \times \beta_1, & \alpha = 1, \\ \beta_\alpha, & \alpha \in \{2, 3, 4\}, \end{cases} \tag{3.4}$$

where $\hat{\chi}^U_{i,j}$ is defined in (3.3) and β_α for $1 \leq \alpha \leq 4$ are parameters to be estimated (see Section 3.4.2.2) using the definition of 'infection' in a financial context that is provided in Section 3.4.1.

3.4 Application to financial crises

In this section we fit the model to the 2008 and to the 2020 financial crises. In Section 3.4.1, we define what is meant for a company in the data set to be 'financially infected'. We then build and compare two different networks representing the financial dependency between the companies in the periods prior to the 2008 financial and to the 2020 financial crises. We study how the model can be used to predict future infections in each case, using recent infection data. We finally assess the importance for predictive accuracy of each layer within our network.

3.4.1 Infection

We define a company in the data set to be infected whenever the volatility of its stock returns over a given period exceeds a predetermined threshold (meaning that the company's stock price is unstable) and its average stock return for the same period is negative.

The volatility for a time horizon $T > 1$ of company i at day t is defined as the standard deviation of the stock returns in the prior T trading days and is calculated as follows:

$$V_{i,t} = \sqrt{\frac{1}{T} \sum_{j=t-T}^{t-1} (x_{i,j} - \mu_{i,t})^2}, \tag{3.5}$$

where $\mu_{i,t}$ is the mean stock return over the same period and $x_{i,j}$ is defined in (3.1). Then a company i is defined to be infected at day t whenever $V_{i,t} \geq \sigma_i$ and $\mu_{i,t} < 0$. In the following analysis we use $T = 21$ trading days (one trading month) and the threshold σ_i to be the 90% quantile of the (empirical) volatility distribution over the whole period for each company.

Using a rolling window of historical returns over the past 21 days is common in risk analysis [9, 134, 198] and suitable for estimating volatility for daily data because it strikes a balance between capturing recent changes in volatility and incorporating sufficient historical data to generate a stable estimate. This balance is especially important given our focus on identifying 'infection': longer periods could include stock price fluctuations whose effect on the market has passed, while short periods are likely to be sensitive to noise. Our choice of σ_i is determined by the 90% quantile of the (empirical) volatility distribution; however, it should be acknowledged that in practice, determining this threshold at a specific time without knowledge of future volatility values may not be feasible. Therefore, the quantile threshold is primarily used as a benchmarking tool to compare and analyse volatility levels across companies in a historical context. An example of a particular company's volatility curve is shown in Figure 3.2a. The red horizontal line corresponds to 90% threshold line. We can see that with a higher threshold σ_i, the final, post-2020 'infection period' (high volatility period) would not be captured.

Figure 3.2b illustrates the number of infected companies per day, as determined using the procedure described above. It can be seen that the highest peaks in the numbers of infections are in two periods: from the year 2008 to the year 2010 (in orange), corresponding to the 2008 financial crisis, and from the beginning of 2020 until July 2021 (in purple), corresponding to the COVID-19 crisis. Furthermore it can be seen that the other infection peaks also correspond to different financial crises including the early 2000s recession

(in blue). Around 2012 we can see a growth in infections (in green) which corresponds to the period during the Cypriot financial crisis and shortly after the start of the Greek government debt crisis. Finally, we can see a slight growth in the number of infections during the period of 2015–2016 (in red), which includes a stock market sell-off when the value of stock prices declined globally.

Figures 3.3a and 3.3b illustrate the number of infected companies in the 2008 and 2020 crises, respectively, along with significant events that occurred during these periods. It can be seen that in the 2008 crisis after the Lehman Brothers' bankruptcy in September 2008 (red vertical line on Figure 3.3a), there is a substantial increase in the number of infected companies. Subsequently, after the TARP (Troubled Asset Relief Program) was implemented in October 2008, the rate at which the companies become infected decreases (purple vertical line on Figure 3.3a) and after the ARRA (American Recovery and Reinvestment Act) was signed into law in February 2009, the companies start recovering (green vertical line on Figure 3.3a). In the 2020 crisis, shortly after the WHO (World Health Organization) declared a global health emergency in March 2020 (blue vertical line on Figure 3.3b), accompanied by national lockdown measures in many countries[2], the number of infected companies increases sharply in a short time period. When the USA and UK governments started offering stimulus packages[3], the rate at which the infections spread declined (around the red vertical line on Figure 3.3b). Finally, in

[2]National emergency was declared in the US on March 13, 2020; the United Kingdom went into lockdown on March 23, 2020; a national lockdown in Italy was imposed on March 9, 2020; nationwide lockdown in France started on March 17, 2020; from March 13, 2020, German states mandated school and kindergarten closures and travel restrictions were put in place in Austria, Denmark, France, Luxembourg and Switzerland; Japan officially declared the COVID-19 outbreak as a national emergency on March 19, 2020.

[3]The Main Street Lending Program (April 9, 2020), Primary Market Corporate Credit Facility (March 23, 2020), CARES Act (March 27, 2020), and Paycheck Protection Program Liquidity Facility (April 9, 2020) were launched in the USA; the Coronavirus Job Retention Scheme (March 1, 2020), Self-Employment Income Support Scheme (March 26, 2020), and Coronavirus (Large) Business Interruption Loan Scheme (March 23, 2020) were launched in the UK

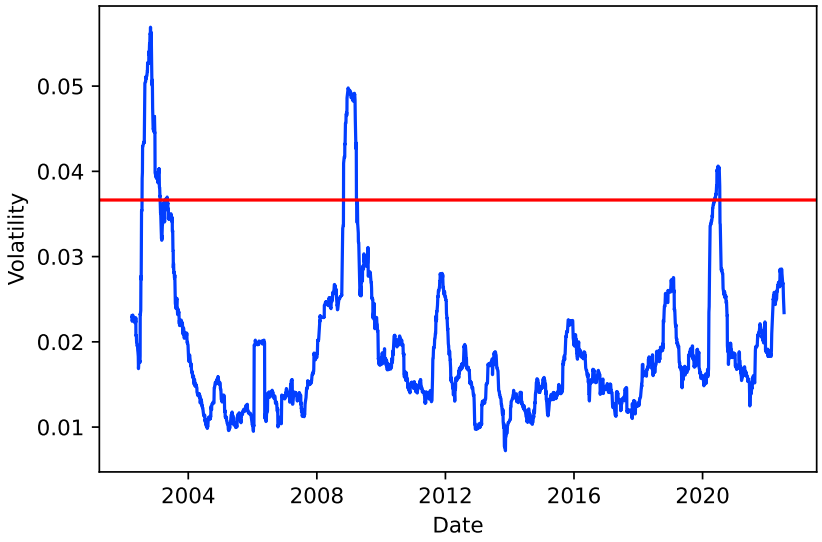

(a)

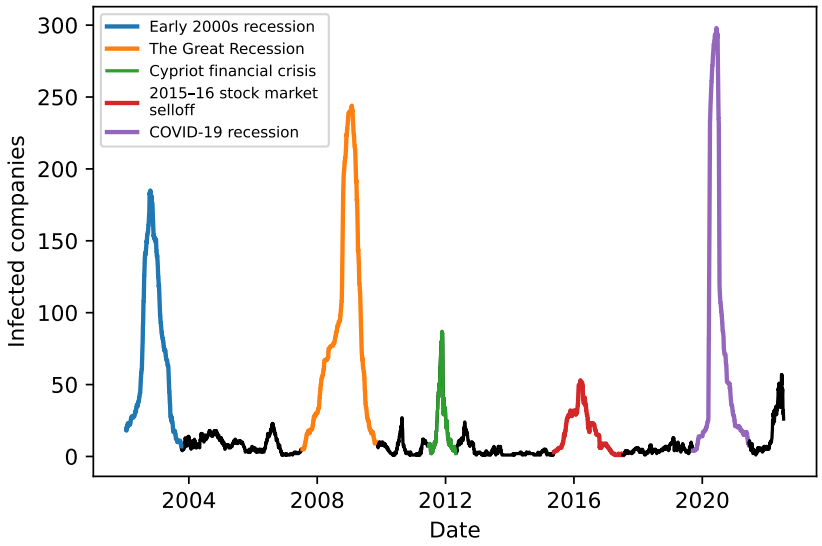

(b)

Figure 3.2: (a) Volatility curve (in blue) and 90% quantile threshold line (in red) for an example company from the dataset. The volatility at day t for the period from 07/02/2002 to 29/02/2020 is calculated using the formula in (3.5) using a rolling window of $T = 21$ trading days; (b) Total number of infected companies per day for the period from 07/02/2002 to 29/02/2020. The five main peaks from left to right correspond to Early 2000s recession, the Great Recession, the Cypriot financial crisis, the 2015-16 stock market selloff and the COVID-19 recession. It can be seen that the highest volatility periods in (a) correspond to the highest peaks in (b).

most of the countries the lockdown restrictions were eased between June and July 2020 (the period around the purple vertical line on Figure 3.3b), leading to recoveries. However, a month later (green vertical line on Figure 3.3b), COVID-19 cases started increasing worldwide[4]. In summary, the analysis of Figures 3.3a and 3.3b reveals the impact of these events on the spread and recovery of infected companies during the 2008 and 2020 crises, and indicates the suitability of our empirical definition of 'infection'.

3.4.2 Predicting future infections

It is of key importance to be able to predict future infections, given past data, for risk prevention and mitigation purposes. This is essential for a number of reasons. Firstly, it enables decision-makers and market participants to make necessary adjustments before the crisis fully unfolds, which can help minimize the damage caused by the crisis, thereby reducing its impact on financial markets, the economy, and society as a whole. Secondly, understanding the potential risks and vulnerabilities in the financial system can allow policymakers and investors to take steps to mitigate these risks such as implementing regulations or reducing exposure to risky assets. Lastly, predicting and preventing financial crises is crucial for ensuring the stability and health of the global financial system, which is vital in protecting investors and promoting sustainable economic growth. By being able to anticipate future infections, we can take the necessary steps to prevent their spread and minimize their impact on the economy and society.

Our study aims to investigate whether it is possible to predict future events

[4]The US confirms more than 50,000 new COVID-19 cases in one day for the first time, the Australian city of Melbourne goes back into lockdown for six weeks after a second outbreak, Florida reports a record 11,458 daily COVID-19 cases, Texas records more than 10,000 daily cases of COVID-19 for the first time, India becomes the third country to record one million cases of COVID-19, the WHO says the Middle East is at a 'critical threshold' with over one million COVID-19 cases.

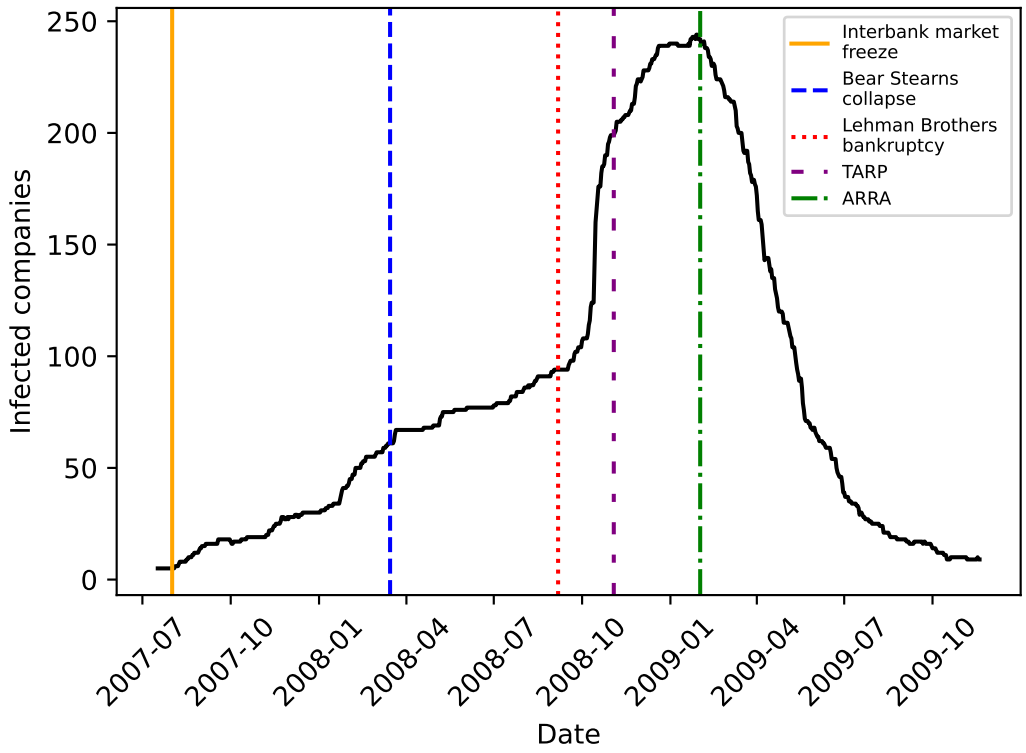

(a)

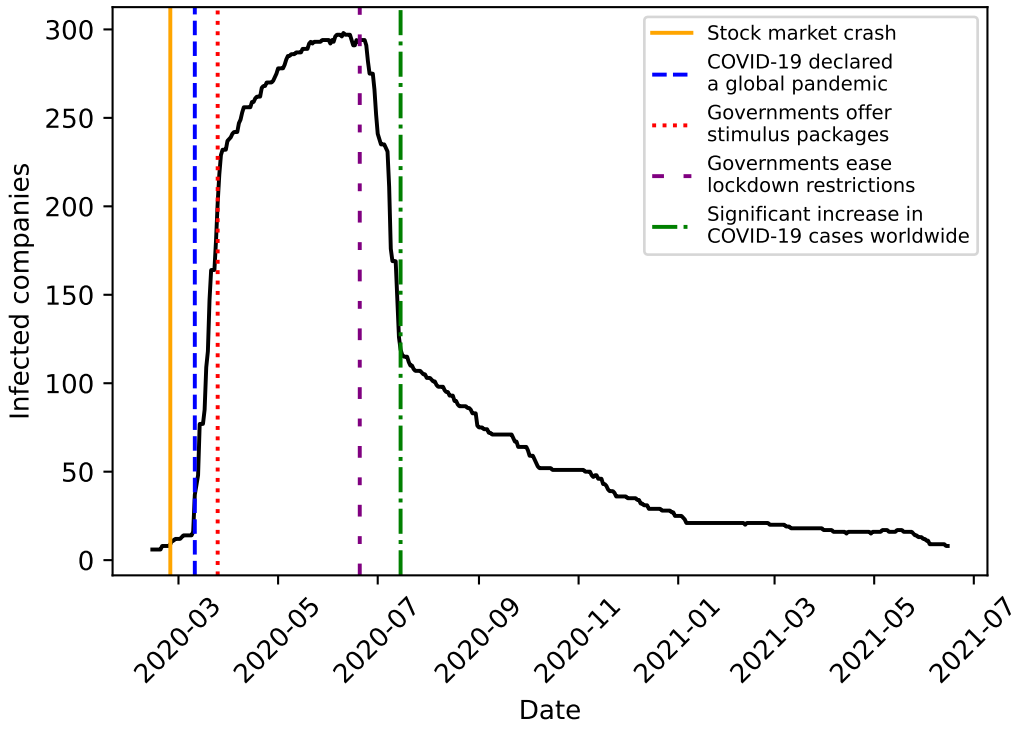

(b)

Figure 3.3: The total number of infected companies within our dataset, as defined in Section 3.4.1, during (a) the 2008 financial crisis and (b) the 2020 financial crisis. The vertical lines show the dates of significant events during each crisis.

during a crisis by analysing transmission and recovery probabilities in the preceding days. Specifically, we examine the accuracy of our model to estimate the number of infected companies in the future, given data from the past n days of the crisis. We evaluate the accuracy of predictions in terms of the total number of infected companies, the number per sector and continent, and identifying the specific companies most likely to be affected.

We construct and compare two distinct networks, one for the 2008 and one for the 2020 financial crises, that represent the dependence structure preceding each crisis, as described in Section 3.4.2.1. These networks are then used to simulate future infections during the corresponding crises. The empirical results shown in Sections 3.4.2.2–3.4.2.4 suggest that the choice of n appreciably impacts the fidelity of prediction, with optimal choice depending on the specific crisis period and the performance metric being considered. In addition, the findings indicate that the multiplex network model is more accurate than the homogeneous mixing model in predicting the future course of a crisis, supporting our assumption that incorporating network structure in the financial network is critical for precise predictions of future financial distress outbreaks. Additionally, the consistent superiority of the full model over the homogeneous mixing model implies that taking into account the various types of connections is crucial to capture the intricate connections between firms that lead to financial crises.

3.4.2.1 The 2008 and the 2020 financial networks

As explained in section 3.3.1, the first network is constructed using all available data before the 2008 financial crisis, which includes the data from 17/01/2002 to 30/06/2007. The second network employs all available data before the 2020 financial crisis, i.e., from 17/01/2002 to 29/02/2020.

We now compare the community structure of the two PMFG networks. For each of the networks we divide the nodes into communities by maximizing the modularity [165] of the network via the Louvain algorithm [153]. More details can be found in Appendix A. Then, for the two sets of communities we estimate the similarity between them using the adjusted mutual information (AMI) score [228]. The AMI takes a value of 1 when the two partitions are identical (perfectly matched), while random partitions, having an expected AMI around 0 on average, can occasionally yield negative values (see Appendix A). The AMI score between the clusterings of the two PMFG networks is 0.2568, suggesting that the community structures of the two graphs are substantially different.

We then perform a clique analysis by adopting the n-clique algorithm of [172] to analyse the community structures. A clique in a graph G is a complete subgraph of G. A clique, in other words, is a subset of a network in which the nodes are more intensively linked to one another than to other members of the network. The maximal clique in a PMFG network consist of 4 nodes, and is also called a 4-clique. By detecting cliques, we can uncover natural clusters or communities of companies that have strong connections or similarities. Table 3.3 shows the structure of the different 3- and 4-cliques in the two PMFGs based on companies' continents and sectors, respectively. The analysis shows that in both networks communities based on continents are more likely to form than communities based on sectors. In addition, the high number of 3- and 4-cliques in which all the companies are in the same continent indicates a strong tendency for continent-based communities. Figures 3.4a and 3.4b illustrate the PMFG networks for the 2008 and 2020 financial crises, respectively, where the companies are coloured by continent. In both figures it can be seen that companies in the same continents tend to form clusters, indicating that the local level transmission is more likely to happen between companies in the

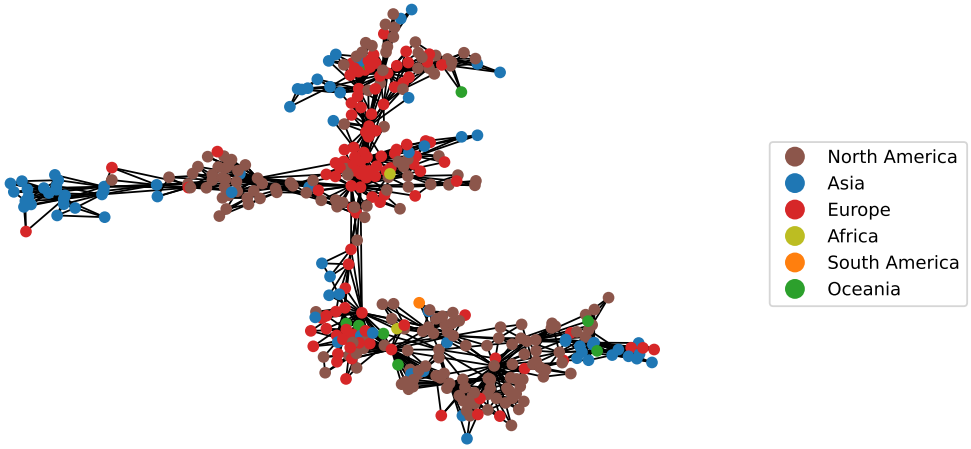

(a)

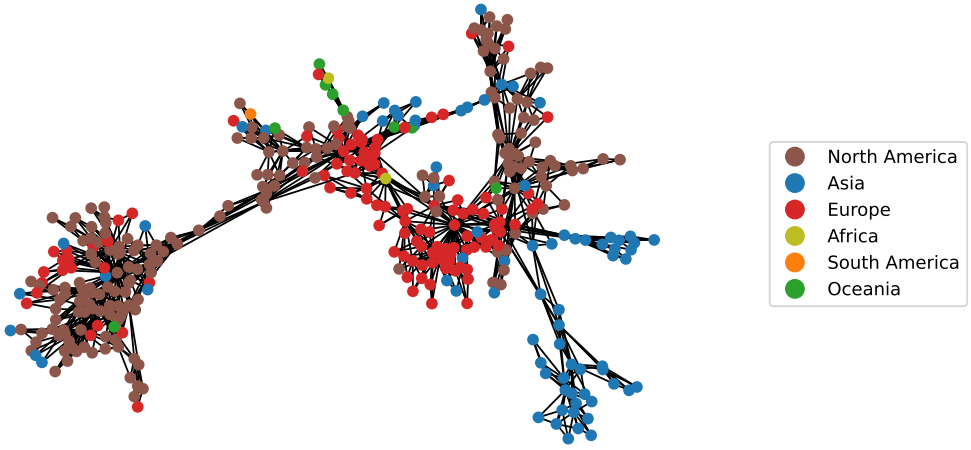

(b)

Figure 3.4: PMFG (tail dependence network) layers in the (a) 2008 and (b) 2020 financial networks, where the companies are coloured by continent. The two networks are constructed by the procedure described in Section 3.3.1.1 using stock price data for the periods from 16/01/2002 to 30/06/2007 and from 16/01/2002 to 29/02/2020, respectively.

same continent. The analysis of these data also suggests that communities based on sectors are more likely to form in the 2020 PMFG than in the 2008 PMFG, due to the higher occurrence of 4-cliques with all nodes in the same sector or three of the nodes in the same sector.

Clique type		Continents		Sectors	
		2008	2020	2008	2020
3-cliques	total number of 3-cliques	25	4	25	4
	all nodes in same continent/sector	14	4	4	0
	two nodes in same continent/sector	11	0	10	2
	all nodes in different continent/sector	0	0	11	2
4-cliques	total number of 4-cliques	372	392	372	392
	all nodes in same continent/sector	174	252	24	58
	three nodes in same continent/sector	133	93	53	91
	two nodes in same continent/sector	63	47	194	162
	all nodes in different continent/sector	2	0	101	81

Table 3.3: Clique analysis of the PMFG networks showing the cliques structure based on the sector or continent in which each company is based.

3.4.2.2 Prediction of the number of infected companies

Our study evaluates the model's accuracy in predicting the number of companies that will be infected or recovered in the future k crisis days, based on the infection data from the past n days, utilizing a 'sliding window' technique. Firstly, we fit the model to the initial data window (data window 1), comprising data from day 1 to $n+1$, obtaining maximum likelihood estimates $\hat{\beta}_i$ of the layer transition probabilities β_i for $1 \leq i \leq 4$ and $\hat{p}$ for the recovery probability p. The detailed methodology for estimating the transmission and recovery probabilities is described in Sections 3.4.2.2.1 and 3.4.2.2.2, respectively. Next, we simulate $N = 10{,}000$ realisations of the estimated SIR model for the upcoming k days (from day $n+2$ to day $n+k+1$, denoted prediction window 1), with the initial data being that from day $n+1$. After each simulation, we record the total number of infected companies, the number

of newly infected companies, and the number of newly recovered companies, and calculate the mean of all simulations as the prediction. We then 'slide the window' forward by one day and refit the model to the period from day 2 to day $n+2$ of the crisis (data window 2), re-estimating $\hat{\beta}_i$ for $1 \leq i \leq 4$ and $\hat{p}$ for the new window. We repeat the above steps for each subsequent data window, with the final prediction window covering the period from day $L-k+1$ to L, where L is the length of the crisis in days. The entire procedure is explained step-by-step in Algorithm 2 step 3a).

Figures 3.5, 3.6 and 3.7 display the model predictions (coloured lines) alongside the observed infections (black lines) at selected time points. Predictions are computed from the mean of all $N = 10{,}000$ simulations for the future $k \in \{10, 20, 30\}$ days of each crisis, respectively, given infection data on the previous $n \in \{1, 30\}$ days. In both figures it can be seen that fitting the model to the previous $n = 1$ crisis days gives the largest error between the actual and the predicted total number of infected individuals after k days for both values of k. For choices of $n > 1$, with greater prediction accuracy, we nevertheless observe large errors at those time points where significant changes in infection or recovery occur. This is natural since predictions are based on data prior to these change-points; it is important to note, however, that such events are often due to extrinsic factors, such as government intervention that could in principle be accommodated within the model. For example, substantial errors are observed in recovery prediction during periods associated with ARRA (2008; see Figs 3.3a, 3.5, 3.6, 3.7) and stimulus packages and lockdown restrictions (2020; see Figs 3.3b, 3.5, 3.6, 3.7). For all other periods in both crises, we obtain good prediction accuracy for suitable choices of k and n, as confirmed by further analysis. The most precise forecasts were observed with $k = 10$, while the least precise predictions were observed with $k = 30$.

In order to compare the predicted and actual numbers of newly infected, to-

Algorithm 2 Model predictions for the future k days

1. Set $t = 1$.
2. Fit the model to the period from day t to day $t + n$ (window t) and estimate $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$ by maximizing $\mathcal{L}_{t,t+n}$ (see equation (3.10)) and $\mathcal{L}^p_{t,t+n}$ (see equation (3.12)), respectively.
3. Run $N = 10{,}000$ SIR simulations on the network from day $t + n + 1$ to day $t + n + k$ (i.e. predict infections in the next k days) using the values of $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$ estimated in Step 2.

   (a) Predict the number of infected (recovered) companies

   - Count the number of newly infected (recovered, respectively) companies in each simulation and take the mean number of the newly infected (recovered, respectively) companies over all N simulations as a prediction for the number of newly infected (recovered, respectively) companies in the period $t + n + 1 \leq t \leq t + n + k$ (prediction window t).
   - Count the total number of infected companies at the end of each simulation and take the mean total number of infected companies over all N simulations as a prediction for the total number of infected companies on the final day of the prediction window t.

   (b) Predict the infected companies' continents/sectors

   - Collect a multiset (set allowing multiple instances of the same element) containing the continents/economic sectors of newly infected companies in each simulation and take it as a prediction for which continents/sectors are likely to be infected in prediction window t. For each simulation compare the predicted to the observed continents/sectors multiset using the Sørensen–Dice similarity coefficient (3.14). The mean over all N simulations is then used as an indicator of prediction quality for prediction window t.

   (c) Predict the infected companies

   - Collect a list of newly infected companies in each simulation and record the proportion of the N simulations in which each company was infected. Then, for every company, consider this proportion as the probability of infection during prediction window t of the crisis.

4. Update $t = t + 1$.
5. If $t > L - k - n$ stop. Else go back to Step 2.

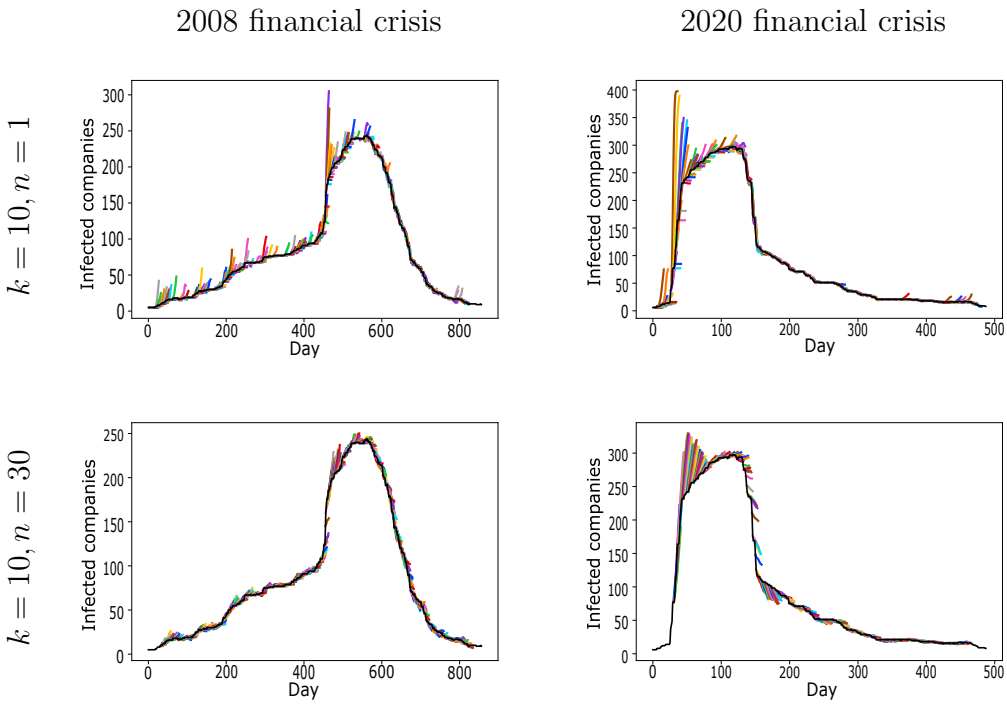

Figure 3.5: The curves (in colours) show the predicted mean total number of infected companies for each sliding window for the next $k = 10$ days over $N = 10{,}000$ simulations, fitting the model to the previous n crisis days for the 2008 financial crisis (left column) and 2020 financial crisis (right column), respectively. The black lines show the observed number of infected companies as determined in Section 3.4.1, providing a reference for comparison with the model predictions.

tal infected and newly recovered companies for each sliding window i (where $1 \leq i \leq L-k-n$), we calculate the absolute difference between the predicted and actual values in each simulation and then we take the average. To ensure accurate evaluation of the model's performance in predicting the number of newly infected (recovered, respectively) companies, we measure predictive accuracy during specific time periods in which these processes occur. Specifically, we consider the period encompassing newly infected (recovered, respectively) companies, which corresponds to the time before (after, respectively) day 600 during the 2008 financial crisis, and before (after, respectively) day 120 in the case of the 2020 financial crisis. The results are shown in Figures 3.8, 3.9 and 3.10. The figures display the distribution of the absolute difference between the predicted and actual number of total infected (row 1), newly

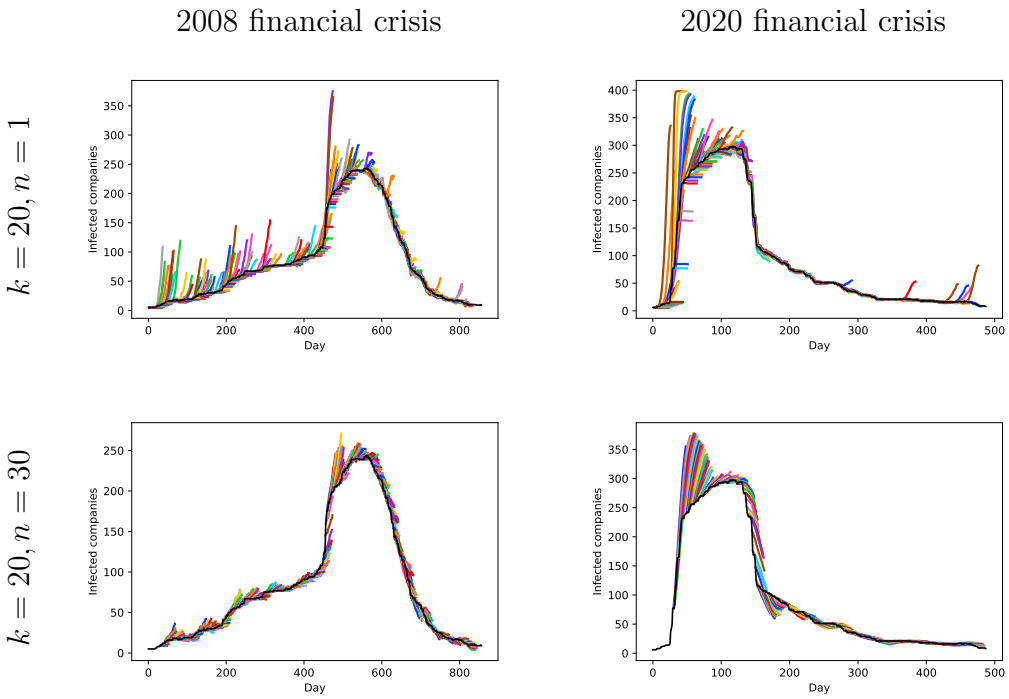

Figure 3.6: The curves (in colours) show the predicted mean total number of infected companies for each sliding window for the next $k = 20$ days over $N = 10{,}000$ simulations, fitting the model to the previous n crisis days for the 2008 financial crisis (left column) and 2020 financial crisis (right column), respectively. The black lines show the observed number of infected companies as determined in Section 3.4.1, providing a reference for comparison with the model predictions.

infected (row 2) and newly recovered (row 3) companies when the model is fitted to the 2008 financial crisis (left column) or the 2020 financial crisis (right column). The white dots connected by white lines indicate the mean absolute difference over all sliding windows. We present results for prediction horizons of $k = 10$, 20 and 30 days, respectively. It can be seen that the trends in mean accuracy are similar for all choices of k. The results indicate that the optimal window size for predicting future infections varies depending on the crisis being analysed. Specifically, for the 2008 financial crisis, the optimal window size is $n = 10$ for predicting both the future total number of infected and number of newly infected companies after k days, while for the 2020 financial crisis, the optimal window size is $n = 3$ for predicting the future total number of infected companies, and $n = 10$ for predicting the future number

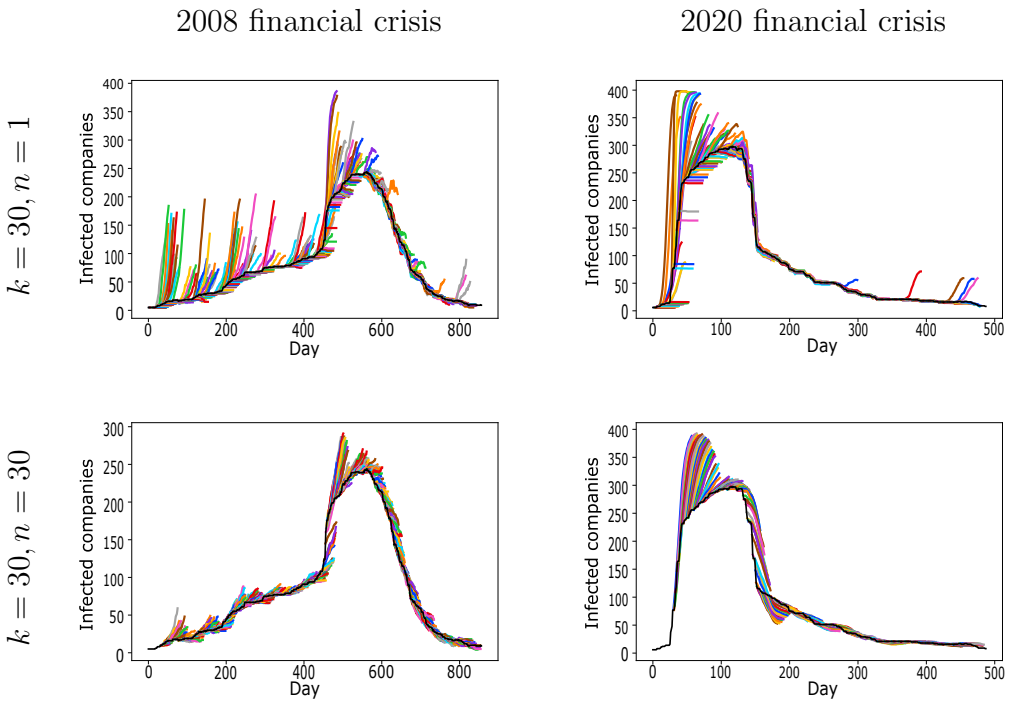

Figure 3.7: The curves (in colours) show the predicted mean total number of infected companies for each sliding window for the next $k = 30$ days over $N = 10{,}000$ simulations, fitting the model to the previous n crisis days for the 2008 financial crisis (left column) and 2020 financial crisis (right column), respectively. The black lines show the observed number of infected companies as determined in Section 3.4.1, providing a reference for comparison with the model predictions.

of newly infected companies. In contrast, a window size of $n = 1$ day for the 2008 crisis and a window size of $n = 30$ days for the 2020 crisis result in the worst predictions. Interestingly, when predicting the number of newly recovered companies in the future k days, for both crises, the worst predictions are obtained when the window size is the largest, i.e. $n = 30$, while the best predictions are obtained when the window size is the smallest, i.e. $n = 1$; we note, however, that the variation with n is not large.

3.4.2.2.1 Estimating the transmission probabilities

In this section we explain the maximum likelihood procedure for estimating the transmission and recovery probabilities. Let us denote the set of infected

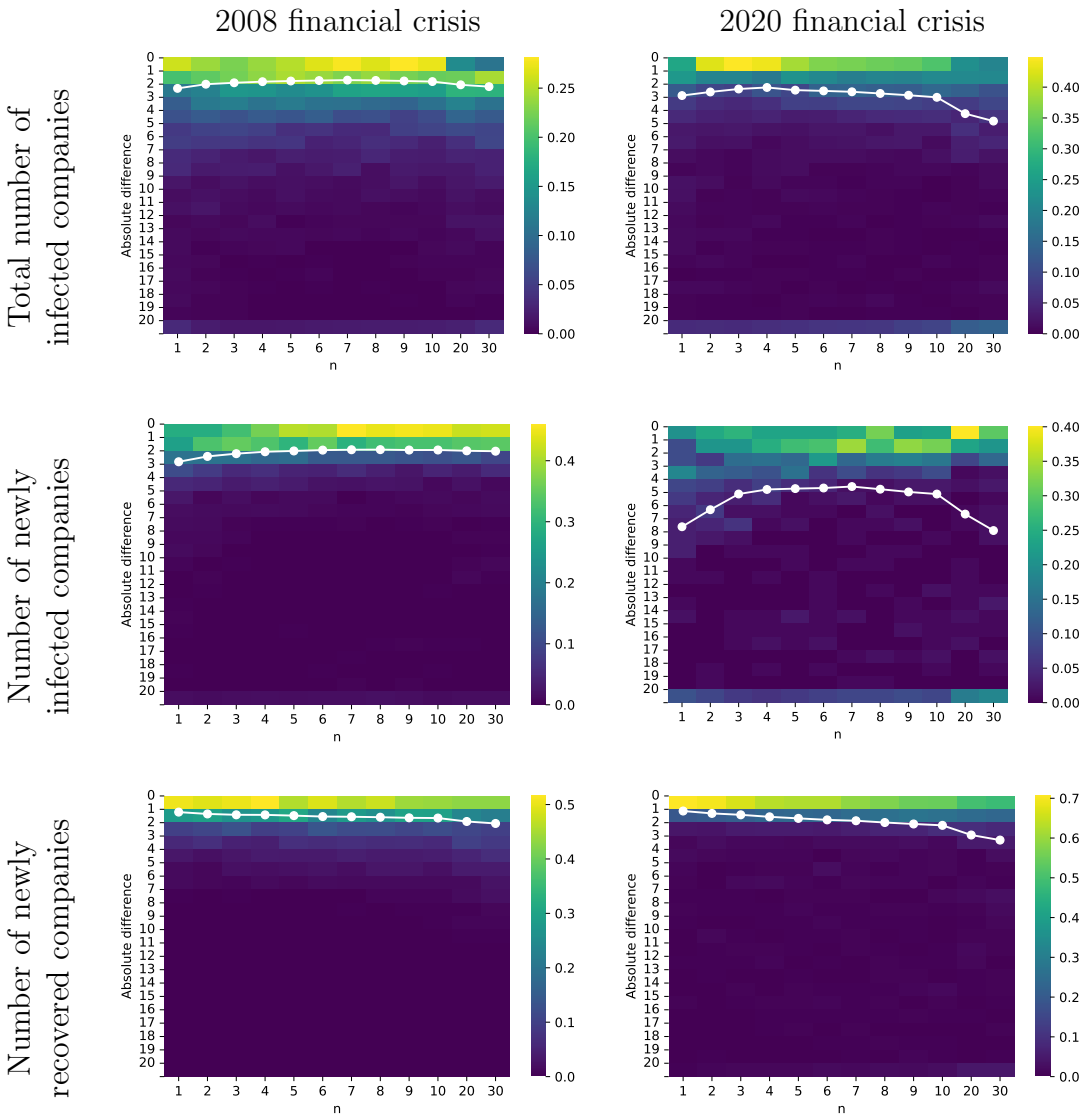

Figure 3.8: Heatmaps showing the distribution of the absolute difference between predicted and actual total number of infected companies after $k = 10$ days (top row), number of newly infected companies in the future $k = 10$ days (middle row) and number of newly recovered companies in the future $k = 10$ days (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days. The white dots indicate the mean absolute difference over all sliding window predictions.

neighbours in layer α of a susceptible company j at time t as $\mathrm{I}_j^{\alpha}(t)$. We know that the probability that company $i \in \mathrm{I}_j^{\alpha}(t)$ does not infect company j through layer α is $1 - w_{i,j}^{[\alpha]}$. Therefore, the probability that company j does not become infected at time $t + 1$ is the probability that none of its infected neighbours

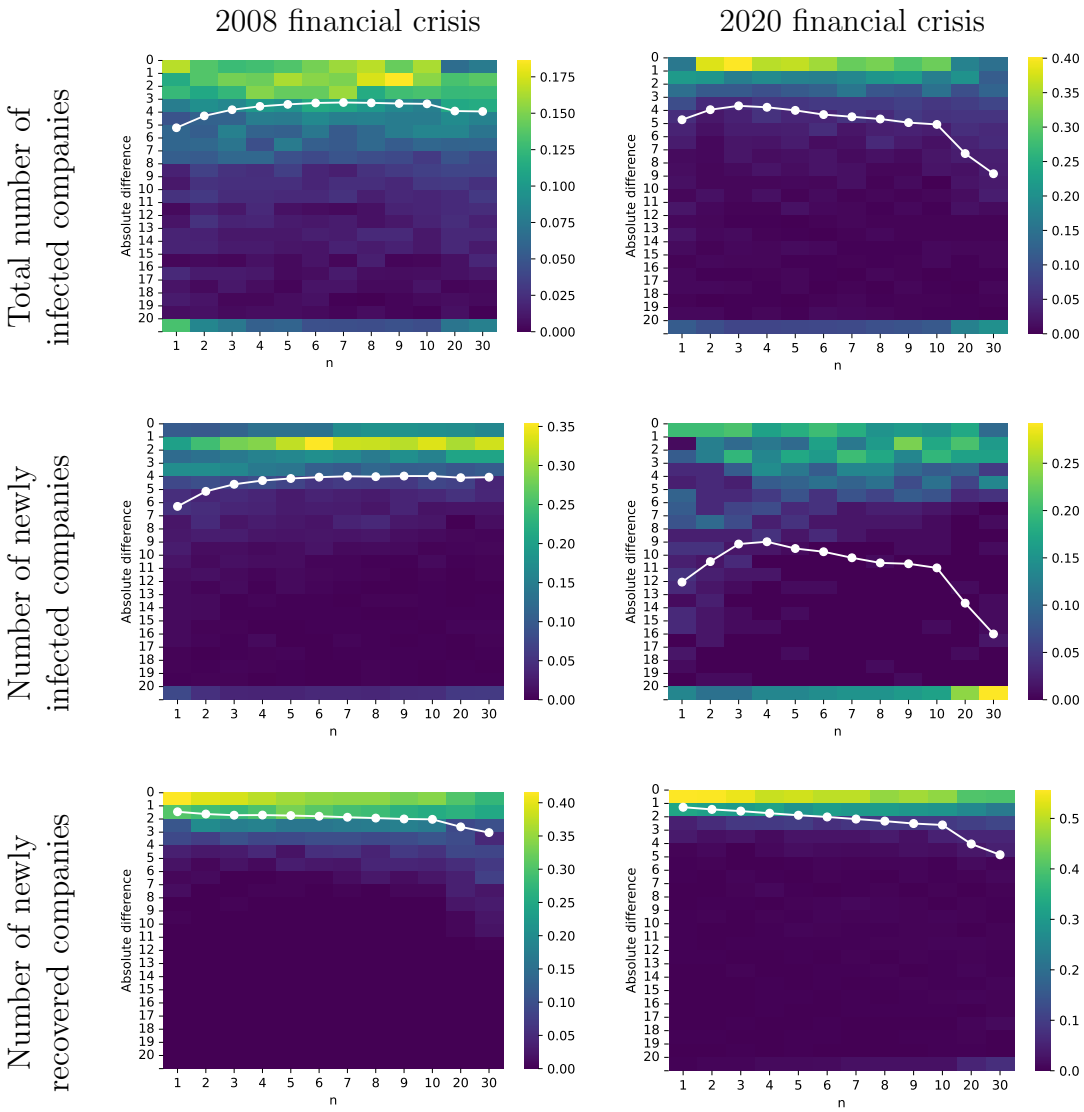

Figure 3.9: Heatmaps showing the distribution of the absolute difference between predicted and actual total number of infected companies after $k = 20$ days (top row), number of newly infected companies in the future $k = 20$ days (middle row) and number of newly recovered companies in the future $k = 20$ days (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days. The white dots indicate the mean absolute difference over all sliding window predictions.

in any of the layers $1 \leq \alpha \leq 4$ at time t infects j:

$$\begin{aligned} &\mathrm{P}(j \text{ does not get infected at time } t+1) \\ &= \prod_{\alpha=1}^{4} \prod_{i \in \mathrm{I}_j^{\alpha}(t)} \mathrm{P}(i \text{ does not infect } j \text{ on layer } \alpha) = \prod_{\alpha=1}^{4} \prod_{i \in \mathrm{I}_j^{\alpha}(t)} (1 - w_{i,j}^{[\alpha]}). \end{aligned} \tag{3.6}$$

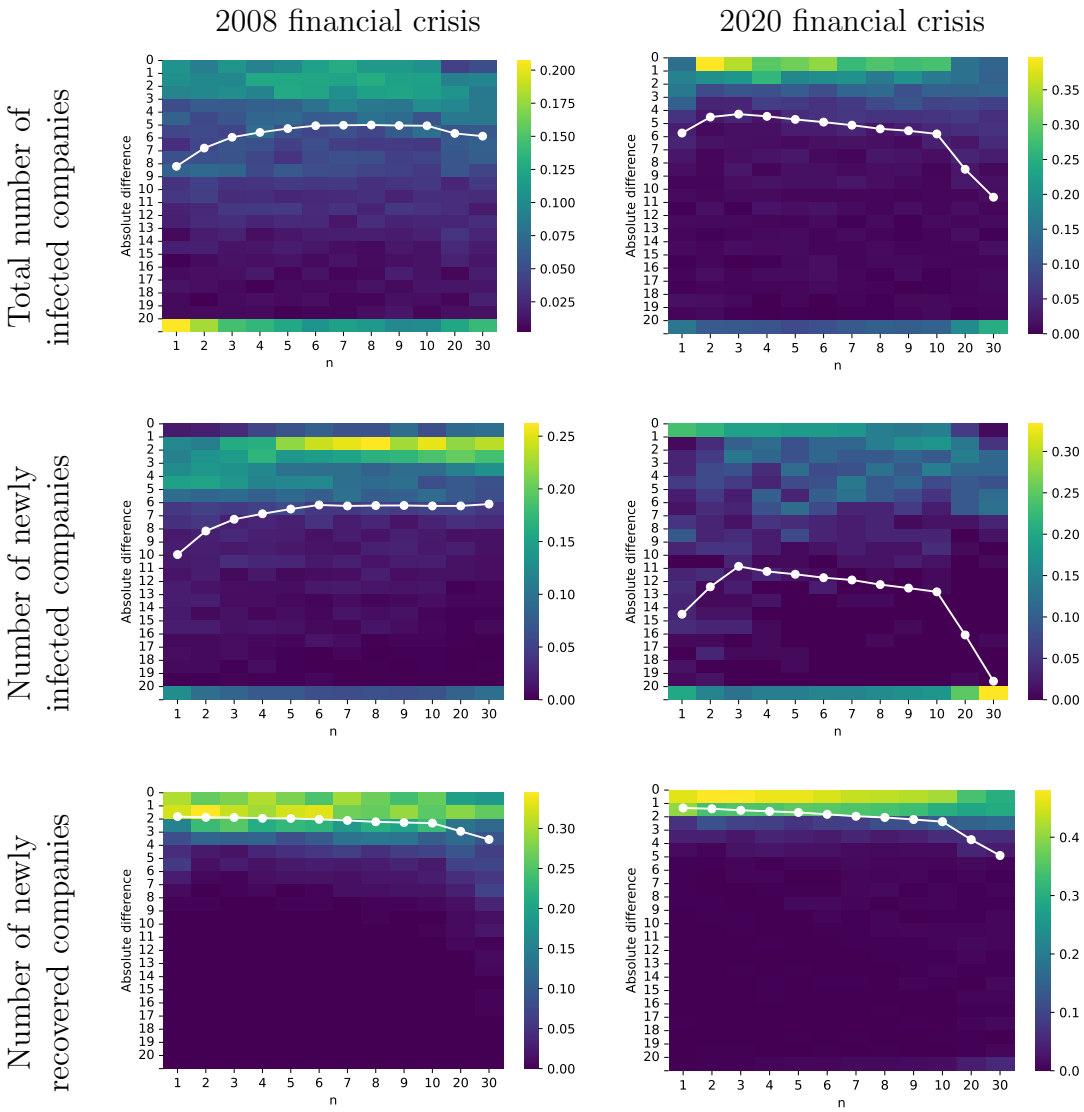

Figure 3.10: Heatmaps showing the distribution of the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days and at a prediction horizon of $k = 30$ days. The white dots indicate the mean absolute difference over all sliding window predictions.

It follows that the probability that company j gets infected at day $t + 1$ is

$$\mathrm{P}(j \text{ gets infected at time } t+1) = 1 - \prod_{\alpha=1}^{4} \prod_{i \in \mathrm{I}_j^{\alpha}(t)} (1 - w_{i,j}^{[\alpha]}). \tag{3.7}$$

Let S_t be the set of susceptible companies at time t and SI_{t+1} denote the

subset of companies which get infected at time $t+1$, so that $\mathrm{S}_{t+1} = \mathrm{S}_t \setminus \mathrm{SI}_{t+1}$. Then the likelihood for the observed infections from time t to time $t+1$ is

$$\begin{aligned}\mathcal{L}_{t,t+1} = &\prod_{j \in \mathrm{SI}_{t+1}} \mathrm{P}(j \text{ gets infected at time } t+1) \\ &\times \prod_{j \in \mathrm{S}_{t+1}} \mathrm{P}(j \text{ does not get infected at time } t+1).\end{aligned} \tag{3.8}$$

Substituting (3.6) and (3.7) into (3.8) we get

$$\mathcal{L}_{t,t+1} = \left\{ \prod_{j \in \mathrm{SI}_{t+1}} \left(1 - \prod_{\alpha=1}^{4} \prod_{i \in \mathrm{I}_j^\alpha(t)} (1 - w_{i,j}^{[\alpha]}) \right) \right\} \left\{ \prod_{j \in \mathrm{S}_{t+1}} \prod_{\alpha=1}^{4} \prod_{i \in \mathrm{I}_j^\alpha(t)} (1 - w_{i,j}^{[\alpha]}) \right\}. \tag{3.9}$$

Suppose we want to estimate the parameters β_i, $1 \leq i \leq 4$ for a period of n days (window length n) from day d to day $d+n$. By substituting $w_{i,j}^{[\alpha]}$ with its representation in terms of β_i, $1 \leq i \leq 4$ (see (3.4)), and writing the likelihood of the observations as

$$\mathcal{L}_{d,d+n} = \prod_{t=d}^{d+n-1} \mathcal{L}_{t,t+1}, \tag{3.10}$$

where $\mathcal{L}_{t,t+1} = \mathcal{L}_{t,t+1}(\beta_1, \beta_2, \beta_3, \beta_4)$ is defined in (3.9), a maximization of (3.10) numerically with respect to $(\beta_1, \beta_2, \beta_3, \beta_4)$ leads to the maximum likelihood estimate $(\hat{\beta}_1, \hat{\beta}_2, \hat{\beta}_3, \hat{\beta}_4)$ of $(\beta_1, \beta_2, \beta_3, \beta_4)$.

3.4.2.2.2 Estimating the recovery probability

Suppose that each company recovers on day t with probability p. Let us denote by II_{t+1} and IR_{t+1} the set of companies which were infected at time t and did not recover at time $t+1$ and the set of companies which were infected at time t but recovered at time $t+1$, respectively. Then the likelihood for the recoveries is

$$\mathcal{L}_{t,t+1}^{p} = (1-p)^{|\mathrm{II}_{t+1}|} \times p^{|\mathrm{IR}_{t+1}|}. \tag{3.11}$$

Hence, the recovery likelihood from day d to day $d+n$ is

$$\mathcal{L}^p_{d,d+n} = \prod_{t=d}^{d+n-1} \mathcal{L}^p_{t,t+1} = (1-p)^{\sum_{t=d}^{d+n-1} |\mathrm{II}_{t+1}|} \times p^{\sum_{t=d}^{d+n-1} |\mathrm{IR}_{t+1}|}. \tag{3.12}$$

The maximum likelihood estimate $\hat{p}$, namely the value of p which maximises (3.12), is then

$$\hat{p} = \frac{\sum_{t=d}^{d+n-1} |\mathrm{IR}_{t+1}|}{\sum_{t=d}^{d+n-1} (|\mathrm{II}_{t+1}| + |\mathrm{IR}_{t+1}|)} = \frac{\sum_{t=d}^{d+n-1} |\mathrm{IR}_{t+1}|}{\sum_{t=d}^{d+n-1} |\mathrm{I}_t|}, \tag{3.13}$$

where I_t is the set of infected companies at time t.

3.4.2.3 Geographic- and sector-specific predictions

In this section, we investigate the model's ability to predict the geographical location and economic sector of the infected companies in the next k days, based on the previous n days' infection data. Rather than counting the number of infected companies, for each simulation we construct a multiset (i.e., a set allowing for multiple instances of each of its elements) that includes the continents or sectors corresponding to the predicted infected companies in that simulation. We then compare each multiset to the observed continents or sectors multiset using the Sørensen–Dice similarity coefficient for multisets, defined as

$$D(A,B) = \frac{2|A \cap B|}{|A| + |B|}. \tag{3.14}$$

Here A and B are multisets, not both empty, $|A|$ and $|B|$ denote the number of elements in A and B, respectively, and if an element appears in both A and B, it is included in the intersection $A \cap B$ with its minimal number of occurrences observed in A and B. The Sørensen–Dice coefficient takes values $D \in [0,1]$ with $D = 1$ indicating identical multisets, and $D = 0$ complete dissimilarity.

By calculating the mean Sørensen–Dice coefficient from all simulations we obtain a measure of performance that reflects the overall effectiveness of the method for each prediction. Figures 3.11, 3.12 and 3.13 illustrate the distribution of Sørensen–Dice coefficients when comparing the predicted and actual continents and economic sectors of newly infected companies in the future $k = 10$, 20 and 30 days, respectively, for different values of n. The left column presents the results when the model is fitted to the 2008 financial crisis, while the right column shows the results for the 2020 financial crisis. The results in each case indicate rather different optimal choices: for 2008, the least accurate predictions are obtained when using only the most recent data ($n = 1$), while for the 2020 crisis, smaller windows are in general preferable with $n = 30$ giving the worst predictions. However, apart from these worst cases, the dependence on n is not strong: for the 2008 data, very little variation in prediction accuracy as a function of n is observed, while for 2020, all choices $1 \leq n \leq 10$ give similar results. The results for all studied values of k are similar, indicating that the model's ability to predict future infected continents and sectors remains stable for most values of the size n of the sliding window.

3.4.2.4 Predicting which specific companies will be infected

In this section, we investigate the performance of the model in predicting which specific companies are likely to be infected in the future k days, using data from the past n days. Predicting which specific companies will be infected is an inherently challenging task, which arises from the intricate interplay of various factors and uncertainties involved in identifying individual companies that may be affected.

To evaluate the accuracy of the model, we employ Algorithm 2 with step 3c.

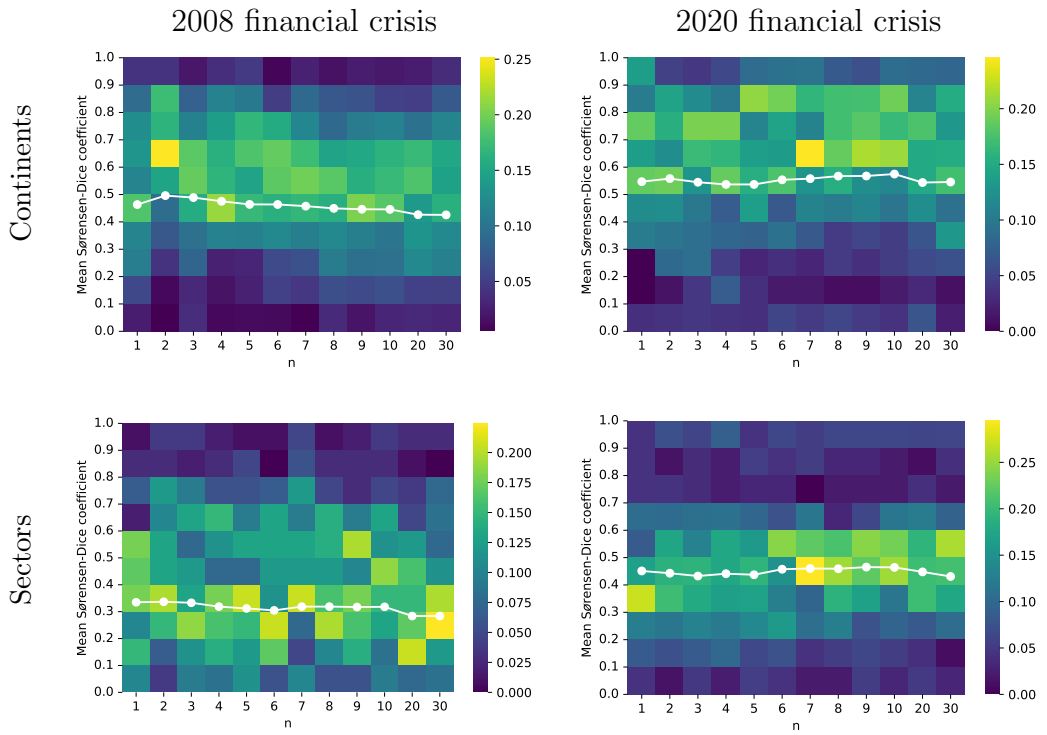

Figure 3.11: Heatmap for the distribution of the mean Sørensen–Dice coefficient between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 10$ days for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days. The white dots indicate the mean over all sliding window predictions.

That is, for each prediction window, we compute the proportion of the N realisations of our model, in which each company is infected. These proportions are then used as estimates of the probabilities of infection for each company during the prediction period t.

To evaluate the accuracy of the model, we employ two different measures: Accuracy and F_1-score. It is important to consider both performance metrics, rather than relying on a single metric, because each metric provides valuable insights into different aspects of the model's performance. The Accuracy is defined as

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + FN + TN},$$

where TP and TN denote the numbers of true positive and true negative pre-

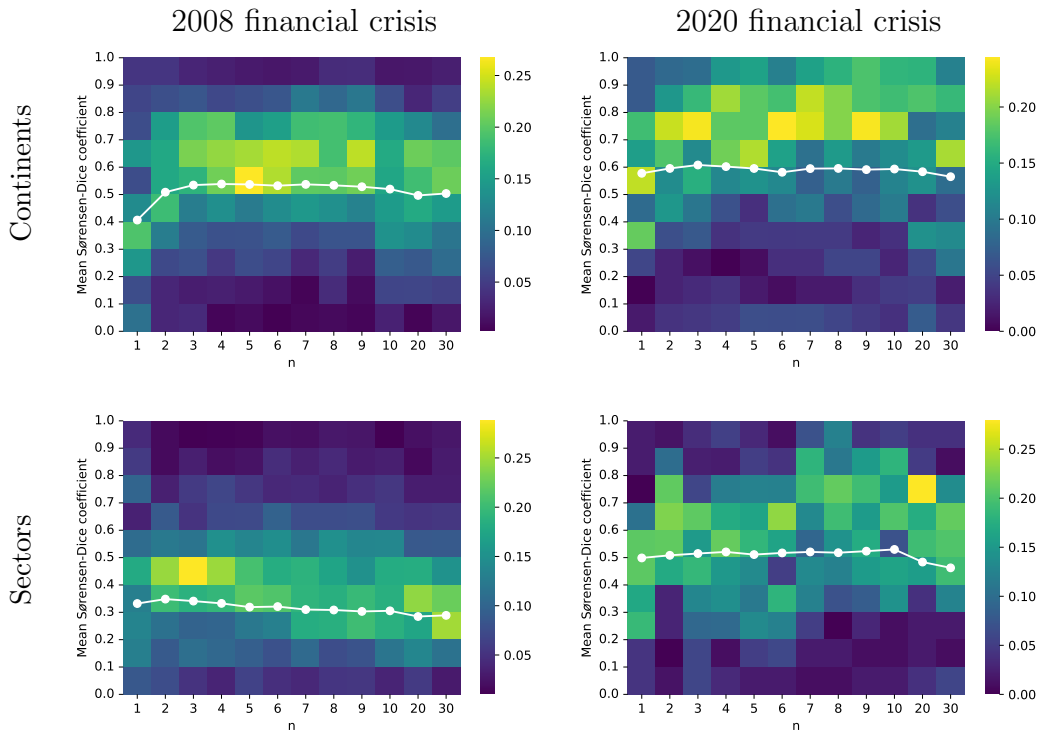

Figure 3.12: Heatmap for the distribution of the mean Sørensen–Dice coefficient between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 20$ days for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days. The white dots indicate the mean over all sliding window predictions.

dictions, respectively (i.e., the numbers of specific correctly predicted companies to become infected or to remain uninfected, respectively), and FP and FN denote the numbers of false positive and false negative predictions, respectively (i.e., the numbers of companies predicted wrongly to become infected or to remain uninfected, respectively). Hence, the Accuracy indicates the ratio of correctly predicted observations to the total observations. The F_1-score is a commonly used performance metric to evaluate the accuracy of a binary classifier or model and is defined as

$$F_1 = \frac{2 \times \text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}} = \frac{2TP}{2TP + FN + FP},$$

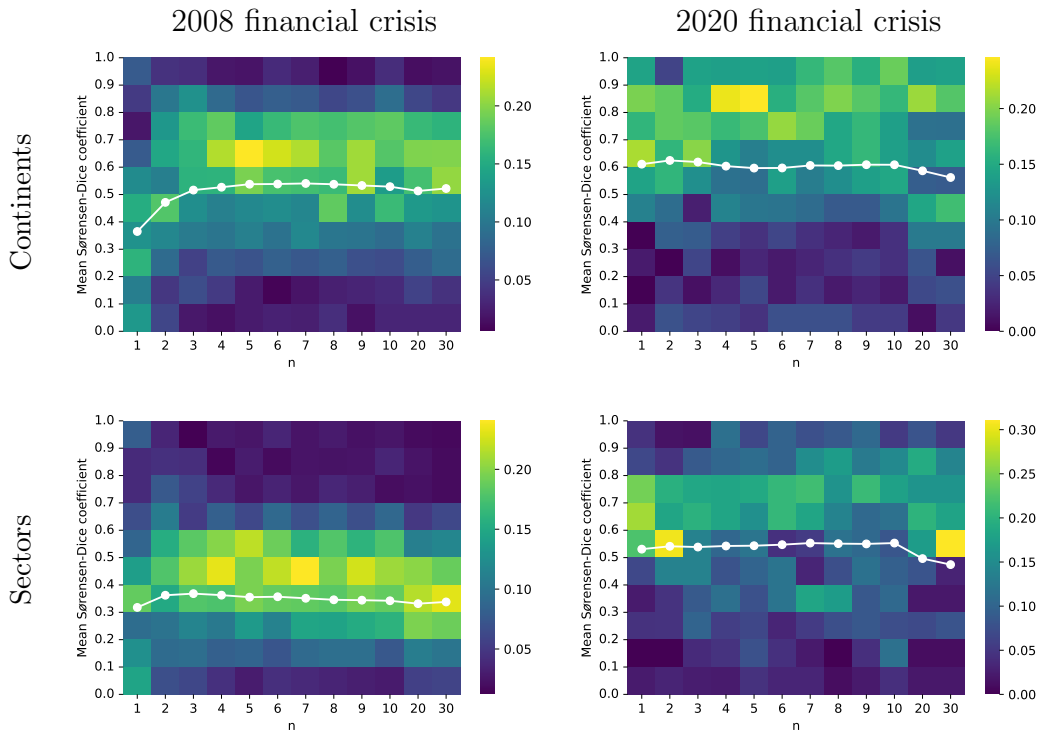

Figure 3.13: Heatmap for the distribution of the mean Sørensen–Dice coefficient between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days. The white dots indicate the mean over all sliding window predictions.

where

$$\text{Precision} = \frac{TP}{TP + FP} \quad \text{and} \quad \text{Recall} = \frac{TP}{TP + FN}.$$

In other words, Recall measures the proportion of actual positive cases correctly identified as positive, while Precision assesses the ratio of correctly predicted positive cases to the total predicted positive cases. The F_1-score combines both Recall and Precision into a single metric to balance the trade-off between correctly identifying positive cases and minimizing false positives. An F_1-score of 1 indicates perfect classification, and a score of 0 indicates that the model does not classify any observation into its correct class.

Later we examine the effect of the window length n on the model's performance, according to the two measures introduced above. Figures 3.20, 3.21

and 3.22 illustrate the mean over all sliding windows of the two different measures for different values of n when the method is applied to the 2008 (left column) or the 2020 (right column) financial crisis for the future $k = 10$, 20 and 30 days, respectively. It can be seen that for both financial crises the Accuracy of the model decreases as n increases. Conversely, the F_1-score is higher for larger window lengths in general (though with similar values being obtained for window choices $4 \leq n \leq 30$). In addition, while the Accuracy for the 2008 financial crisis is approximately 5% higher than the Accuracy in the 2020 financial crisis for all values of n, the F_1-score for the 2008 financial crisis is around 15% lower than the one for 2020 financial crisis. The reason for this difference in performance metrics could be attributed to changes in the characteristics and complexity of the two financial crises, as well as variations in the data available for modelling. In particular, having less data for the 2008 financial crisis may limit the model's ability to capture relevant dynamics and dependencies between the companies, leading to lower predictive accuracy.

Overall, our findings suggest that the optimal window length n depends on the specific performance metric that is being considered. Choosing an appropriate window length can impact the model's ability to predict financial crises. For instance, a longer window, while providing valuable historical context and potentially enhancing the F_1-score, may exhibit lower Accuracy in predicting rapidly evolving crises. As a result, it is essential for practitioners and researchers to consider thoughtfully their goals and prioritize the specific metric that holds the most significance when customizing the window length for their models.

3.4.3 Assessing the importance of the layers

In this section we study the importance of each layer within our model, as defined by $w_{i,j}^{[\alpha]},\ \alpha \in \{1,2,3,4\}$ in (3.4). We compare the performance of six different networks: (i) the full network, comprising the PMFG, continents, sectors and global layer; (ii) the network without PMFG, i.e., $\alpha \in \{2,3,4\}$; (iii)–(v) duplex networks comprising the global layer and one other, i.e., $\alpha \in \{1,4\}$, $\alpha \in \{2,4\}$, or $\alpha \in \{3,4\}$; (vi) the global layer only, i.e., $\alpha = 4$. We remark that the latter corresponds to the homogeneous mixing population case, which assumes that the probability of transmission is the same between all companies.

We first compare in Figures 3.14, 3.15 and 3.16 how the different multiplex networks perform, compared to the global layer only, in predicting the total number of infected companies in the future $k = 10$, 20 and 30 days, respectively. For each of the six networks, we compute the mean absolute difference between the predicted and actual total number of infected companies after k days, averaged over all sliding windows. The calculation is performed across different values of n. To assess the 'improvement' achieved by each network in comparison to the homogeneous mixing population model (which comprises only the global layer), we compute the difference in average absolute discrepancies between the predicted and actual total number of infected companies for each n comparing the global layer network with each of the other five networks. Our results demonstrate that for both financial crises the full network outperforms the other network structures and gives the highest improvement in predicting the total number of infected companies after k days, for all studied values of k, compared to the homogeneous mixing population model. Using the global layer alone (homogeneous mixing population model) gives the least accurate predictions since each other network produces positive im-

provements. Moreover, the second best results are achieved when using the network comprising of the global and PMFG layers, indicating the importance of the PMFG layer. We remark that the six network models display similar accuracy of prediction in the case of new-recoveries, as is to be expected, since the recovery probability is independent of network structure (data not shown).

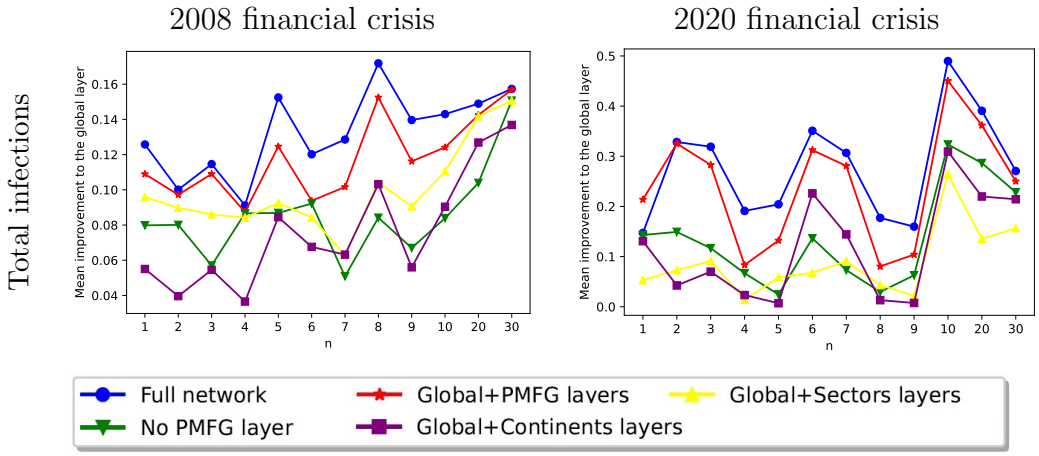

Figure 3.14: Comparison between the total number of infected companies at a prediction horizon of $k = 10$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right columns), using the infections data from the previous n days, in comparison to the homogeneous mixing population model.

We then employ the same methodology as described in Section 3.4.2.3 and compare accuracy of the results obtained from fitting all six models, according to the Sørensen–Dice coefficient. Figures 3.17, 3.18 and 3.19 illustrate the comparison between the mean Sørensen–Dice coefficient over all sliding windows between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 10$, 20 and 30 days, respectively, for both the 2008 (left column) and the 2020 (right column) financial crises. The results demonstrate that for all studied values of k and n, the full model, containing all four layers, consistently yields the highest mean Sørensen–Dice coefficient for predicting both the continents and sectors in which newly infected companies will emerge. Conversely, employing only

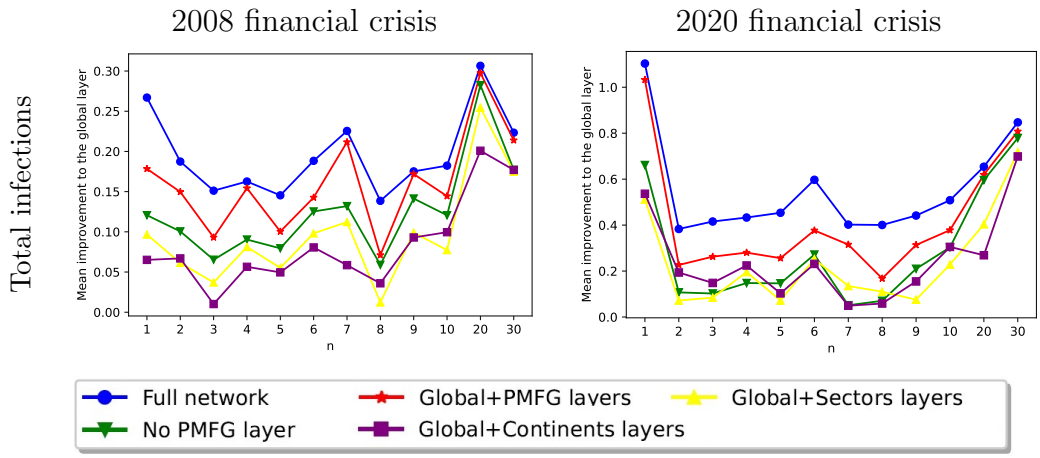

Figure 3.15: Comparison between the total number of infected companies at a prediction horizon of $k = 20$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right columns), using the infections data from the previous n days, in comparison to the homogeneous mixing population model.

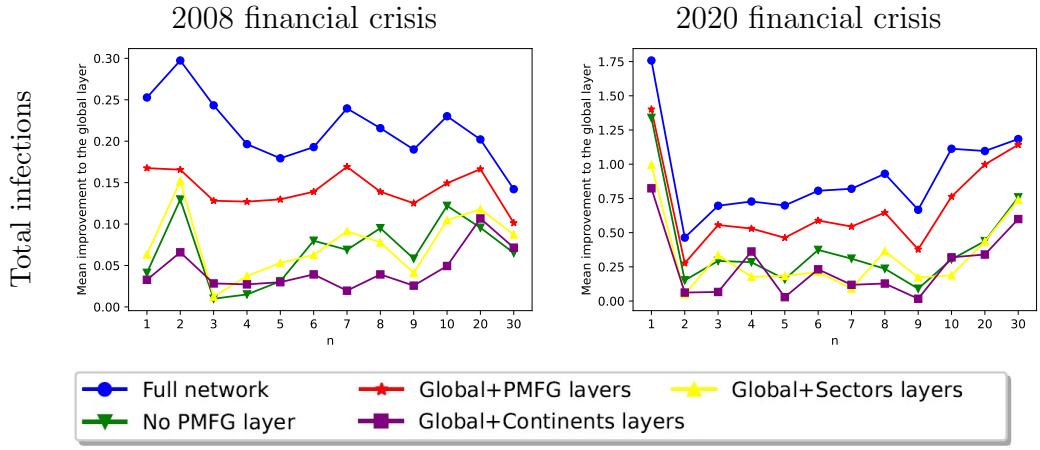

Figure 3.16: Comparison between the total number of infected companies at a prediction horizon of $k = 30$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right columns), using the infections data from the previous n days, in comparison to the homogeneous mixing population model.

the global layer produces the lowest mean Sørensen–Dice coefficients across all combinations of n and k. Furthermore, the second-best results for all combinations of n and k are consistently observed when utilizing the network comprising only the global and PMFG layers. Adding each of the continents and sectors layers, in addition to the global layer, improves the quality of the predictions. This means that each of the layers within our model, and par-

ticularly the PMFG layer, includes information which improves the model's predictive power.

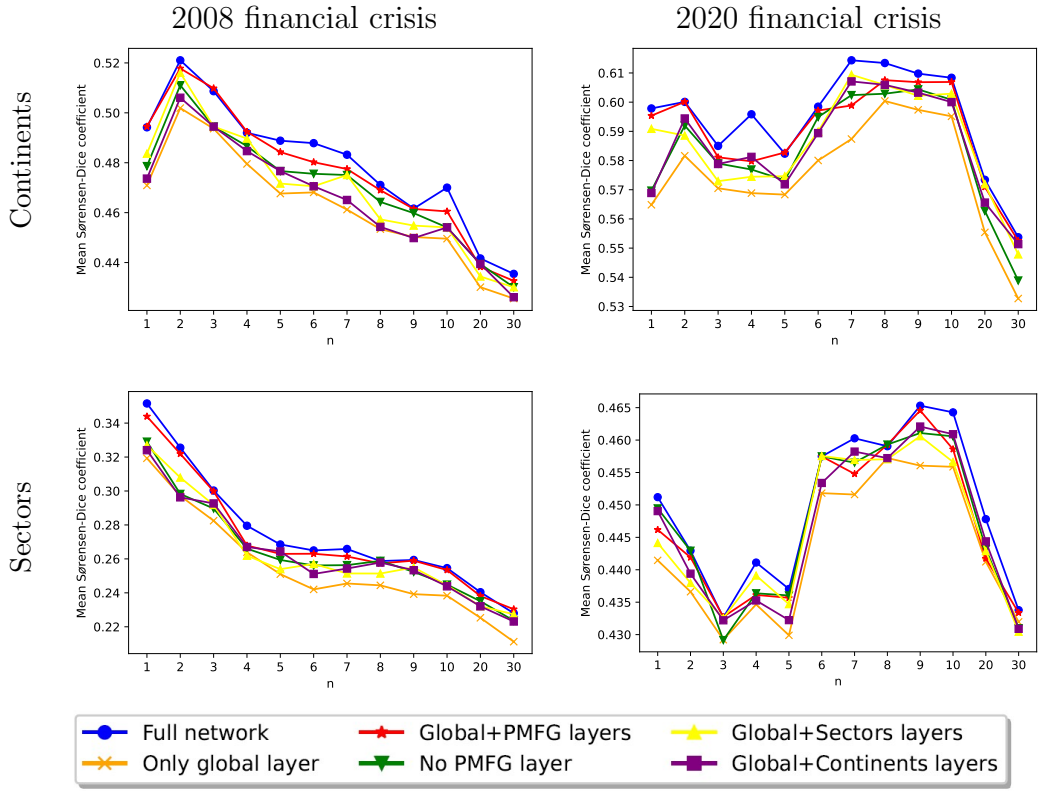

Figure 3.17: Comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the 2008 (left columns) and the 2020 (right column) financial crises, using the infections data from the previous n days when using six different network models.

We then compare the two different scores introduced in Section 3.4.2.4: Accuracy and F_1- score. Figures 3.20, 3.21 and 3.22 illustrate the comparison of the mean scores over all sliding windows for different values of n when $k = 10$, 20 and 30, respectively, for all six networks. The results demonstrate that the full model consistently outperforms the homogeneous mixing population model in all scenarios. In particular, when examining the 2020 financial crisis, the full model's Accuracy surpasses that of the homogeneous mixing population model by nearly 10%.

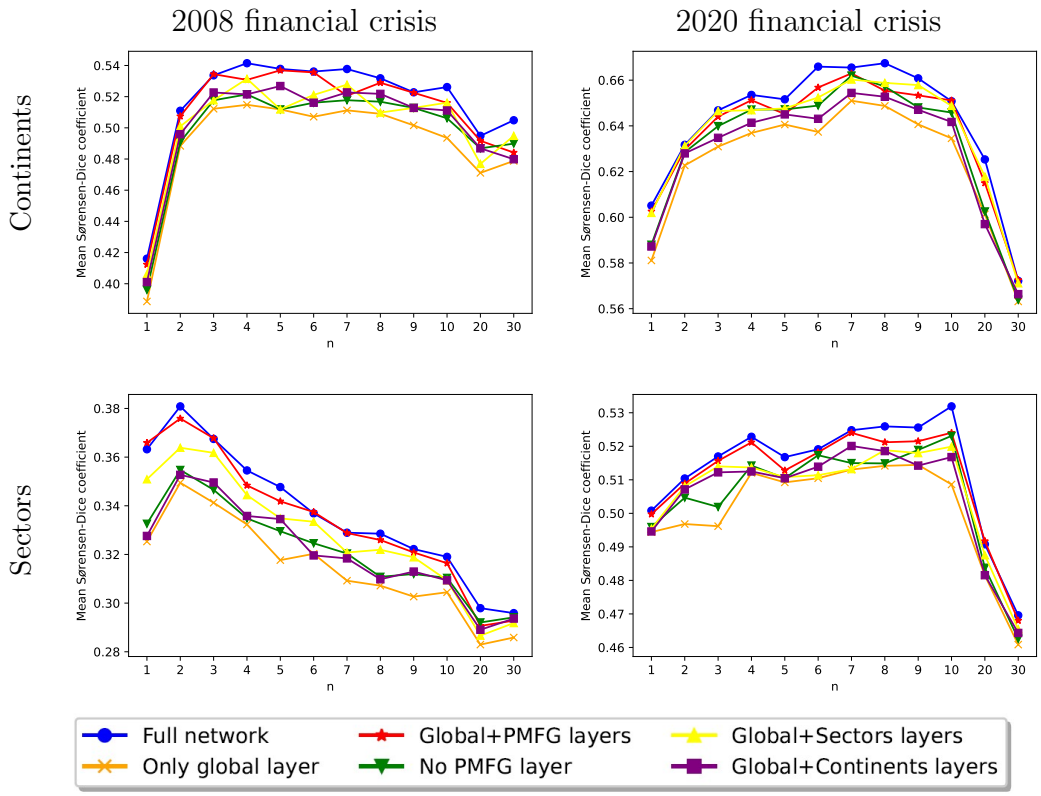

Figure 3.18: Comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the 2008 (left columns) and the 2020 (right column) financial crises, using the infections data from the previous n days when using six different network models.

Similar trends are observed when examining the mean F_1-score. Specifically, in the 2008 financial crisis the mean F_1-score of the full model improves that of the random model by around 5%, while an increase of nearly 10% is observed in the 2020 financial crisis. Furthermore, the network's performance is substantially improved when the model includes both the global and PMFG layers, resulting in the second highest scores. The results are consistent for all values of k, which demonstrates the superiority of the full model over the homogeneous mixing population model, but also highlights the importance of incorporating the PMFG network for achieving more accurate predictions. The maximum F_1-score is achieved when $k = 30$ and for the 2008 financial crisis it reaches the value of 0.08, whereas during the 2020 financial crisis,

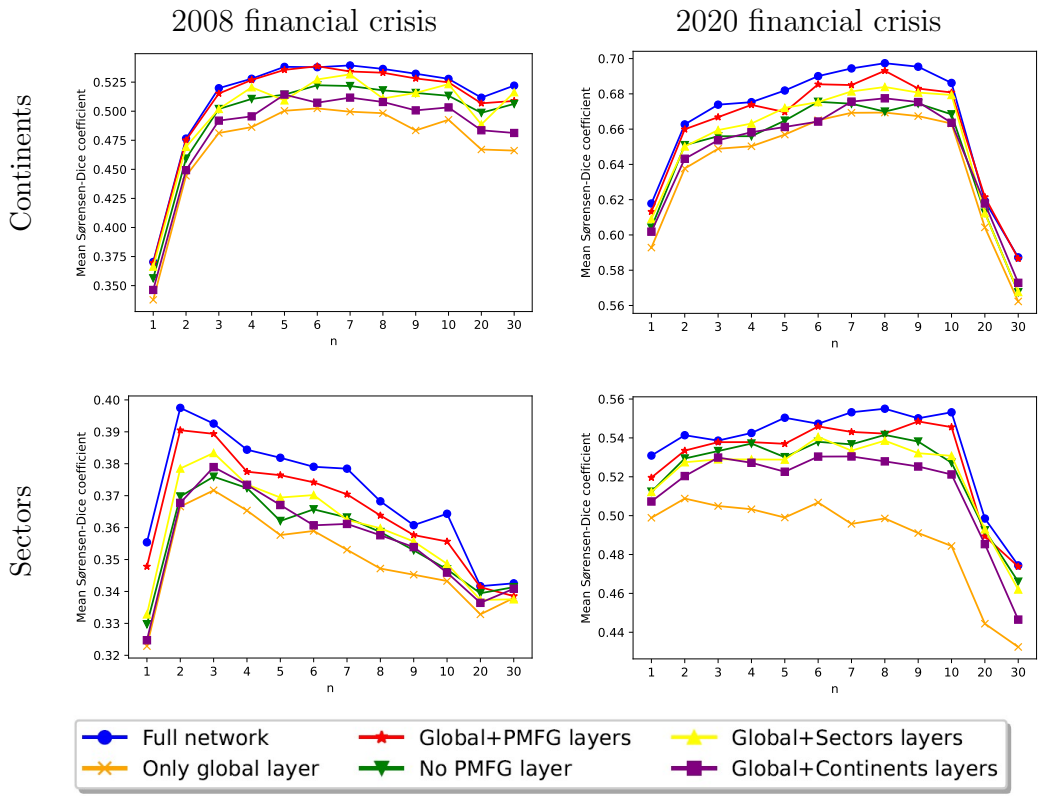

Figure 3.19: Comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the 2008 (left columns) and the 2020 (right column) financial crises, using the infections data from the previous n days when using six different network models.

it reaches 0.24. These values are, of course, too low for practical prediction: our work constitutes a proof-of-concept rather than an immediately applicable method in this context. In addition, the observed improvements in Accuracy and F_1-score suggest that the additional information incorporated through the multilevel structure holds potential for enhancing predictive models in future research efforts.

Finally, we compare the models' computational run-time. We compare both the time for estimating the MLE parameters and also the computational time per simulation. Figure 3.23 shows two plots: the left plot illustrates the average time in seconds per iteration for the MLE estimation for different values of n and the right plot shows the average time per simulation in seconds

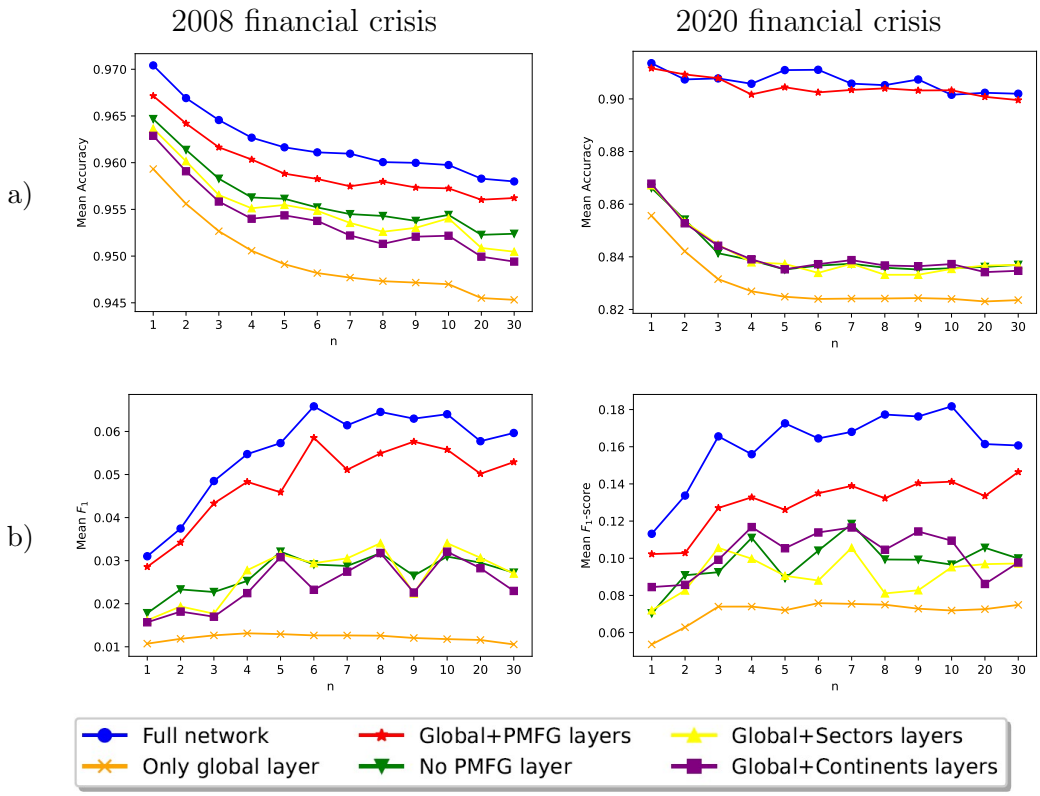

Figure 3.20: Comparison between the mean a) Accuracy and b) F_1-score in the future $k = 10$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right column), using the infections data from the previous n days to predict the set of individual infected companies when using six different network models.

when simulating the SIR model for the future k days for different values of k. It can be seen that in all cases using the full network requires the most computational time, while using only the global layer takes the least time. The computational time is based on the 2020 financial crisis.

To determine which model provides the best trade-off between accuracy and computational time, we define the model's Efficiency as the ratio of its performance, in terms of the relevant accuracy measure (Accuracy or F_1-score), to its computational time[5] (per MLE iteration or simulation). Figure 3.24 illustrates the Efficiency of the six different networks when the MLE computational time (left column) and the time per SIR simulation (right column)

[5]Computations were performed on a machine with an AMD Ryzen 5 processor and 8 GB RAM.

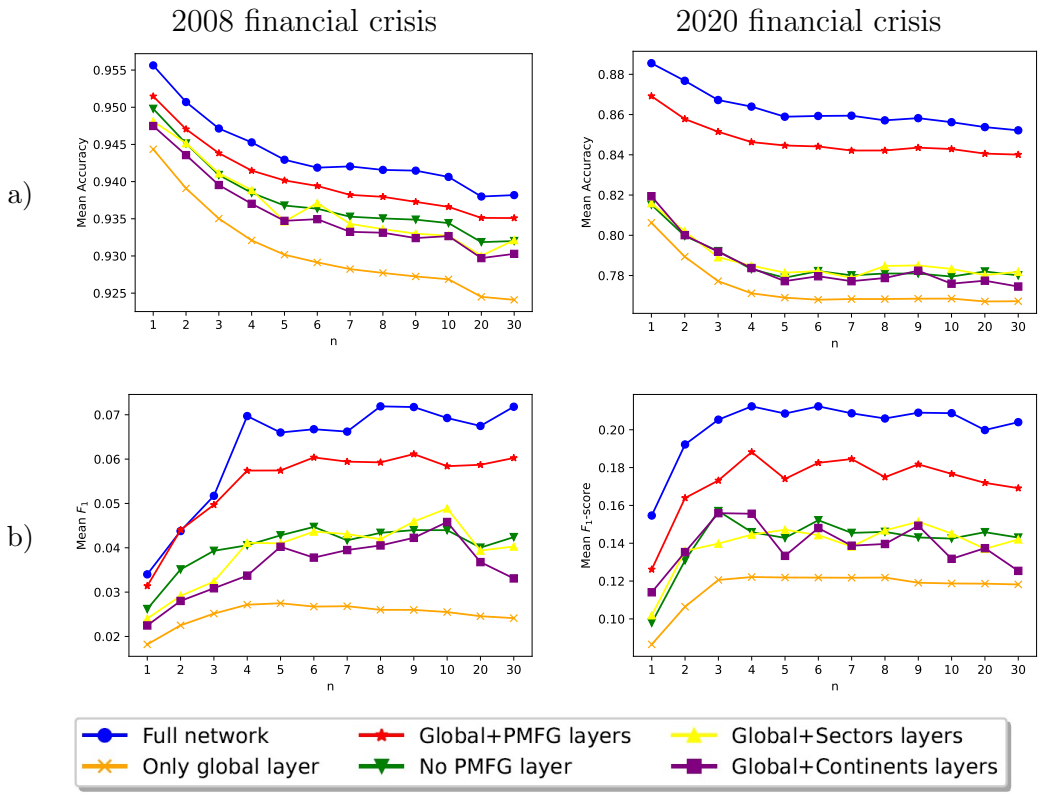

Figure 3.21: Comparison between the mean a) Accuracy and b) F_1-score in the future $k = 20$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right column), using the infections data from the previous n days to predict the set of individual infected companies when using six different network models.

are taken into account, that is mean Accuracy per iteration/simulation and mean F_1-score per iteration/simulation. Note that the computational time of the SIR simulations does not depend on n, just on k, so to estimate the Efficiency we use the mean Accuracy and F_1-score results and the computational time per simulation for $k = 30$.

It can be seen that while if we consider Accuracy (top row) the homogeneous mixing population model shows highest Efficiency score for all n, the second best model for the trade-off between computational time and Accuracy is the model including the global and the PMFG layers. In addition, when we consider the F_1-score, we can see that the model consisting of the global and the PMFG layer yields highest Efficiency. Overall, the optimal choice of model

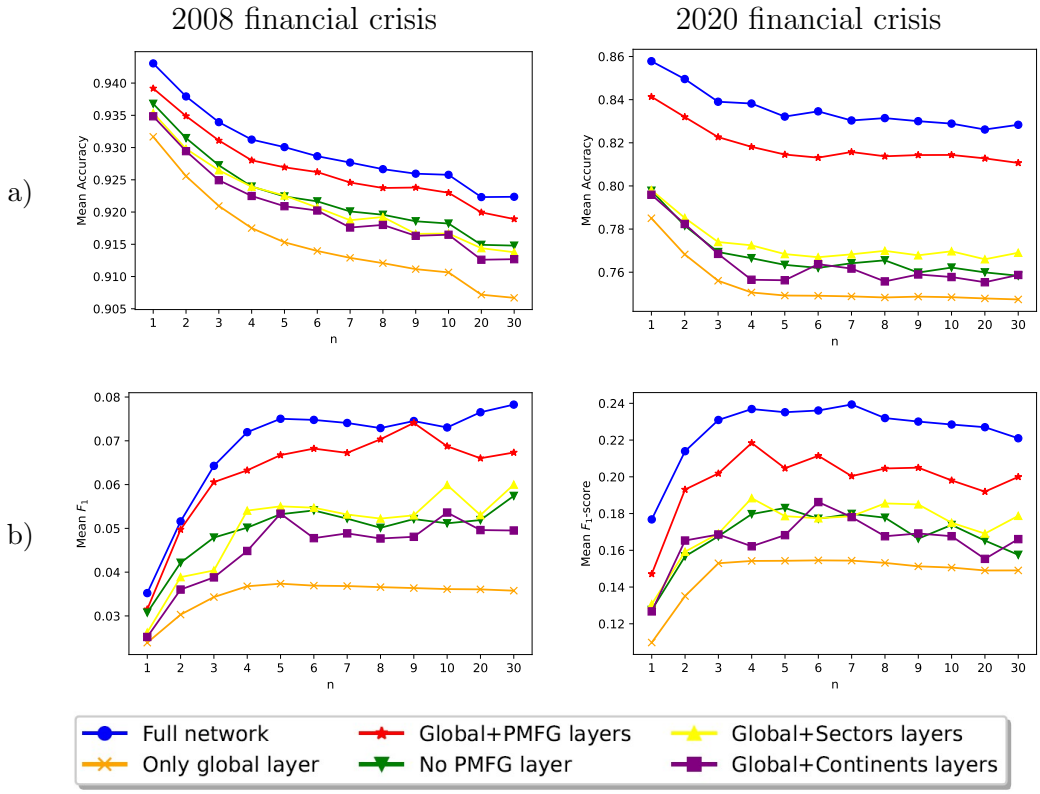

Figure 3.22: Comparison between the mean a) Accuracy and b) F_1-score in the future $k = 30$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right column), using the infections data from the previous n days to predict the set of individual infected companies when using six different network models.

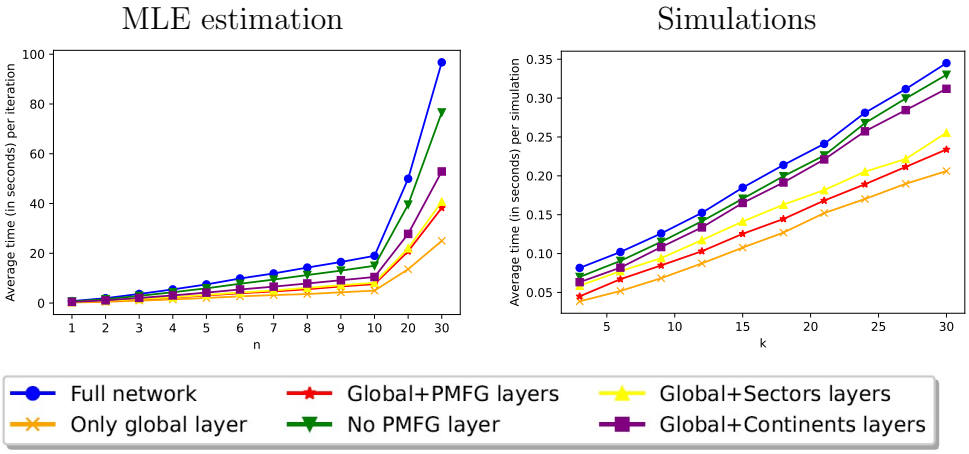

Figure 3.23: The average computational time (in seconds) per iteration for MLE estimation fitting the model to the n previous crisis days (left plot) and per simulation for the future k crisis days (right plot) for the six different networks.

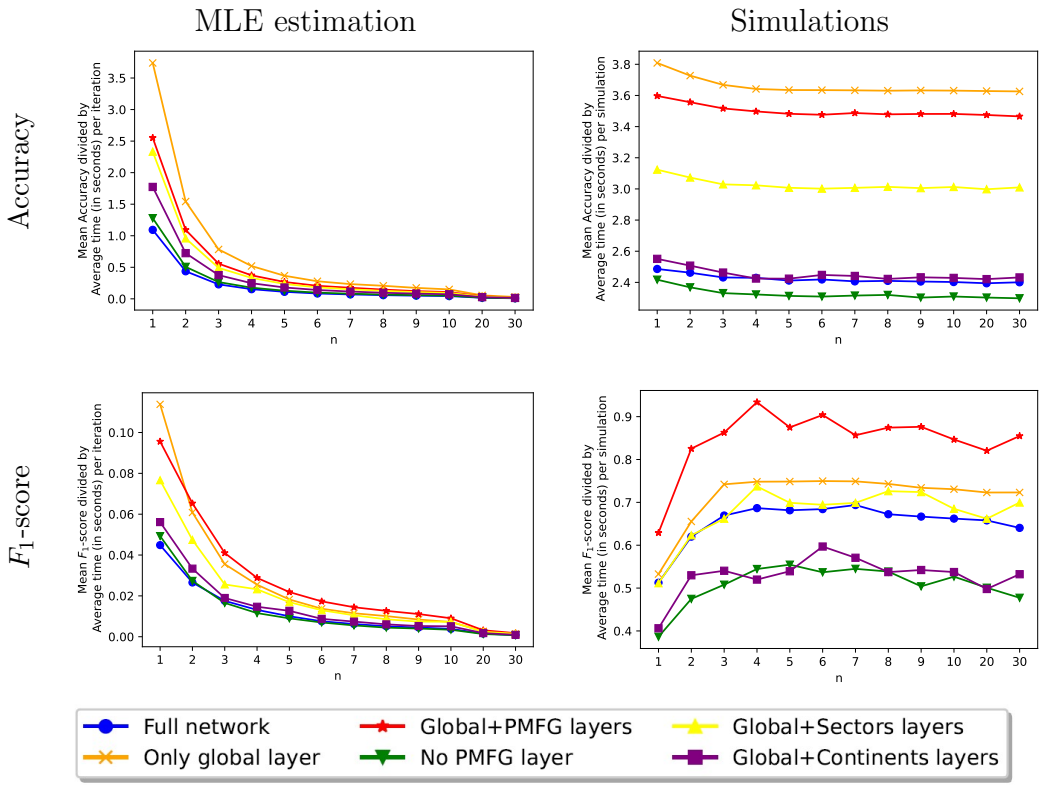

Figure 3.24: Comparison between the mean Accuracy (top row) and mean F_1-score (bottom row) per computational time for MLE iteration (left column) and for one simulation (right column).

depends on whether one prioritizes accuracy or computational efficiency. If the model's predictive power is the main interest, then the full network and the network including the global and the PMFG layers give the best results. In addition, if one is interested in computational efficiency, while preserving high accuracy, the network consisting of the global and PMFG layers might again be the best choice.

3.5 Discussion and conclusion

This chapter proposes a novel framework to analyse the spread of financial crises. We integrate stock price, geographical, and economic sector data to provide a four-layer multiplex network on which a discrete-time SIR model

is simulated, so as to predict the spread of financial risk through interconnections between companies. Specifically, by fitting infection and recovery parameters on each layer of our network to historic stock data through a maximum likelihood approach, we predict future infection dynamics.

We investigate and evaluate the utility of our approach through application to two recent financial crises: the 2008 crisis, initiated by the subprime mortgage market and the 2020 crisis, associated with the COVID-19 pandemic. In each case, we examine the ability of our model to estimate dynamically future infection risk over a horizon of k days, given data from the prior n days. Using a range of accuracy measures, we analyse the dependence of prediction accuracy on k and n, in terms of total number of infections, as well as sector- and location-specificity. Thereby we demonstrate that interactions among companies within and across sectors and continents in the financial network play a substantial role in the spread of financial crises and their incorporation into the model improves the prediction of future outbreaks of financial distress. By comparison with a homogeneous mixing assumption in particular, we highlight the importance of understanding and accounting for the complex interdependencies between companies in financial systems for risk prediction.

While our model offers valuable insights into the spread of financial crises, it is essential to recognize its limitations and constraints. One significant limitation of our model is its reliance on historical stock price data, which only gives an incomplete view of the financial stability of a company. The accuracy and reliability of our predictions heavily depend on the availability and quality of the data, which may vary across different companies, sectors, and regions. Moreover, our model operates under several basic assumptions, such as the division into susceptible, infected, and recovered companies. While these assumptions simplify the complexity of financial contagion dynamics, they also impose constraints on the model's applicability and may not fully

capture the nuances of real-world scenarios. Despite incorporating multiple layers representing stock prices, geographical locations, and economic sectors, our model overlooks other potentially important factors influencing financial contagion, such as macroeconomic indicators, regulatory policies, investor sentiment, and systemic risk factors. We emphasize that the primary goal of this research is to present our novel framework combining extreme value theory, financial network construction and SIR modeling for the spread of financial risk in networks rather than to undertake comprehensive prediction. Despite its limitations, the incorporation of the multilevel network structure has shown potential in enhancing prediction power and capturing the interdependencies among companies driving financial contagion dynamics.

Overall, our results suggest that the proposed framework is effective in predicting risk spread, this information potentially being useful in terms of risk prevention and mitigation. In particular, our predictive approach updates in real time as new data become available, making it potentially useful in crisis management as the crisis unfolds. Moreover, the framework is universal in nature and hence could be applied to analyse any contagious financial crisis, not only in a network of companies but also in a network of countries and economic sectors. In addition, our results agree with the existing research which consistently shows the importance of including economic sector [147] and geographical location [53] information in predicting a company's future performance. However, [86] suggests that industry-level analysis may not always provide a quantity of information that results in substantial improvement of future profitability, therefore suggesting that efficiently incorporating further, more granular information might be difficult.

As we highlight in our analyses, predictive accuracy suffers during periods of rapid change. We note, however, that such events are typically associated with extrinsic factors, such as government financial stimulus packages or lockdown

periods. Natural future work includes incorporating such significant changes into the model. In the next chapters, we address this by using a change point detection approach, as well as observations weighting approach, where more recent data have higher importance than an older data.

Chapter 4

Enhancing financial crisis prediction: Integrating change point detection for exogenous event identification

As noted previously, our model's predictive accuracy suffers during and shortly after significant events, such as lockdowns or major economic policy changes like stimulus packages. This chapter explores the integration of change point detection (CPD) techniques to improve the model's adaptability to major market events. Exogenous shocks, such as geopolitical tensions and natural disasters, can lead to substantial changes in stock prices and market dynamics. By implementing the ChangeFinder [208] algorithm for change point detection, we aim to identify and incorporate such events, so that the model can adapt quickly. Once a change point is detected, the model adjusts by re-estimating its parameters based on the data following the detected change. The process involves continuously monitoring the data for anomalies. When a

change point is identified, the model discards the older data before the change point and focuses on the more recent data. Compared to the model introduced in the previous chapter, incorporating change points enhances forecasts in predicting the number of affected companies, as well as in company-, sector-, or location-specific predictions. The results suggests that incorporating change points into the model generally yields better performance, particularly in improving the predictions of the number of infected companies.

4.1 Introduction

Detecting significant changes in financial time series data is a critical aspect of quantitative analysis, especially in understanding the dynamics of market behavior and mitigating financial risks. Financial markets are highly susceptible to various exogenous and endogenous events that can cause abrupt shifts in asset prices, volatility, trading patterns and investor decisions [57]. Accurately identifying these changes is critical for investors, financial institutions and policymakers as it allows them to adjust their approaches in response to changing market circumstances.

Many events during the 2008 and the 2020 financial crises had significant effects on the financial markets. For instance, in the 2008 financial crisis, governments worldwide implemented various decisions and interventions to stabilize financial markets and mitigate impact of the crisis. In the United States, central banks, including the Federal Reserve, implemented aggressive interest rate cuts to stimulate economic activity and ease borrowing costs [26]. Lower interest rates aimed at encouraging spending, investment and borrowing to stimulate economic growth [213]. During the 2020 financial crisis, government decisions regarding lockdowns, travel restrictions and quarantine

measures significantly impacted various industries, leading to market disruptions, supply chain interruptions and shifts in consumer behaviour. Sectors like travel, hospitality, entertainment and retail experienced substantial declines in revenue during the crisis, impacting stock prices and overall market performance [96]. In response to the market tension, governments globally introduced massive fiscal stimulus packages, providing financial support to individuals, businesses and industries severely affected by lockdowns and economic disruptions [233].

Other events, such as wars and geopolitical tensions [47], natural disasters [87] and protests [56] can also influence the financial market significantly. For example, the terrorist attacks on September 11, 2001, led to a temporary closure of financial markets in the United States [4, 52, 167, 207]. When the markets reopened, there was a significant drop in stock prices. The attacks led to increased volatility in the global financial market, particularly affecting airline-, insurance- and tourism-related stocks. In addition, following the earthquake and tsunami in Japan in 2011, the Fukushima nuclear disaster led to widespread disruptions in Japan's economy [17]. Japanese stocks and the yen experienced severe volatility as concerns about nuclear safety, economic damage and supply chain disruptions appeared. Industries such as technology, automotive and manufacturing were particularly affected by disruptions in the supply chain [123]. Brexit, the United Kingdom's decision to leave the European Union (EU), also had a substantial impact on the financial markets [112]. The uncertainty surrounding Brexit negotiations led to significant fluctuations in the value of the British pound (GBP). In addition, UK-focused stocks, particularly those reliant on EU trade or with substantial exposure to the domestic economy, were especially sensitive to Brexit developments [34].

These examples demonstrate how non-financial events such as wars, natural disasters, geopolitical tensions and pandemics can have profound effects on

financial markets, leading to increased volatility, changes in asset prices, as well as disruptions across various industries and sectors. In this chapter we enhance the model introduced in Chapter 3 by incorporating change point detection techniques. The approach introduced in this chapter enables the dynamic integration of significant events in real time, contributing to a more robust and adaptive model. By utilizing change point detection, we can automatically identify and adapt to these significant changes as soon as they happen, allowing our model to incorporate their impact. This approach enables us to capture and account for the effects of these events as they occur, enhancing the model's ability to adapt to evolving market conditions and improve its predictive capabilities in a dynamic and automated manner.

The remainder of the chapter is organized as follows. In Section 4.2 we introduce the ChangeFinder algorithm, used for online change point detection. Then in Section 4.3 we incorporate change point detection into the model using the ChangeFinder algorithm to detect significant changes and adapt to them in real time, and then compare these predictions with the model introduced in Chapter 3. Finally, in Section 4.4 we discuss our results, the limitations of our research and provide directions for future developments.

4.2 Change point detection

Change point detection is a technique for identifying the points at which the probability distribution of a time series change. Change point detection algorithms are usually classified as 'online' or 'offline'. The 'offline' algorithms take into account the whole available data series at once, where the main goal is to detect all change points. 'Online', or real-time, algorithms, on the other hand, operate concurrently with the activity they are monitoring, analysing

each data point as it becomes available with the aim of identifying a change point as soon as possible after it occurs, ideally before the next data point arrives. In practice, no change point detection algorithm can actually work in real time, since it must first examine the new data before determining whether a change point happened between the old and new data points. In our study we utilize an online change point detection algorithm, called ChangeFinder [208].

4.2.1 ChangeFinder

Change point detection has been employed in statistical analysis and data mining for a number of functions, including trend analysis [12], edge detection in image processing [217] and anomaly detection [195]. Since the autoregressive (AR) models constitute the simplest widely-used class of models to represent time series, we begin this section by describing a typical change point detection method based on the AR model.

4.2.1.1 AR Model

Autoregressive (AR) models are an essential concept in time series analysis and forecasting. They are widely used in many fields, including economics and finance [190]. AR models represent the relationship between an observation and previous observations in a time series. The fundamental principle is that the current value of a time series can be expressed as a linear combination of its past values and some random noise.

Let $(z_1, ..., z_n)$ be a time series with mean 0. Then, the k-th order AR model,

denoted by $AR(k)$ can be defined as

$$z_t = \sum_{i=1}^{k} \alpha_i z_{t-i} + \epsilon_t, \tag{4.1}$$

where α_i are the regression parameters, which indicate the degree with which z_t is correlated with z_{t-i}, and ϵ_t are independent identically distributed (iid) random variables generated from a Normal distribution with mean 0 and variance σ^2. Let $\mathbf{x}_{1:n} = (x_1, ..., x_n)$ be the time series that we observe, where $x_t = z_t + \mu$. Let $\mathbf{x}_{1:t} = (x_1, ..., x_t)$ and $\mathbf{x}_{t+1:n} = (x_{t+1}, ..., x_n)$ denote the time series before (and including) and after time t, respectively. Then, assuming the k-th order AR model, the conditional probability density function of x_t is given by

$$P(x_t|\mathbf{x}_{t-k:t-1}, \theta) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x_t - \omega_t)^2}{2\sigma^2}\right]. \tag{4.2}$$

Here, ω_t is given as follows:

$$\omega_t = \sum_{i=1}^{k} \alpha_i (x_{t-i} - \mu) + \mu, \tag{4.3}$$

and $\theta = (\alpha_1, \alpha_2, ..., \alpha_k, \mu, \sigma)$ are the model parameters. Then, if $\hat{\omega}_t$ is the estimated value of ω_t obtained by substituting the estimated model parameters $\hat{\theta} = (\hat{\alpha}_1, ..., \hat{\alpha}_k, \hat{\mu}, \hat{\sigma})$ in (4.3), the model-fitting error for $\mathbf{x}_{1:n}$ is given by

$$I(\mathbf{x}_{1:n}) = \sum_{t=1}^{n} (x_t - \hat{\omega}_t)^2. \tag{4.4}$$

Here, the model fitting error $I(\mathbf{x}_{1:t})$ is calculated using maximum likelihood on the data only from time 1 to time t, i.e. the parameter θ is estimated given the observations sequence $\mathbf{x}_{1:t}$. Hence, if $I(\mathbf{x}_{1:t}) + I(\mathbf{x}_{t+1:n})$ is significantly smaller than $I(\mathbf{x}_{1:n})$, then t is a change point. In other words, the procedure assesses if it is significantly better to fit two separate models, one from time 1 to t and the other from time $t+1$ to n, than a simpler full model from time 1 to time n.

Although the algorithm is simple, it has a computational complexity of $O(n^2)$, so it is impractical for online change point detection, especially for large time series, where data arrives sequentially and the algorithm needs to process it in real-time [116]. Moreover, in the above formulation we assume that we have only one change point, with time series being stationary before and after it. However, this assumption can be overly simplistic and unrealistic. The SDAR model, which we introduce in the following section, deals with the situation in which we have multiple change points, with stochastic processes between them changing slowly [208].

4.2.1.2 SDAR Model

The SDAR (Sequentially Discounting Auto Regression model learning) algorithm can be used to address the computational problem mentioned above [208]. The SDAR algorithm is employed for adapting and discounting data in an online fashion, utilizing the AR model. For an estimate θ_t of θ given $\mathbf{x}_{1:t}$, the SDAR algorithm is employed to calculate in an online manner the value of θ that maximizes the following quantity:

$$\sum_{i=k+1}^{t} (1-r)^{t-i} \log P(x_i|\mathbf{x}_{i-k:i-1}, \theta), \tag{4.5}$$

where $r \in (0, 1)$ is a discounting factor. A smaller r indicates a stronger impact from past data. For values of $t < k + 1$, default values of θ_t are assumed. As mentioned in [116], incorporating the discounting of older data, the SDAR algorithm is useful for real-time learning with non-stationary time series data. Furthermore, its computation cost is $O(n)$, making it a favourable choice for online change point detection.

4.2.1.3 ChangeFinder algorithm

The ChangeFinder algorithm utilizes the SDAR algorithm to detect change points in real time. It consists of a two-step learning process which combines the detection of outliers and change point detection in a time series. The first learning step of the algorithm uses the SDAR method to calculate an outlier score for each data point, which shows how much it deviates from the learned model, where a higher score suggests a higher likelihood of being an outlier (see Step 2 in Algorithm 3). In more detail, a series of probability density functions $\{P_t = P(x_t|\mathbf{x}_{t-k:t-1})|t = k+1, k+2, \dots\}$, learned incrementally every time a new data point x_t is input, are constructed by the ChangeFinder algorithm as follows. Firstly, estimates θ_t, of θ, given $\mathbf{x}_{1:t}$, are obtained by maximizing equation (4.5), where P is defined in (4.2). For each input x_t $(t \geq k+1)$ an outlier score $S(x_t)$ is then calculated by taking the logarithmic loss score $S(x_t) = -\log P(x_t|\mathbf{x}_{t-k:t-1}, \theta_{t-1})$. Then the outlier scores over a fixed-length window T are calculated and as the window moves along the data, the scores within each window are averaged. This process generates a new time series $\{y_t\}$ of these averaged scores. In this way the influence of isolated outliers in the time series is reduced (see step 3 in Algorithm 3). The second learning step involves repeating the same procedure on the new time series $\{y_t\}$. That is, calculate an outlier score $S(y_t)$ for each y_t by using the SDAR method to estimate θ_t. Then estimate the change point or anomaly score, showing the probability of a change point, by taking the average over a sliding window with fixed length T (see step 4 in Algorithm 3). Figures 4.1 and 4.2 illustrate a flowchart and an illustration of the two learning phases of the ChangeFinder algorithm, respectively. The step-by-step method is described in Algorithm 3.

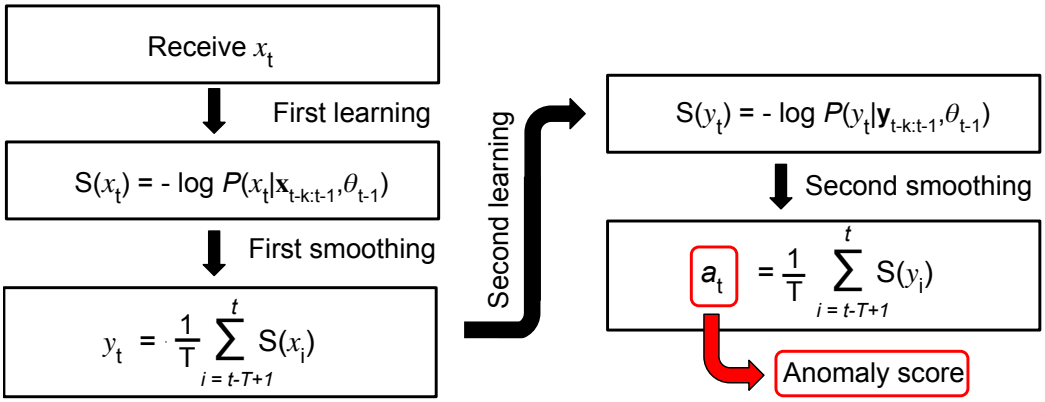

Figure 4.1: Flowchart of ChangeFinder. Adopted from [116].

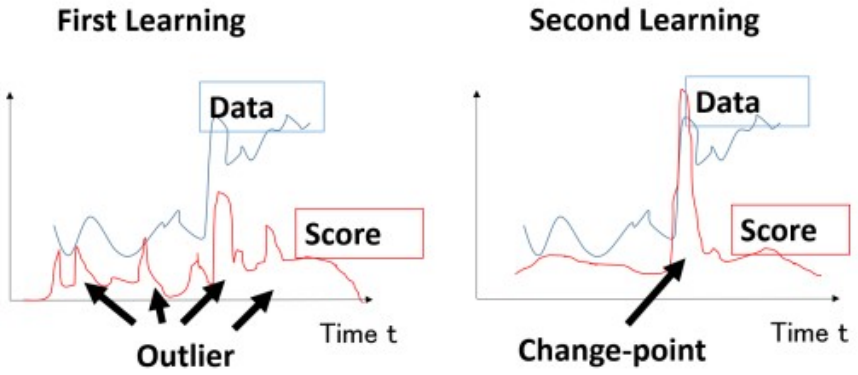

Figure 4.2: Two phase learning ChangeFinder. Image taken from [116]. The plot on the left illustrates the sequence $\{y_t\}$, as defined in equation (4.7), while the plot on the right illustrates the anomaly scores $\{a_t\}$, as defined in equation (4.9).

Algorithm 3 ChangeFinder algorithm

1. Receive a new data point x_t at time t.

2. For each $t \geq k+1$, the outlier score $S(x_t)$ at x_t is calculated,

$$S(x_t) = -\log P(x_t|\mathbf{x}_{t-k:t-1}, \theta_{t-1}), \tag{4.6}$$

   where the values of θ_{t-1} are estimated using the SDAR method, maximizing equation (4.5), with P given in (4.2). Here, a higher score $S(x_t)$ indicates that x_t is an outlier with higher probability.

3. For each t, a rolling mean of the outlier scores within a time window with fixed length T is computed. In other words, a sequence of moving averages of the outlier scores y_t for $t = T+1, T+2, \ldots$ is computed:

$$y_t = \frac{1}{T} \sum_{i=t-T+1}^{t} S(x_i). \tag{4.7}$$

   Dummy values are chosen for y_t, where $t \leq T$. This process is needed in order to reduce the influence of isolated outliers contained in the time series $(x_1, ..., x_t)$.

4. For each t, an outlier score $S(y_t)$ at y_t is calculated,

$$S(y_t) = -\log P(y_t|\mathbf{y}_{t-k:t-1}, \theta_{t-1}), \tag{4.8}$$

   where θ_{t-1} is estimated using the SDAR method, maximizing equation (4.5), with P given in (4.2).

5. A smoothing step is applied to get a sequence of moving averages a_t for $t = 0, 1, 2, ...$:

$$a_t = \frac{1}{T} \sum_{i=t-T+1}^{t} S(y_i). \tag{4.9}$$

   Here, a_t represents the change point or anomaly score at time t. A higher value of a_t indicates a higher possibility of change point.

4.3 Incorporating change points into the model

In this section we focus on implementing the ChangeFinder algorithm within our model framework. Our goal is to identify and assimilate abrupt shifts or significant events in real-time data and adjust the model's parameters accordingly by considering only the data after a change point is detected and disregarding the data before. By doing so, we allow the model to identify and react to significant shifts, such as market crashes or other economic events, ensuring timely updates and more reliable predictions. The results, shown in Section 4.3.3, indicate that compared to the model without change point detection, our enhanced model demonstrates improved predictive accuracy. The predictions are especially improved when predicting the number of infected companies. In particular, while the model without change point detection performs poorly during and after significant events, incorporating change points into the model significantly enhances the results around these critical moments. In the next section we show how we can incorporate such significant changes into the model, yielding higher prediction accuracy.

4.3.1 Change point detection

In this section we incorporate the ChangeFinder algorithm into our model. Let $I(t)$ denote the total number of infected companies on day t in each of the two crises. To detect significant changes in the probability of infection and incorporate them into the model, we examine the time series $\{I(1), I(2), ..., I(L)\}$, where L is the length of the studied crisis in days. We fit the ChangeFinder algorithm with order $k = 1$, discounting rate $r = 0.5$ and rolling window with length $T = 7$ for the smoothing steps, to both the 2008 and the 2020 financial crises. In particular, the parameter $k = 1$ is chosen because, firstly, it

keeps the computational burden low. Secondly, with $k = 1$ the model adapts quickly to sudden changes, or in other words, if a sudden spike or drop in the time series happens, the ChangeFinder algorithm can detect it with minimal latency. In addition, when T is small, outliers and change points can be detected immediately after they appear, but might be difficult to discriminate from one another. On the other hand, if T is large, it leads to time delay for detecting change points, but outliers are filtered. Choosing the parameter $T = 7$ strikes a balance between time delay and outlier filtering. Finally, a moderate r value balances between sensitivity to changes and noise in the results.

For completeness, Figures 4.3 and 4.4 present the anomaly scores for different combinations of the ChangePoint parameters $r \in \{0.1, 0.5, 0.9\}$ and $T \in \{3, 7, 30\}$, when the algorithm is run on the 2008 and the 2020 financial crises, respectively. The solid red vertical lines illustrate points with high probability of being a change point, whereas the other vertical lines correspond to the significant events that happened during each of the two crises. It can be seen that the smaller the window size T and the larger the discount rate r, the noisier the anomaly scores get. In addition, when $T = 30$ we lose important data from the beginning of the crisis as the ChangeFinder algorithm is then not able to detect changes in the first 30 days of the crisis. Also, in the 2020 financial crisis, when $T = 3$ or when $r = 0.9$ it can be seen that in all cases there is no change point detected after governments offer stimulus packages (vertical red dotted line) and after the significant increase of COVID-19 cases (vertical green line). Hence, the value of $r = 0.5$ is chosen such that reduces the noise in the anomaly scores while still allowing for the detection of critical events.

Using the values $k = 1$, $r = 0.5$ and $T = 7$ we obtain the results shown in Figures 4.5a and 4.5b, for the 2008 and 2020 financial crises, respectively. The

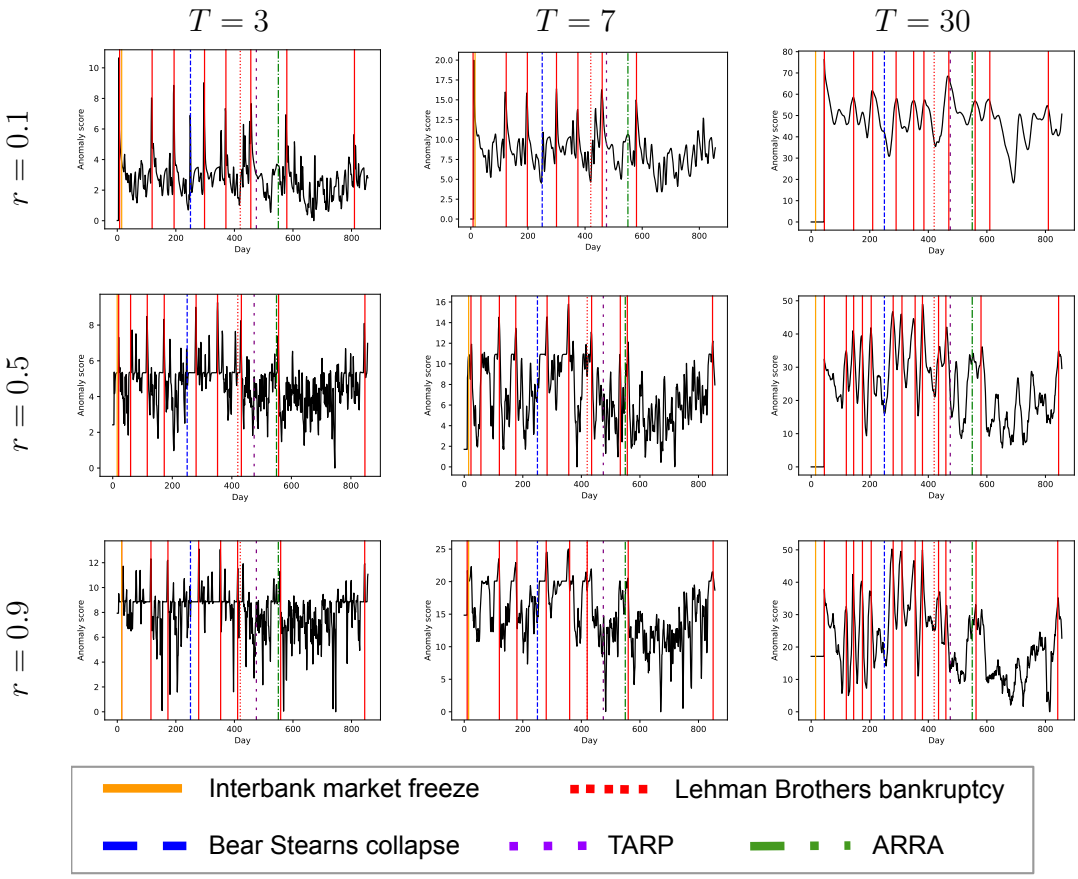

Figure 4.3: The figures illustrate the anomaly scores output, using the ChangeFinder algorithm with different combinations of the discounting rate r and the rolling window size T. The ChangeFinder algorithm is run on the 2008 daily infections data. The solid red vertical lines illustrate points with high probability of being a change point, whereas the other vertical lines correspond to the significant events that happened during the crisis.

top plot in each figure shows the total number of infected companies in black and the solid red vertical lines indicate the days where a high probability of a change point is detected by the ChangeFinder algorithm, or in other words the days with highest anomaly scores, illustrated in the bottom plot in each figure. The other coloured vertical lines correspond to significant events that happened during each crisis. In both figures it can be seen that most of the time shortly after the real events, a high probability of a change point has been detected. This demonstrates the effectiveness of the ChangeFinder algorithm in identifying significant changes in the data.

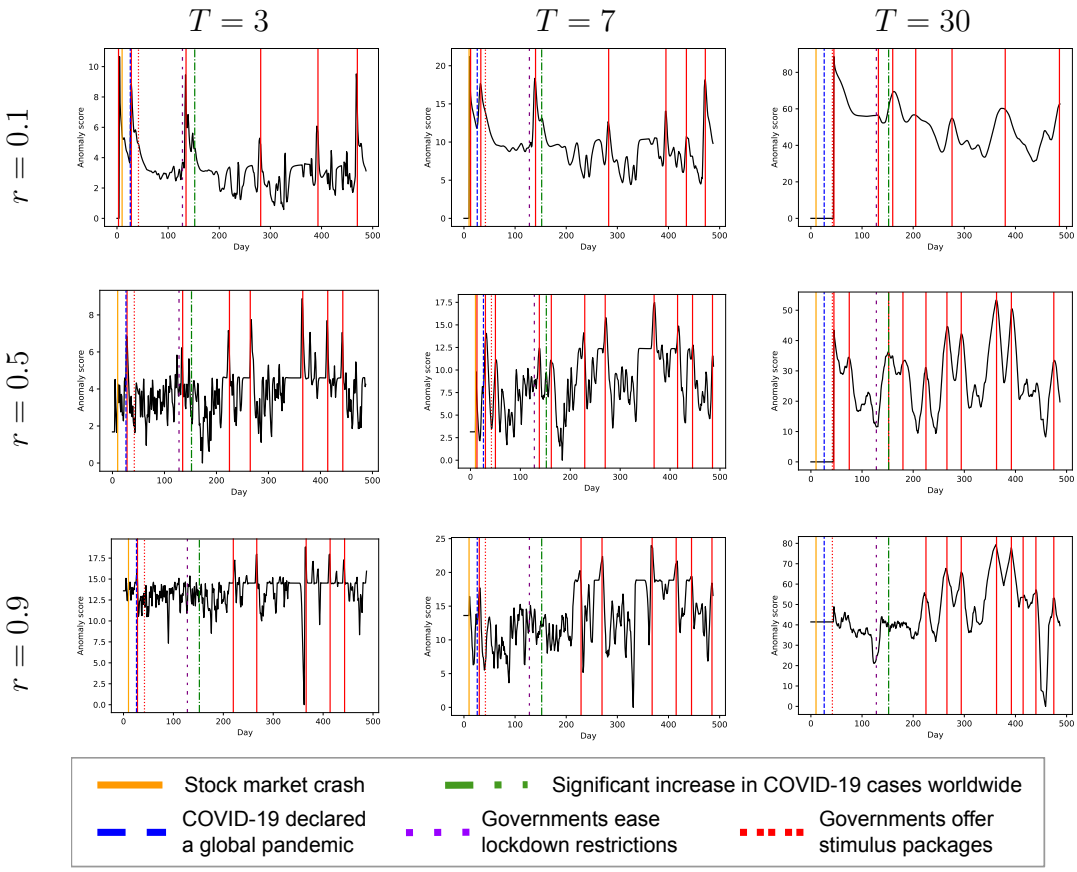

Figure 4.4: The figures illustrate the anomaly scores output, using the ChangeFinder algorithm with different combinations of the discounting rate r and the rolling window size T. The ChangeFinder algorithm is run on the 2020 daily infections data. The solid red vertical lines illustrate points with high probability of being a change point, whereas the other vertical lines correspond to the significant events that happened during the crisis.

4.3.2 Reestimating the parameters

While in the previous chapter we studied the effect of the sliding window size on the prediction accuracy, here we introduce a way to reestimate the model's parameters immediately after a change point is detected. In the previous chapter we fitted the model to the previous n crisis days for different values of n. However, here, utilizing the ChangeFinder algorithm, we can detect a significant change in a time series a day after it occurs. Instead of fitting the model to all the previous n crisis days, if a change point is detected during the past n days of a crisis, we discard the data before the change

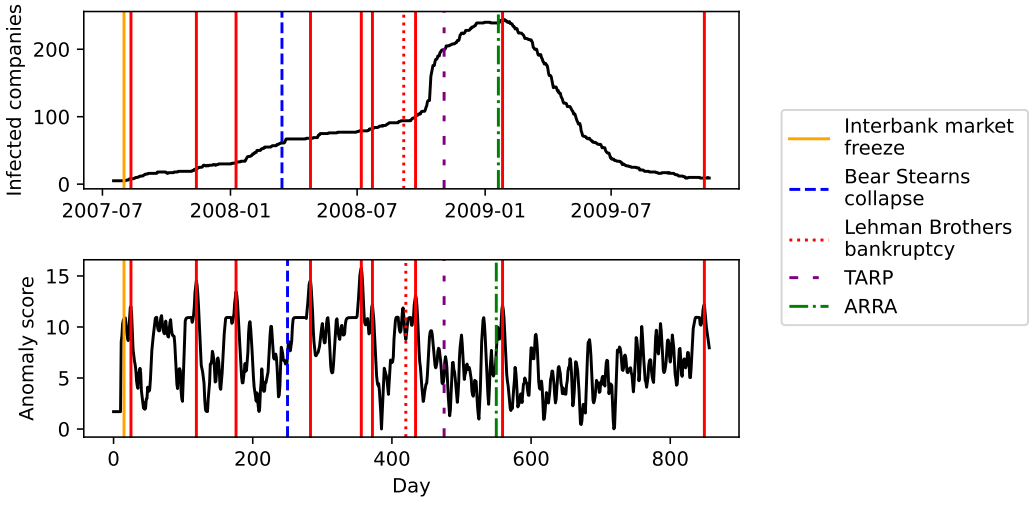

(a)

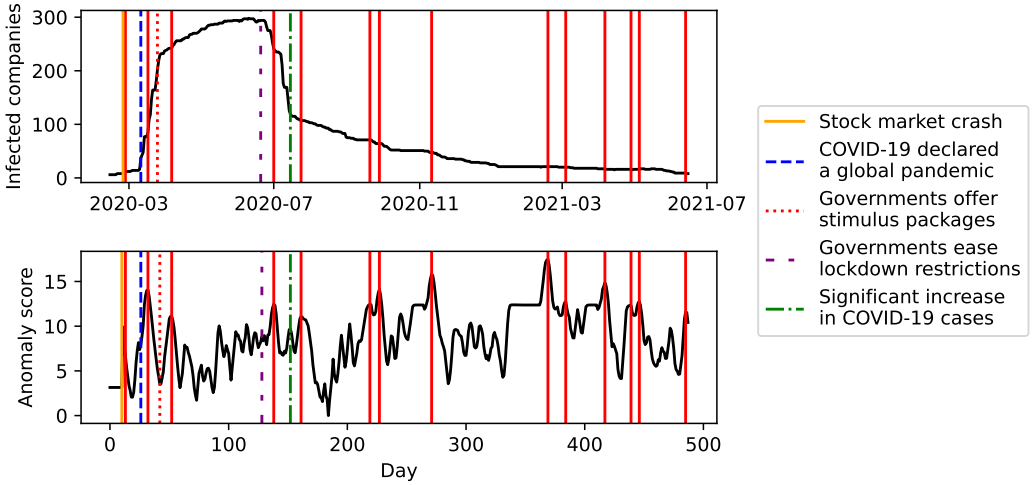

(b)

Figure 4.5: Change point detection on the total number of infected companies in the a) 2008 and b) 2020 financial crises, respectively. The bottom plots illustrate the anomaly scores, as defined in (4.9). The solid red vertical lines show the highest anomaly scores, while the other coloured vertical lines indicate some of the most important events during both crises. It can be seen that often change points (corresponding to high anomaly scores) are detected shortly after a real-life event has happened.

point has happened and fit the model exclusively to the days following (and including) the change point, utilizing Algorithm 4. If more than one change point is detected during the past n days, we discard all the data before the last (most recent) change point. In the following section we evaluate the model incorporating change points (we henceforth denote this as the change point model) and compare its performance with that of the model from the previous chapter.

4.3.3 Evaluating the change point model

In this section we compare the model introduced in Chapter 3 with the change point model. We first compare in Figures 4.6, 4.7 and 4.8 how the change point model performs, in comparison with the original model from Chapter 3, in predicting the total number of infected companies (top row), the number of newly infected companies (middle row) and the number of newly recovered companies (bottom row) in the future $k = 10$, $k = 20$ and $k = 30$ days, respectively, for different values of n. The results reveal that the model incorporating the change point detection algorithm (shown in orange) consistently outperforms the model without the change points (shown in blue), with higher improvement for higher values of n. Notably, when $n = 1$, both models yield similar performance, since a change point can be detected at least one day after its occurrence. Consequently, when $n = 1$, a detected change point aligns with the first day of the sliding window. In addition, Figure 4.9 illustrates the mean absolute difference between actual and predicted total number of infected companies per prediction window, when the two models are fitted to the past $n = 30$ crisis days of the a) 2008 and b) 2020 financial crises. Each prediction window has length $k = 30$ days, so the first prediction window starts at day 31 and ends at day 60 (with model fitted to the first $n = 30$

Algorithm 4 Model predictions for the future k days: CPD included

1. Set $t = 1$; set threshold value $\mathcal{T}$
2. Set $t_c = t$ (t_c is used to record and update the change points' times)
3. For $i \in \{t, ..., t + n\}$ calculate the anomaly score a_i, using the ChangeFinder algorithm. If a_i is a local maxima and $a_i > \mathcal{T}$ (if the anomaly score is over the threshold $\mathcal{T}$), then i is classified as a change point. Set $t_c = i$.
4. Fit the model to the period from day t_c (disregard all data points before the change point) to day $t + n$ (window t) and estimate $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$.
5. Run $N = 10{,}000$ SIR simulations on the network from day $t + n + 1$ to day $t + n + k$ (i.e. predict infections in the next k days) using the values of $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$ estimated in Step 4.

   (a) Predict the number of infected companies

   - Count the number of newly infected (recovered, respectively) companies in each simulation and take the mean number of the newly infected (recovered, respectively) companies as a prediction for the number of newly infected (recovered, respectively) companies in the period from day $t + n + 1$ to day $t + n + k$ (prediction window t).
   - Count the total number of infected companies at the end of each simulation and take the mean total number of infected companies as a prediction for the total number of infected companies on the final day with a prediction window t.

   (b) Predict the infected companies' continents/sectors

   - Collect a multiset (set allowing multiple instances of the same element) containing the continents/economic sectors of newly infected companies in each simulation and take it as a prediction for which continents/sectors are likely to be infected in prediction window t. For each simulation compare the predicted to the observed continents/sectors multiset using the Sørensen–Dice similarity coefficient (3.14). The mean over all N simulations is then used as an indicator of prediction quality for prediction window t.

   (c) Predict the infected companies

   - Collect a list of newly infected companies in each simulation and record the proportion of the N simulations in which each company was infected. Then, for every company, consider this proportion as the probability of infection during prediction window t of the crisis.

6. Update $t = t + 1$.
7. If $t > L - k - n$ stop. Else go back to Step 2.

crisis days and SIR simulations running from day 31 to day 60). It can be seen that while the model without change point detection performs poorly during and after significant events, denoted by vertical lines, incorporating change points into the model significantly enhances the results, specifically around these critical moments.

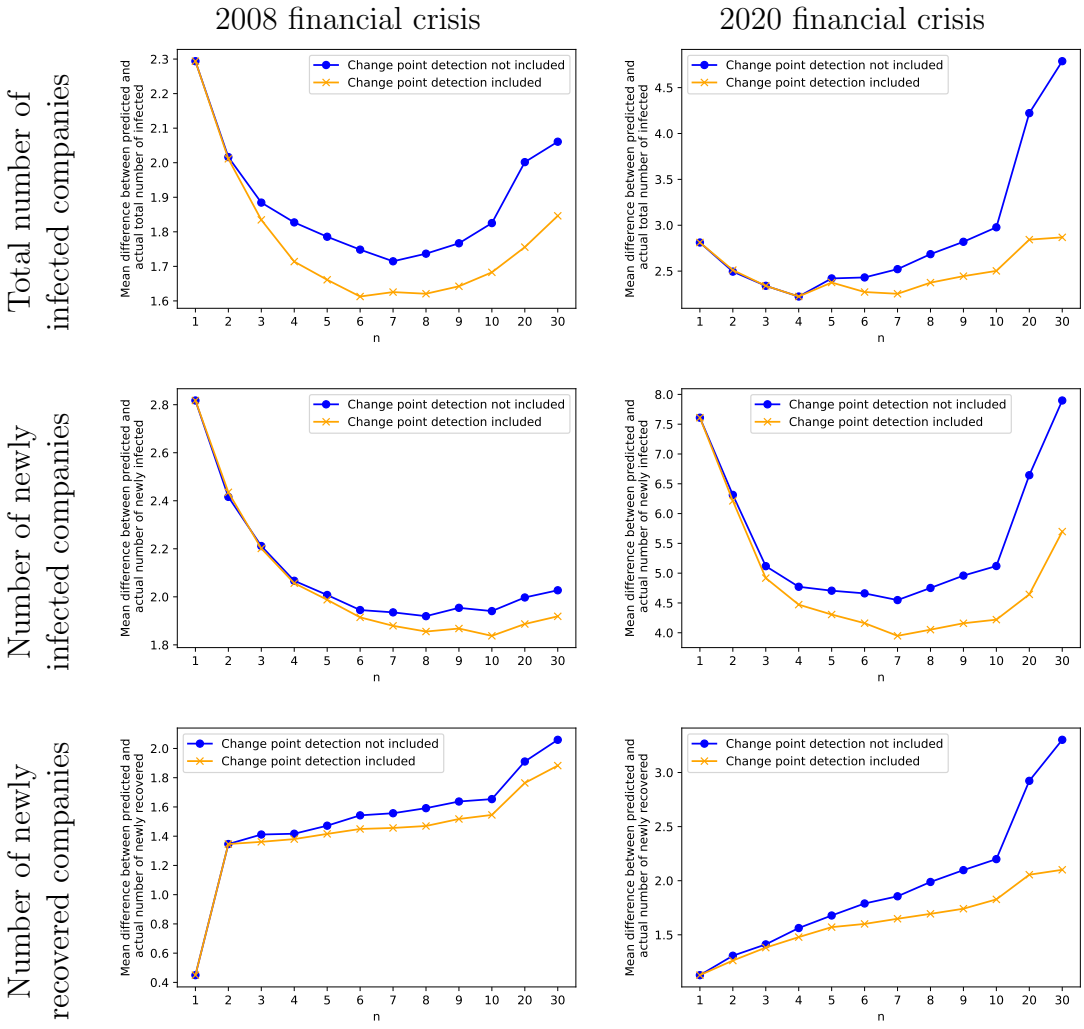

Figure 4.6: Comparison between the full model from Chapter 3 and the change point detection model of the absolute difference after $k = 10$ days between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days.

We next compare how the two models perform in predicting the infected companies' continents and sectors. Figures 4.10, 4.11 and 4.12 illustrate the

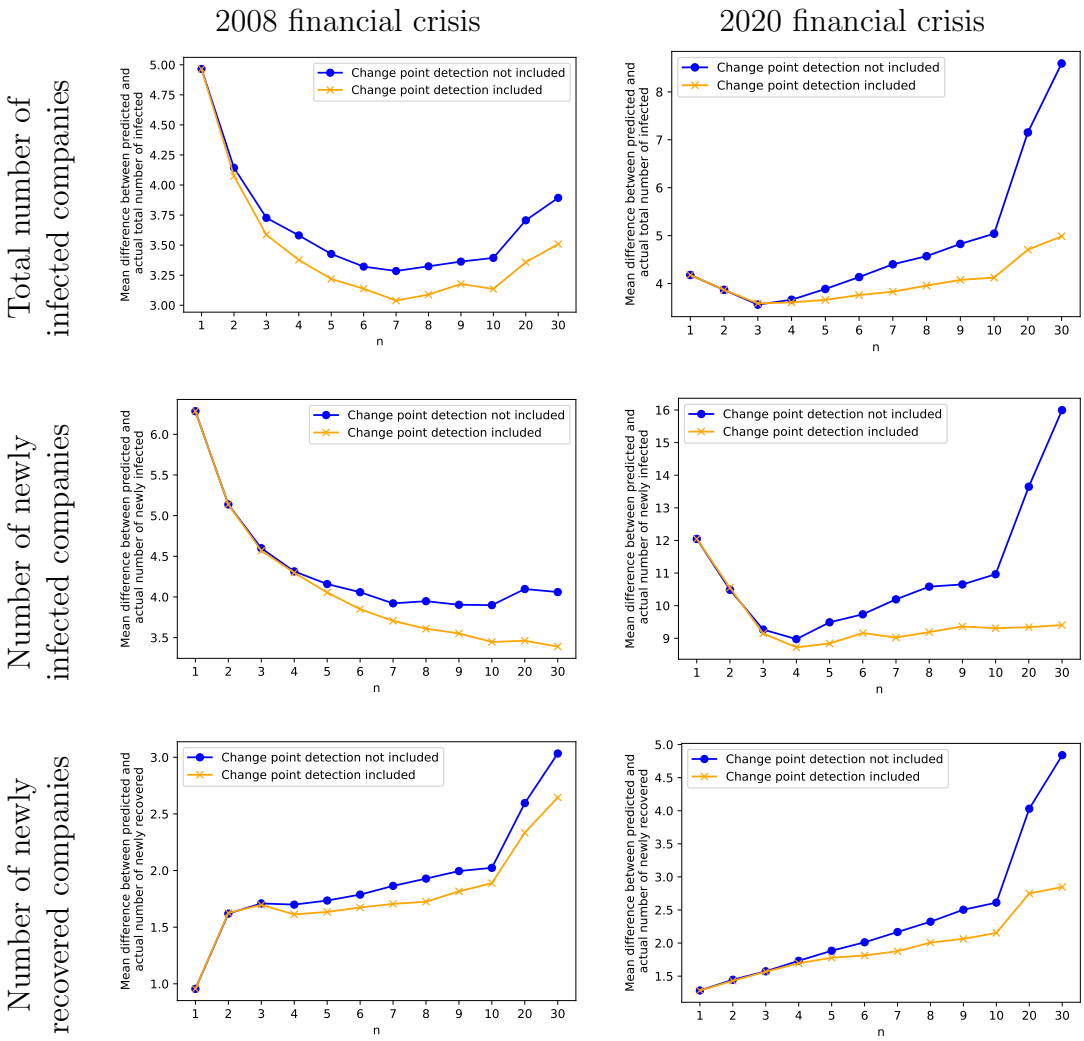

Figure 4.7: Comparison between the full model from Chapter 3 and the change point detection model of the absolute difference after $k = 20$ days between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days.

comparison between the mean Sørensen–Dice coefficient over all sliding windows between predicted and actual continents (top row) and sectors (bottom row) in the future $k = 10$, $k = 20$ and $k = 30$ days, respectively, for different values of n. It can be seen that for all values of k (10, 20, and 30 days into the future) and high values of n, the change point model performs consistently better than the model without change point detection incorporated, with higher improvement for higher values of n. This suggests that

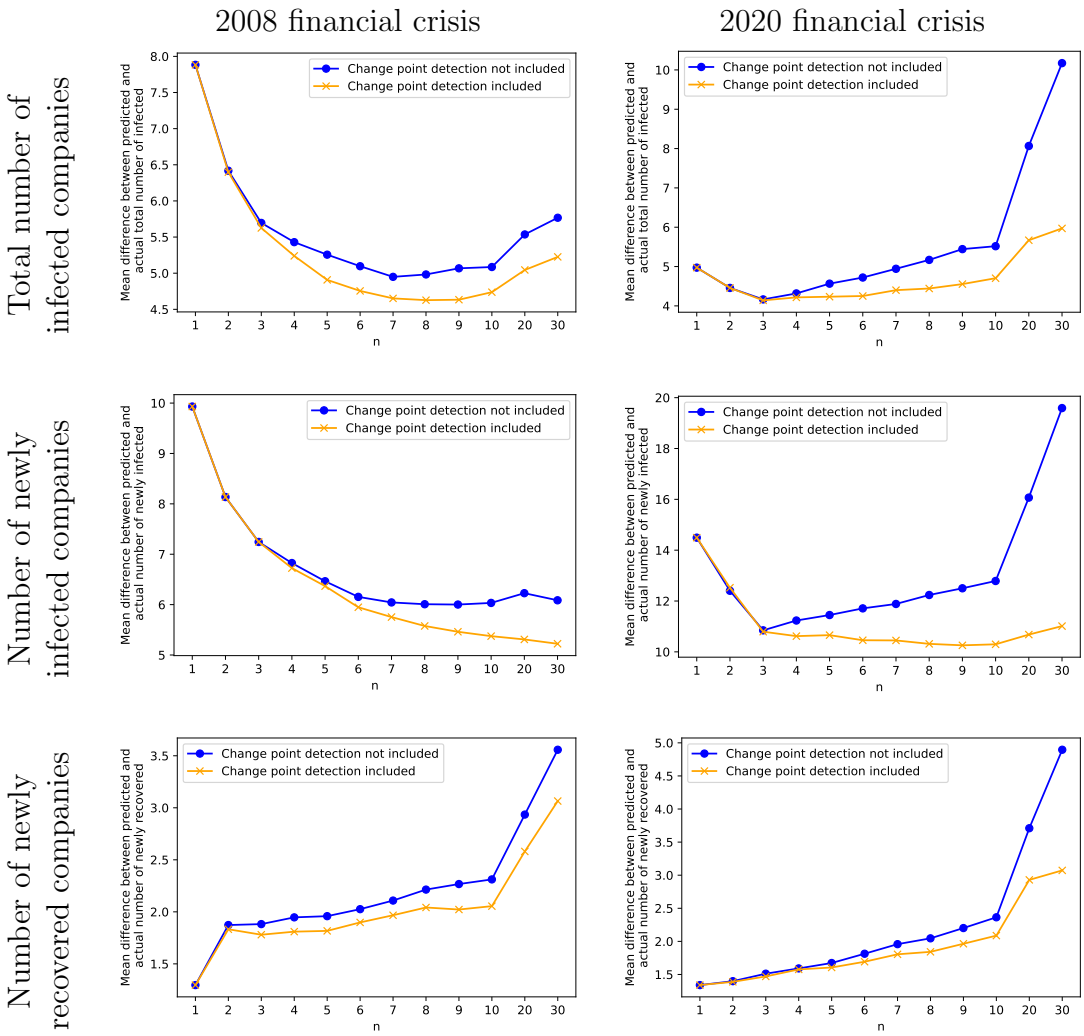

Figure 4.8: Comparison between the full model from Chapter 3 and the change point detection model of the absolute difference after $k = 30$ days between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous n days.

integrating change point detection improves the model's ability to accurately predict the geographical and sectoral spread of infection when the number of observations (n) is sufficiently large. The improvement is more pronounced for higher values of n, indicating that the change point model benefits from more data to detect significant shifts and adjust parameters accordingly. In addition, for small n there is little to no gain by ignoring part of the data as there is little amount of data anyway. However, it can also be seen that for

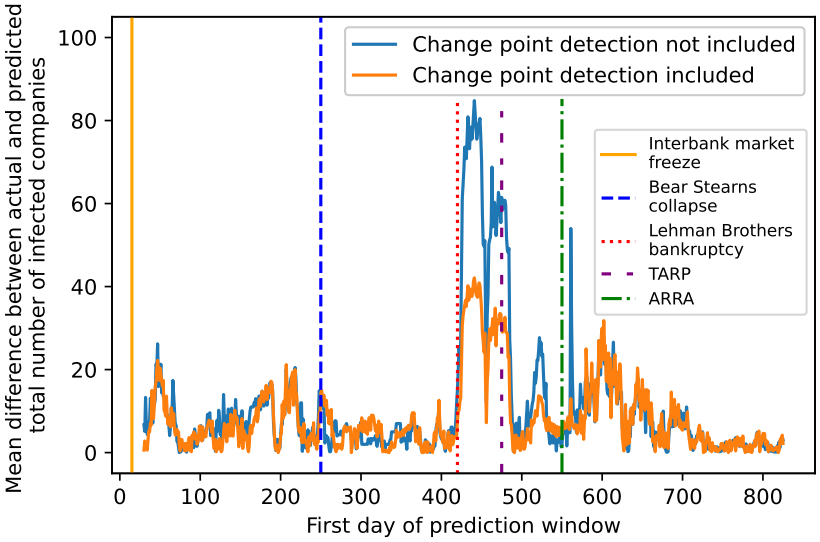

(a)

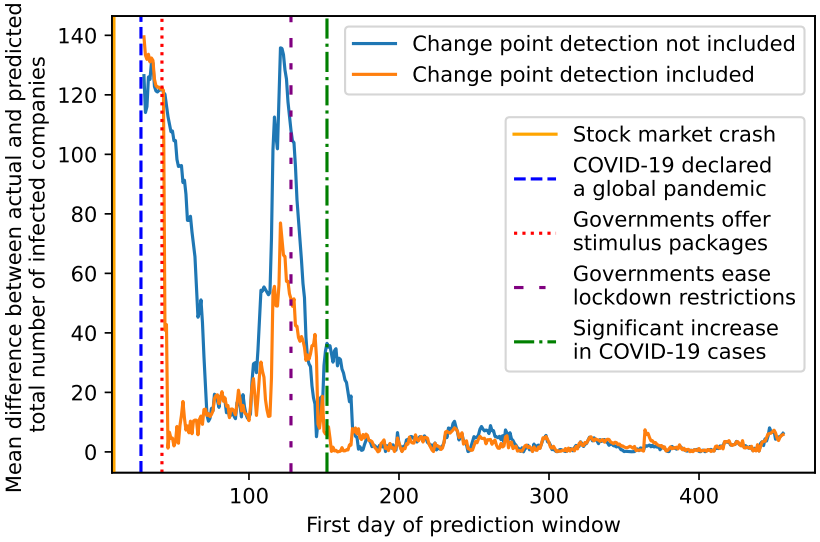

(b)

Figure 4.9: Comparison between the full model from Chapter 3 (in blue) and the change point detection model (in orange) of the mean absolute difference between predicted and actual total number of infected companies for the future $k = 30$ days per prediction window of the a) 2008 and b) 2020 financial crises. The models have been fitted to the past $n = 30$ days, hence the first day of the first prediction window is 31. The coloured vertical lines indicate the important events during each crisis.

smaller values of n, the change point model sometimes yields a lower mean Sørensen–Dice coefficient compared to the model without change point detection. The reason for this could be that when the model is fitted to a smaller window (i.e., a smaller number of data points n), detecting a change point within this window can result in the exclusion of a substantial portion of data. Consequently, a significant amount of information is lost, which can adversely affect the model's accuracy and predictive performance. Overall, while the change point model shows clear advantages for larger values of n, its performance can be less reliable with smaller n. This highlights the importance of considering the window size n when choosing to incorporate change point detection into the model.

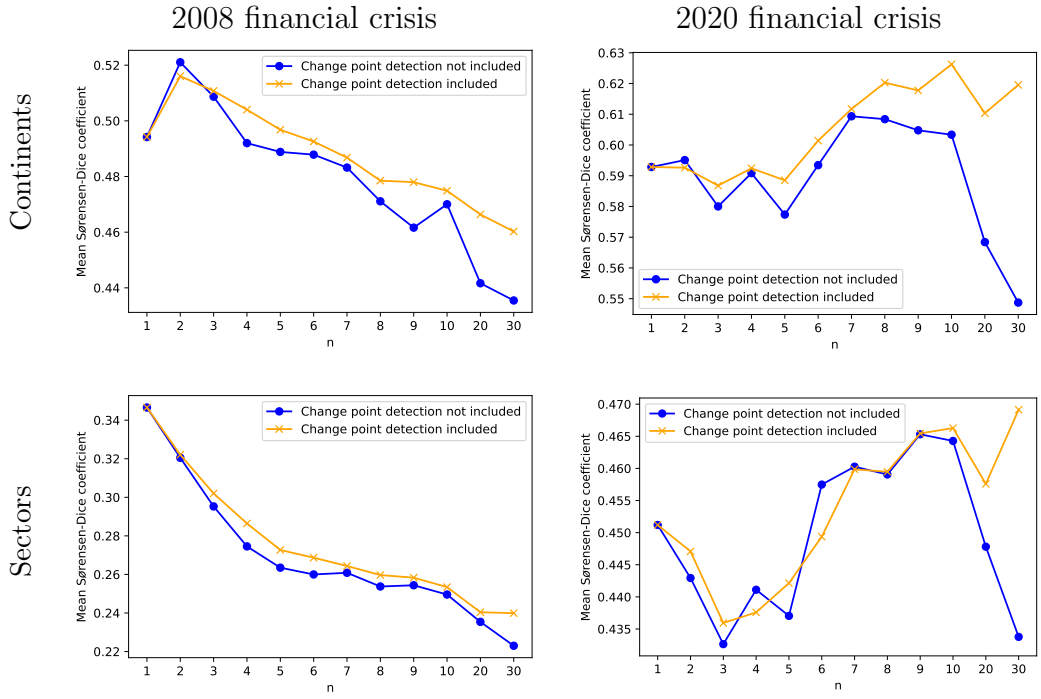

Figure 4.10: Comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 10$ days for the 2008 (left columns) and the 2020 (right column) financial crises, using the infections data from the previous n days when using the full model from Chapter 3 and the change point detection model.

We finally compare the Accuracy and the F_1-scores, introduced in Section

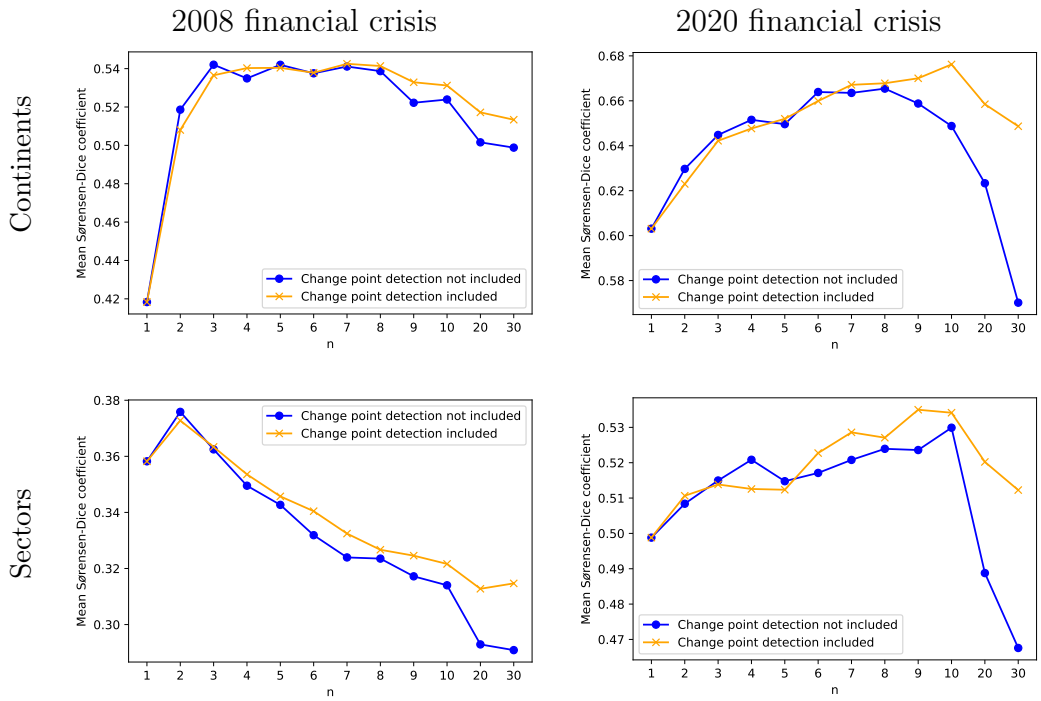

Figure 4.11: Comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 20$ days for the 2008 (left columns) and the 2020 (right column) financial crises, using the infections data from the previous n days when using the full model from Chapter 3 and the change point detection model.

3.4.2.4. Figures 4.13, 4.14 and 4.15 show that for all considered values of k the change point model has higher Accuracy for almost all values of n. This indicates that the change point model is generally better at correctly predicting the infected companies when looking at the overall proportion of correct predictions. However, when considering the F_1-scores for smaller values of n the model without change points consistently outperforms the change point model. The lower F_1-scores for the change point model suggest that while it may make accurate overall predictions (reflected in higher Accuracy), it might struggle more with correctly identifying the true positive cases. The explanation for this is that The F_1-score is, as explained earlier in Section 3.4.2.4, a harmonic mean of precision (the proportion of predicted infected companies that are actually infected) and recall (the proportion of actual infected companies that the model predicts as infected). For larger values of n the change

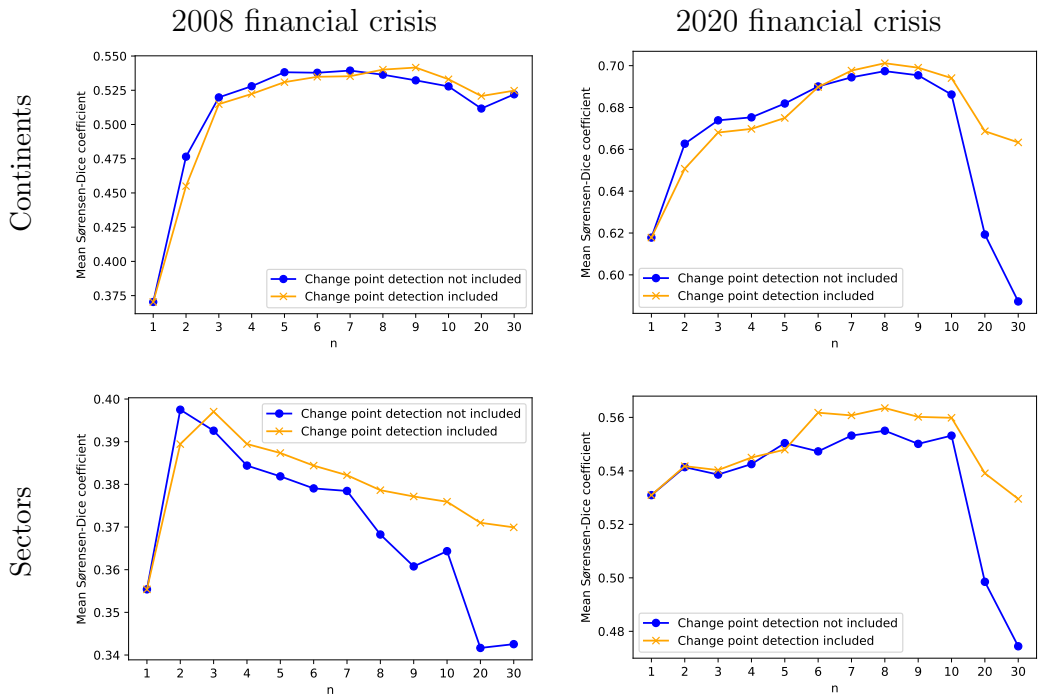

Figure 4.12: Comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the 2008 (left columns) and the 2020 (right column) financial crises, using the infections data from the previous n days when using the full model from Chapter 3 and the change point detection model.

point model generally yields higher F_1-score values. The results suggest that in certain instances information preceding a detected change point may hold significant relevance for future predictions.

4.4 Discussion and conclusion

In this chapter, we introduced and incorporated the ChangeFinder algorithm into our model for change point detection in real time. Through a comparative analysis, we demonstrated the enhanced predictive performance of the change point model over the original model from the previous chapter. The improvement across various metrics affirms the effectiveness of integrating change points for a more accurate and adaptive predictive model.

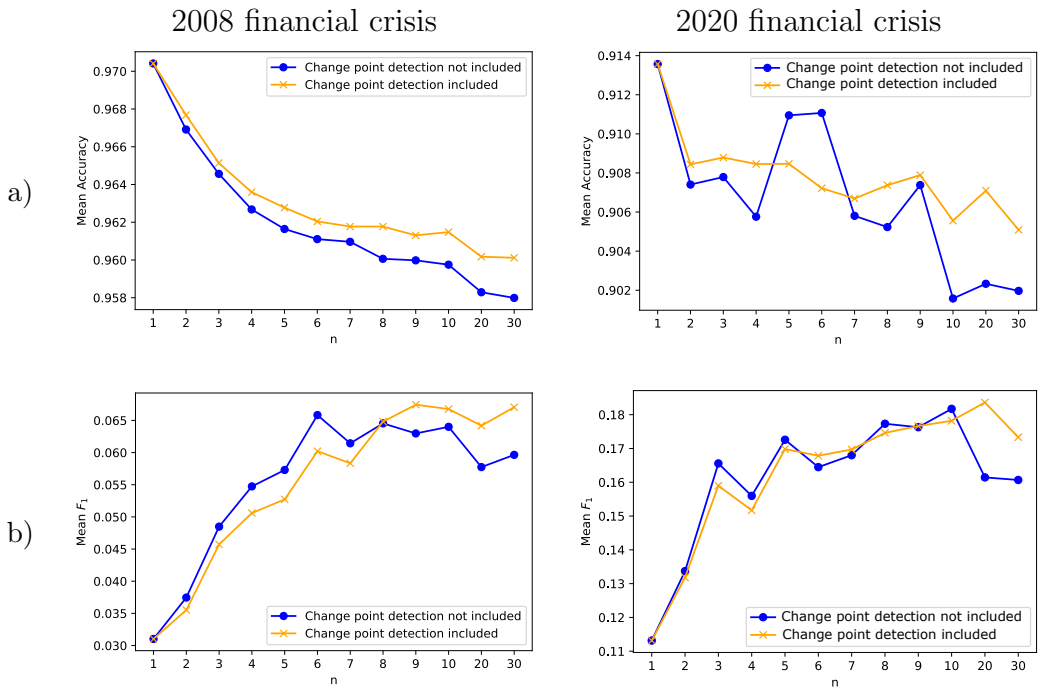

Figure 4.13: Comparison between the mean a) Accuracy and b) F_1-score in the future $k = 10$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right column), using the infections data from the previous n days to predict the set of individual infected companies when using change point detection model (in orange) vs the model without incorporating change points (in blue).

While the integration of change points proves advantageous, it is important to acknowledge the potential drawbacks. Firstly, the necessity to choose a threshold (step 1 in Algorithm 4) could impact the algorithm's sensitivity to change points. Another drawback of the model is that it does not consider data prior to a detected change point meaning that potentially valuable information present in older data may be overlooked. While the model's focus on recent data aids in capturing immediate trends, and change points in particular, the omission of historical insights may impact its ability to discern more persistent patterns. To address these limitations, in the next chapter, we introduce an observation weighting approach, which assigns higher weights to more recent observations while still retaining and incorporating older data to a lesser extent. This method aims to balance the need to quickly respond to new data with keeping valuable historical information, leading to higher

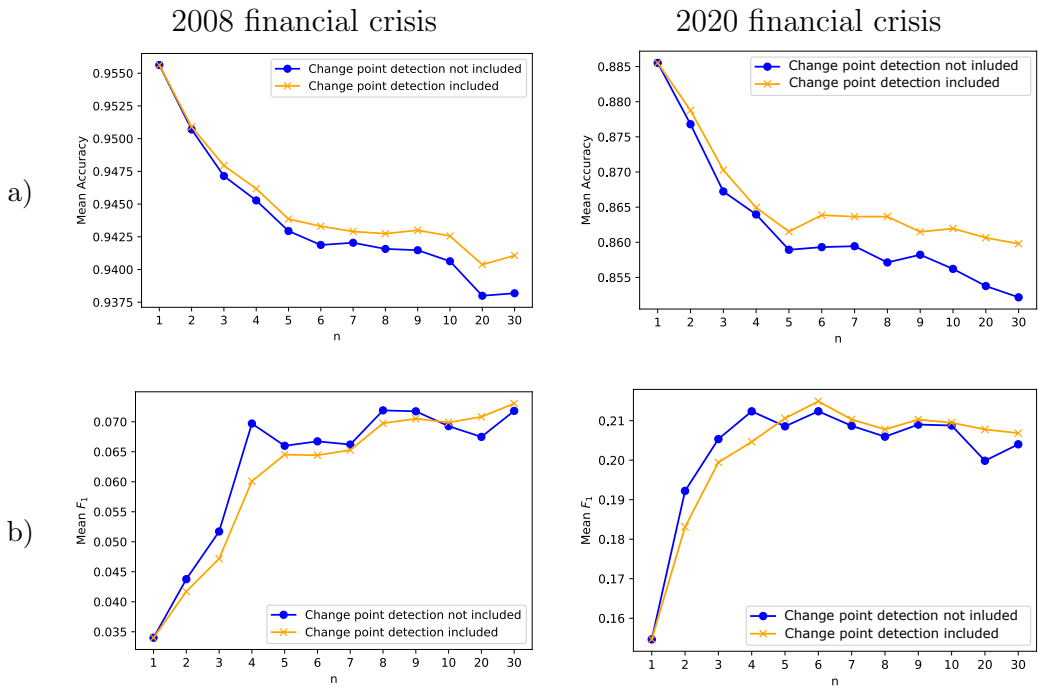

Figure 4.14: Comparison between the mean a) Accuracy and b) F_1-score in the future $k = 20$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right column), using the infections data from the previous n days to predict the set of individual infected companies when using change point detection model (in orange) vs the model without incorporating change points (in blue).

predictive accuracy.

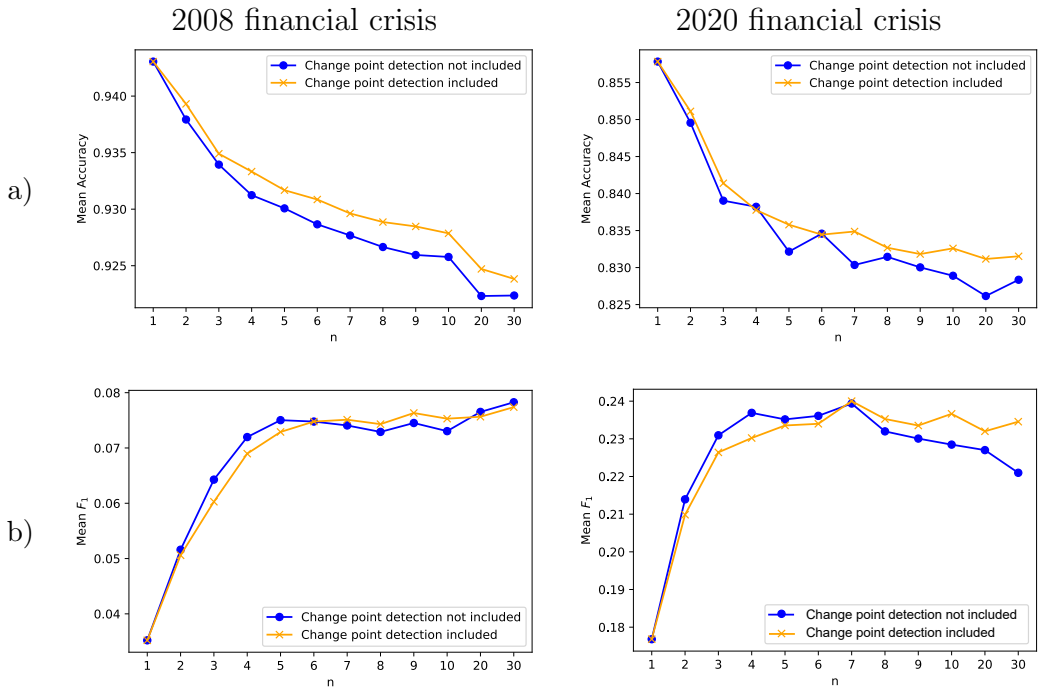

Figure 4.15: Comparison between the mean a) Accuracy and b) F_1-score in the future $k = 30$ days for the 2008 financial crisis (left column) and the 2020 financial crisis (right column), using the infections data from the previous n days to predict the set of individual infected companies when using change point detection model (in orange) vs the model without incorporating change points (in blue).

Chapter 5

Observation weighting in predictive modelling

5.1 Introduction

Many time series models from econometrics and empirical finance primarily employ maximum likelihood methods for estimation. The conventional maximum likelihood principle assumes equal weight for all sample observations. However, adjustments of the weighting scheme can improve forecast accuracy. For instance, in the modelling of financial time series, assigning greater weight to recent observations, as opposed to the default equal weights, has been shown to enhance risk prediction [206]. This concept arises naturally and is not new to the existing body of literature. The method of time-based weighted maximum likelihood (TiWML) has been used to forecast financial returns [150]. It is particularly useful in forecasting of financial time series as financial markets are dynamic in nature, with trends and patterns changing over time. Giving more weight to recent data allows the models to adapt to shifts in market conditions. In addition, the TiWML technique has found

various applications in risk management and financial modelling [174], particularly in scenarios of market volatility, as recent volatility shocks are well recognised to be good indicators of volatility in subsequent periods, or rapid change, as previous data might become obsolete [206].

In this chapter we aim to enhance the model's predictive accuracy by applying different weighting schemes to the original model introduced in Chapter 3. In particular, in Section 5.2 we employ two popular TiWML schemes, geometric and hyperbolic, in place of the MLE method described in Chapter 3. The resulting model is applied to the 2020 financial crisis. The results indicate that introducing both weighting schemes improves the model's accuracy when predicting the number of infected companies. In particular, the geometric weighting scheme consistently outperforms the hyperbolic weighting scheme, leading to substantially higher improvement in predictive accuracy, especially for higher values of the window length n.

Then, in Section 5.3 we combine the change point detection model (CPD model), introduced in Chapter 4, and the hyperbolic and geometric TiWML schemes. In particular, we first detect change points in the data using the procedure described in the previous chapter, discarding all the data preceding the change point. We then weight the data after (and including) the change point using either of the two weighting schemes. The results indicate that the CPD + TiWML model gives more accurate predictions for the number of infected companies, particularly when using the geometric weighting scheme.

5.2 TiWML

The TiWML approach is based on the assumption that recent observations may better reflect the current market conditions and trends. Hence, the

method gives higher weight to recent data and lower weight to older data. To implement the weighting scheme a vector of weights $(\tau_1, ..., \tau_n)$ is chosen, such that $\sum_{i=1}^{n} \tau_i = 1$, where n is the window size (the number of previous crisis days we use to fit the model), as introduced in Chapter 3. Then, each component of the log-likelihood function, associated with day t is multiplied by τ_t. In particular, the likelihood function of the observed new infections, associated with the period of the crisis from day d to day $d + n$, introduced in equation (3.10), becomes

$$\mathcal{L}_{d,d+n} = \prod_{t=d}^{d+n-1} (\mathcal{L}_{t,t+1})^{\tau_t}, \tag{5.1}$$

where $\mathcal{L}_{t,t+1}$ is defined in (3.9). Therefore, the log-likelihood function becomes

$$\ell_{d,d+n} = \log \mathcal{L}_{d,d+n} = \sum_{t=d}^{d+n-1} \tau_t \log \mathcal{L}_{t,t+1}. \tag{5.2}$$

Similarly, the recovery likelihood function, introduced in equation (3.12), becomes

$$\mathcal{L}^p_{d,d+n} = \prod_{t=d}^{d+n-1} (\mathcal{L}^p_{t,t+1})^{\tau_t}, \tag{5.3}$$

where $\mathcal{L}^p_{t,t+1}$ is defined in (3.11). Then, the recovery log-likelihood function becomes

$$\ell^p_{d,d+n} = \log \mathcal{L}^p_{d,d+n} = \sum_{t=d}^{d+n-1} \tau_t \log \mathcal{L}^p_{t,t+1}. \tag{5.4}$$

In this section we will study two of the most popular weighting schemes used in the literature [150], namely geometric weighting, for which $\tau_t \propto \rho^{n-t}$, and hyperbolic, given by $\tau_t \propto (n - t + 1)^{\rho-1}$. Here the parameter ρ indicates the strength of the more recent observations in comparison with older ones. In particular, a value of $\rho > 1$ ($\rho < 1$) indicates that more recent observations have greater (less) relative importance than past observations. A value of $\rho = 1$ indicates that all observations have the same weight. Figure 5.1

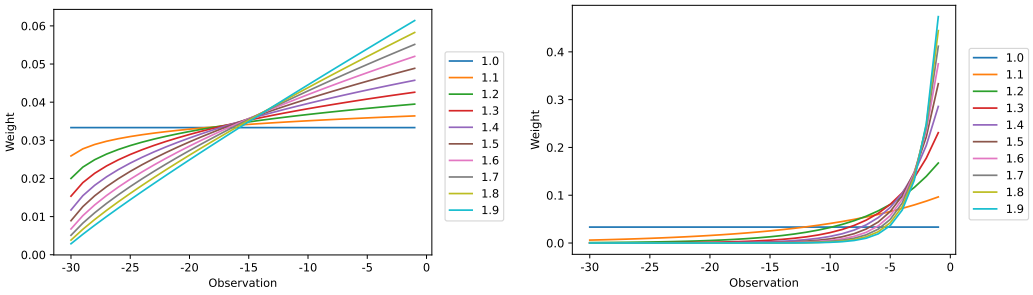

(a) (b)

Figure 5.1: a) Hyperbolic and b) geometric weighting schemes for window length $n = 30$ and different values of ρ. The day/observation furthest in the past (-30) has the lowest weight (importance), while the most recent day/observation (-1) is given the highest weight (importance) depending on the value of ρ.

illustrates the hyperbolic and geometric weighting schemes for window size $n = 30$ and different values of $\rho > 1$. It can be seen that hyperbolic weighting assigns weights that decrease gradually with distance in time, ensuring all observations retain some influence, whereas geometric weights drop off rapidly, potentially reducing the contribution of distant observations to near insignificance. In the following section we will employ both weighting methods and study how the value of ρ influences the model's predictive power. The step-by-step procedure is described in Algorithm 5.

5.2.1 TiWML: Evaluation and comparison

5.2.1.1 Prediction of the number of infected companies

In this section we compare how the hyperbolic and geometric weighting schemes influence the model's accuracy in predicting the number of infected companies in the future k crisis days, where the model is fitted to the previous n crisis days, similarly to the fitting procedures shown in Chapters 3 and 4. We present results for a prediction horizon of $k = 10$, $k = 20$ and $k = 30$ days in Figures 5.2, 5.3 and 5.4 display the absolute difference between the predicted

Algorithm 5 Model predictions for the future k days: TiWML

1. Set $t = 1$; set weighting parameter ρ; choose a weighting scheme
2. For $i \in \{t, ..., t+n-1\}$ calculate the weighting coefficients τ_i, based on the chosen weighting scheme and the value of ρ, such that $\sum_{i=t}^{t+n-1} \tau_i = 1$.
3. Fit the model to the period from day t to day $t+n$ (window t) and maximizing the weighted likelihood functions $\mathcal{L}_{t,t+n}$ (see (5.1)) and $\mathcal{L}^p_{t,t+n}$ (see (5.3)) estimate $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$, respectively.
4. Run $N = 10{,}000$ SIR simulations on the network from day $t+n+1$ to day $t+n+k$ (i.e., predict infections in the next k days) using the values of $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$ estimated in Step 3.

   (a) Predict the number of infected companies

   - Count the number of newly infected (recovered, respectively) companies in each simulation and take the mean number of the newly infected (recovered, respectively) companies as a prediction for the number of newly infected (recovered, respectively) companies in the period from day $t+n+1$ to day $t+n+k$ (prediction window t).
   - Count the total number of infected companies at the end of each simulation and take the mean total number of infected companies as a prediction for the total number of infected companies on the final day with a prediction window t.

   (b) Predict the infected companies' continents/sectors

   - Collect a multiset (set allowing multiple instances of the same element) containing the continents/economic sectors of newly infected companies in each simulation and take it as a prediction for which continents/sectors are likely to be infected in prediction window t. For each simulation compare the predicted to the observed continents/sectors multiset using the Sørensen–Dice similarity coefficient (3.14). The mean over all N simulations is then used as an indicator of prediction quality for prediction window t.

   (c) Predict the infected companies

   - Collect a list of newly infected companies in each simulation and record the proportion of the N simulations in which each company was infected. Then, for every company, consider this proportion as the probability of infection during prediction window t of the crisis.

5. Update $t = t + 1$.
6. If $t > L - k - n$, where L is the number of crisis days, stop. Else go back to Step 2.

and actual number of total infected (row 1), newly infected (row 2), and newly recovered (row 3) companies for different values of $\rho \in [1, 1.9]$, using the hyperbolic (left column) and geometric (right column) weighting schemes, for a prediction horizon of $k = 10$, $k = 20$ and $k = 30$ days, respectively. It can be seen that introducing both weighting schemes leads to more accurate predictions, with higher values of ρ consistently leading to greater improvement in predictions, compared to the non-weighted scenario ($\rho = 1$) for all values of k. In particular, the highest improvement for all three metrics of interest is achieved when $\rho = 1.9$. Specifically, for all values of k, when $\rho = 1.9$, the hyperbolic weighting scheme yields around 20% improvement, compared to the non-weighted scenario, in predicting the number of newly infected companies for large values of n. In contrast, the geometric weighting scheme improves the predictions by approximately 80%. In addition, the hyperbolic weighting scheme achieves its highest improvement in predicting the number of newly recovered companies when $k = 30$ and $\rho = 1.9$, leading to 30% improvement. However, the geometric weighting scheme outperforms it notably, achieving more than 200% improvement at the higher values of ρ and n. Similarly, for all values of k, with large n and ρ, the geometric weighting consistently achieves more than 110% improvement in predicting the total number of infected companies. The hyperbolic weighting scheme on the other hand does not exceed 20% improvement. In addition, in comparison with the hyperbolic weighting, the geometric weighting scheme results in substantially better predictions, even for small values of ρ. Moreover, one can notice that for higher values of ρ, the prediction accuracy of the model when using the geometric weighting scheme is nearly independent of n. Figures B.1, B.2 and B.3 in Appendix B present a visual representation of the percentage improvements mentioned above.

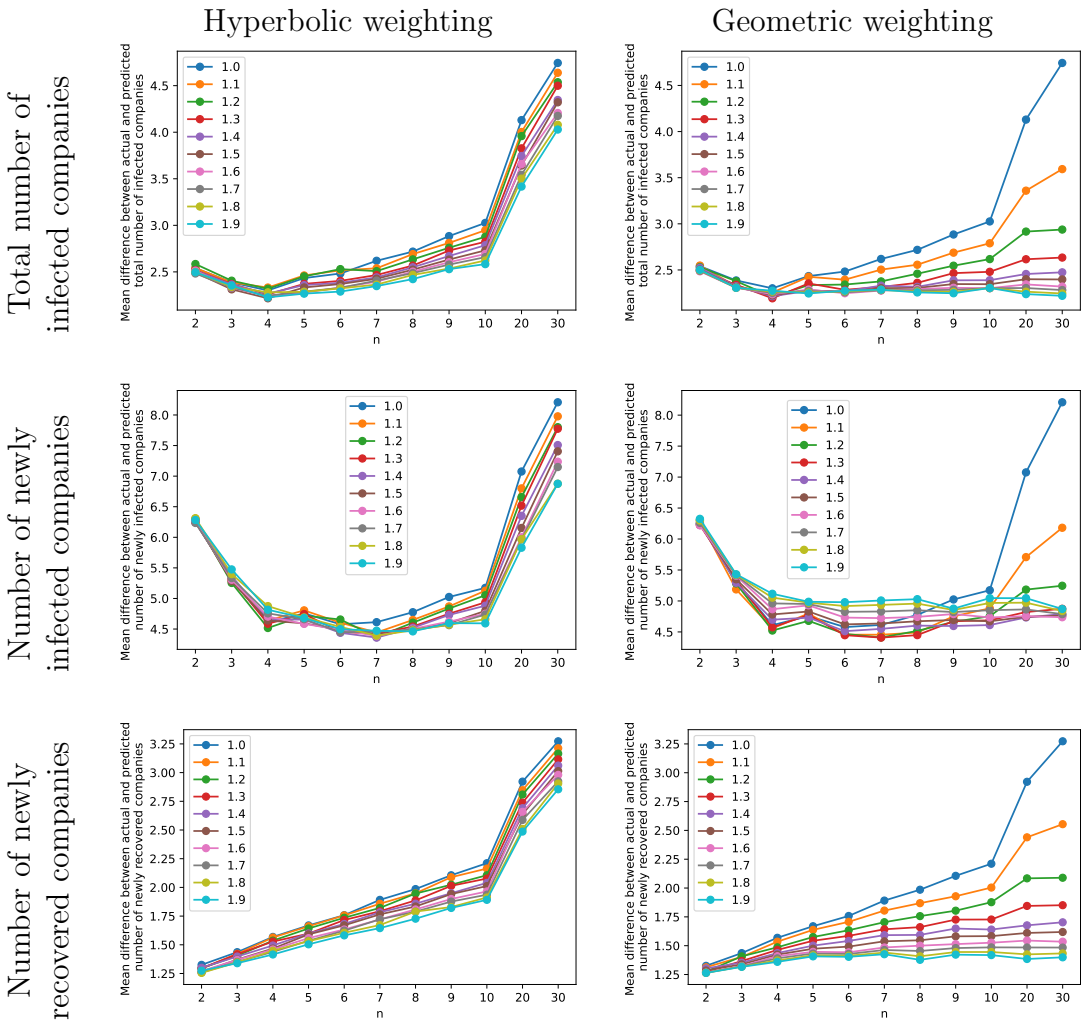

Figure 5.2: The plots illustrate the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2020 financial crisis, using the hyperbolic (left column) and geometric (right column) weighting schemes on the infections data from the previous n days and at a prediction horizon of $k = 10$ days.

5.2.1.2 Geographic- and sector-specific prediction

In this section, we study how the two weighting schemes perform in predicting the infected companies' continents and sectors. Figures 5.5, 5.6 and 5.7 illustrate the results for $k = 10$, $k = 20$ and $k = 30$, respectively. It can be seen that while in Section 5.2.1.1 introducing both weighting schemes was beneficial for predicting the number of infected companies, in the prediction of infected companies' continents and sectors for most integer values of $n \in [1, 10]$ it

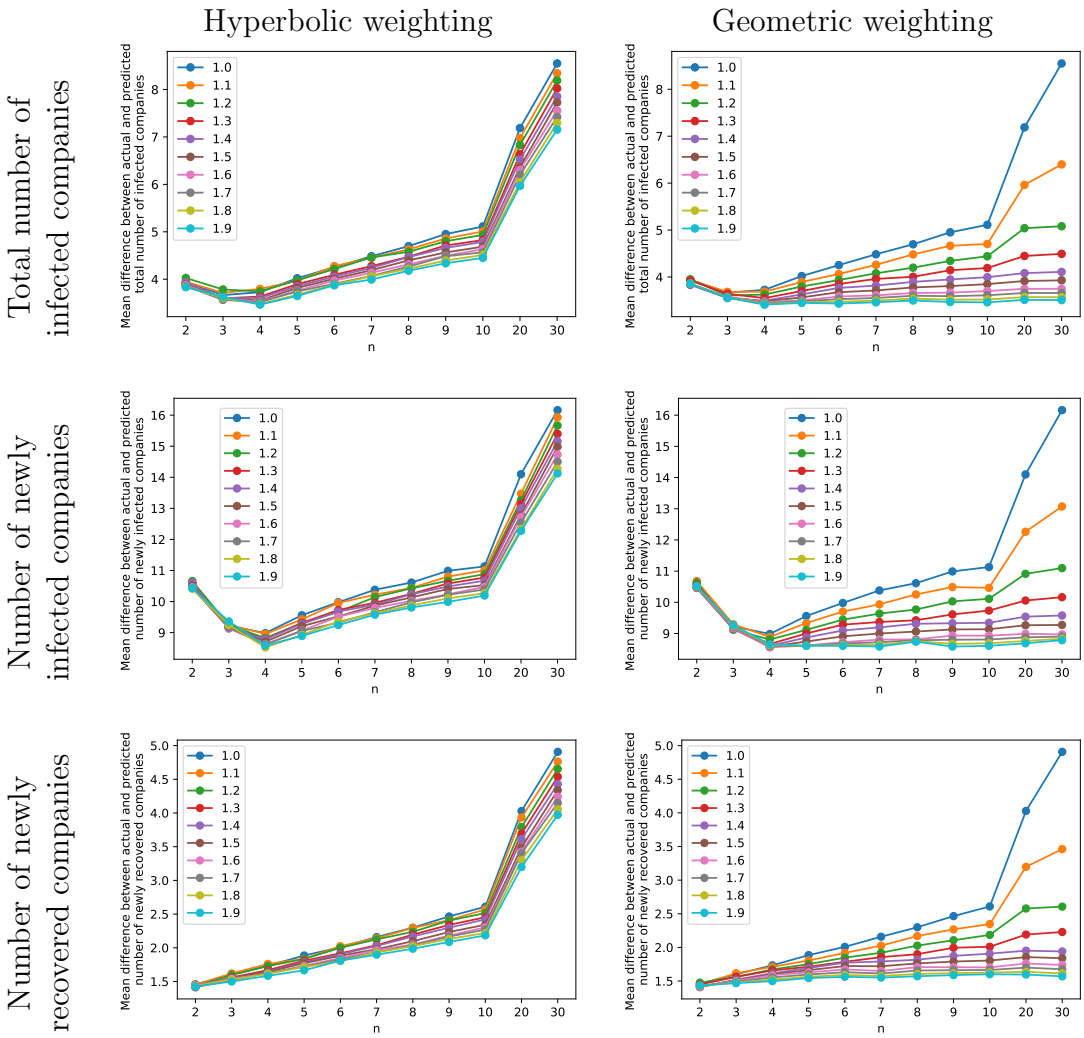

Figure 5.3: The plots illustrate the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2020 financial crisis, using the hyperbolic (left column) and geometric (right column) weighting schemes on the infections data from the previous n days and at a prediction horizon of $k = 20$ days.

consistently leads to a lower mean Sørensen–Dice coefficient, compared to the original model from Chapter 3. In more detail, the greatest improvement over the non-weighted model from Chapter 3 in predicting the infected companies' continents and sectors is around 4% for the hyperbolic and roughly 10% for the geometric weighting scheme. For a visual representation of these percentage improvements, refer to Figures B.4 , B.5, and B.6 in Appendix B. In addition, highest improvements are achieved when $n = 30$ for both weighting schemes,

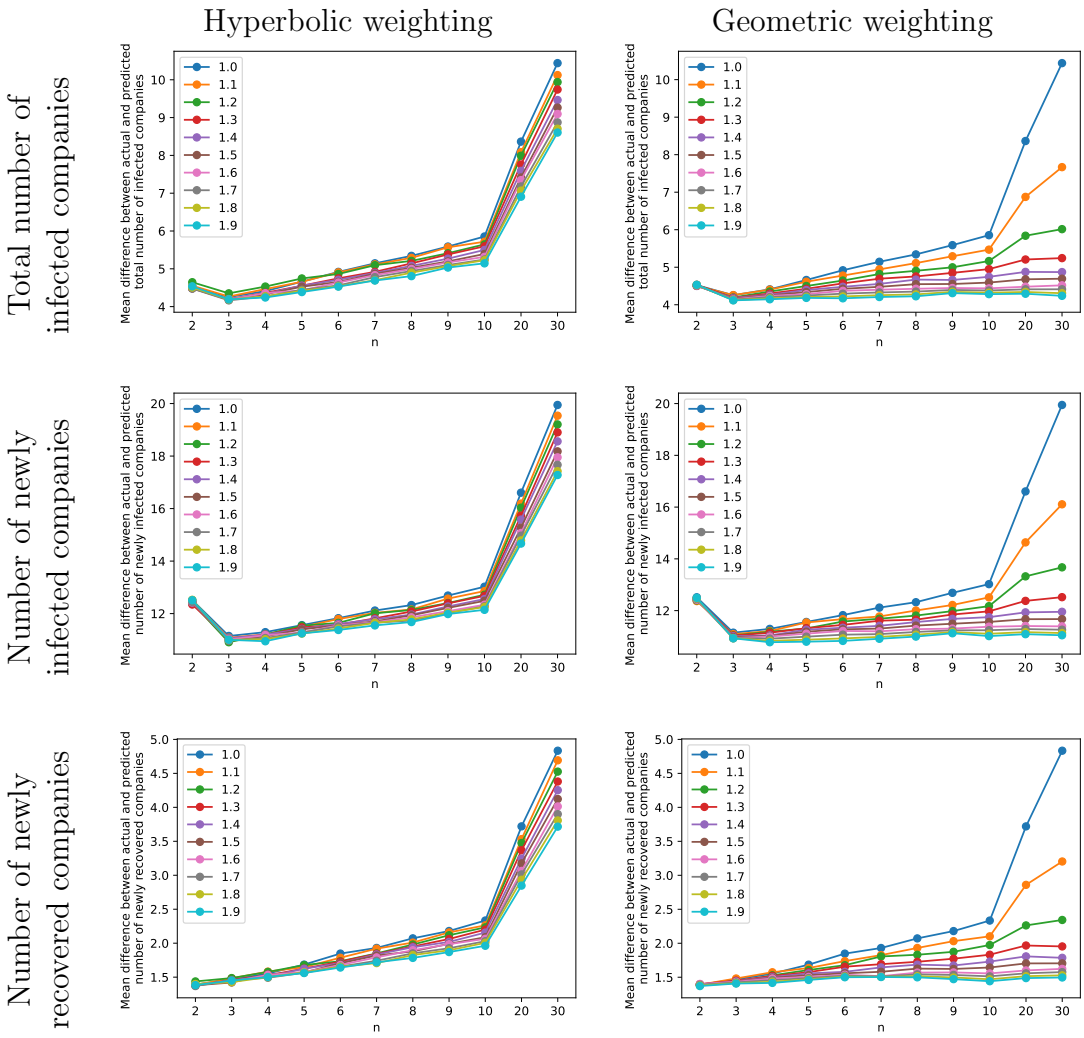

Figure 5.4: The plots illustrate the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2020 financial crisis, using the hyperbolic (left column) and geometric (right column) weighting schemes on the infections data from the previous n days and at a prediction horizon of $k = 30$ days.

with more accurate results achieved using the geometric weighting scheme. Overall, while in predicting the number of infected companies we could see a clear pattern, with higher values of ρ leading to greater improvement, here there is not such a consistent trend.

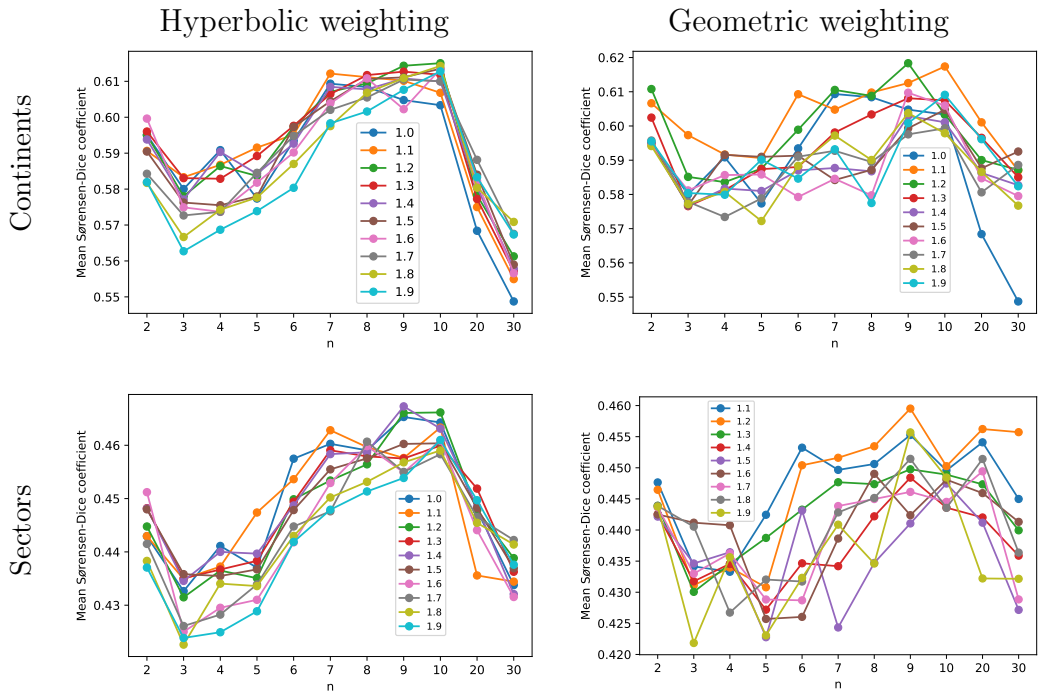

Figure 5.5: The plots show a comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 10$ days for the the 2020 financial crisis, using the infections data from the previous n days when using hyperbolic (left column) and geometric (right column) weighting schemes.

5.2.1.3 Predicting the specific infected companies

In this section we compute the Accuracy and the F_1-scores, as introduced in Section 3.4.2.4, for the hyperbolic and geometric weighting schemes across different values of ρ. Figures 5.8, 5.9 and 5.10 display the mean Accuracy (top row) and the mean F_1-score (bottom row), averaged over all sliding windows, for a prediction horizon of $k = 10$, $k = 20$ and $k = 30$ days, respectively. It can be seen that for all values of k and large values of n, when considering the mean Accuracy of the predictions, applying either weighting scheme leads to higher Accuracy, in comparison with the original model from Chapter 3 (equivalent to $\rho = 1$). In particular, higher values of ρ yield higher Accuracy. Specifically, the highest relative improvement in Accuracy achieved by the hyperbolic weighting scheme, in comparison with the non-weighted scenario,

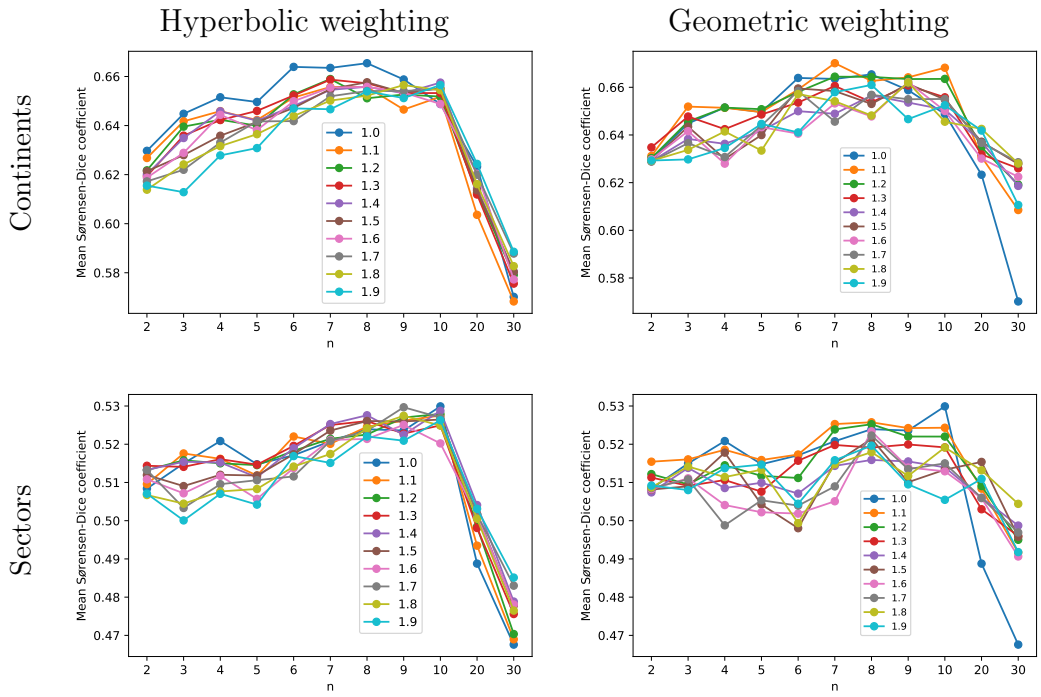

Figure 5.6: The plots show a comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 20$ days for the the 2020 financial crisis, using the infections data from the previous n days when using hyperbolic (left column) and geometric (right column) weighting schemes.

is the modest 0.4% when $k = 20$ and $n = 30$. In contrast, the geometric weighting scheme results in a higher relative improvement of 1.2% under the same conditions.

However, when comparing the mean F_1-scores, there is no consistent trend. Specifically, when $k = 10$, higher values of ρ lead to worse predictions for most values of n. When $k = 20$ and $k = 30$ introducing either weighting scheme generally leads to poorer predictions for $n \leq 7$, obviously so with the geometric weighting scheme. On the other hand, for $n \geq 8$ using a weighting with higher ρ results in better predictions, with higher ρ typically producing higher F_1-scores. In particular, for $k = 30$ and $n \geq 8$, the relative improvement in the F_1-score increases progressively, starting at 2% for $n = 8$ and rising to 6% for $n = 30$. A similar trend can be observed for $k = 20$ with

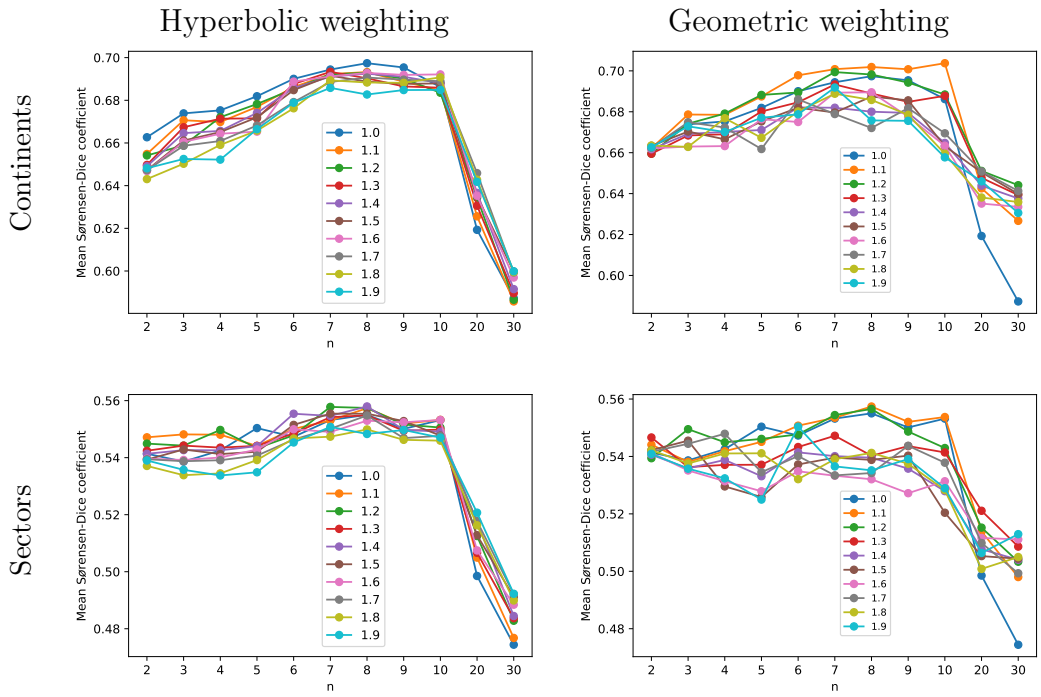

Figure 5.7: The plots show a comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the the 2020 financial crisis, using the infections data from the previous n days when using hyperbolic (left column) and geometric (right column) weighting schemes.

improvement starting from 1% for $n = 8$ and reaching more than 4% when $n = 30$. These observations suggest that the benefit of applying weighting schemes is more apparent when there is a sufficient amount of historical data (larger n). Overall, the results suggest that the choice of ρ and n should be carefully chosen based on the specific prediction horizon k and the metric of interest. For completion, Figures B.7, B.8 and B.9 in Appendix B present the percentage improvement in Accuracy and F_1-scores when using the TiWML model with hyperbolic and geometric weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (the original model from Chapter 3).

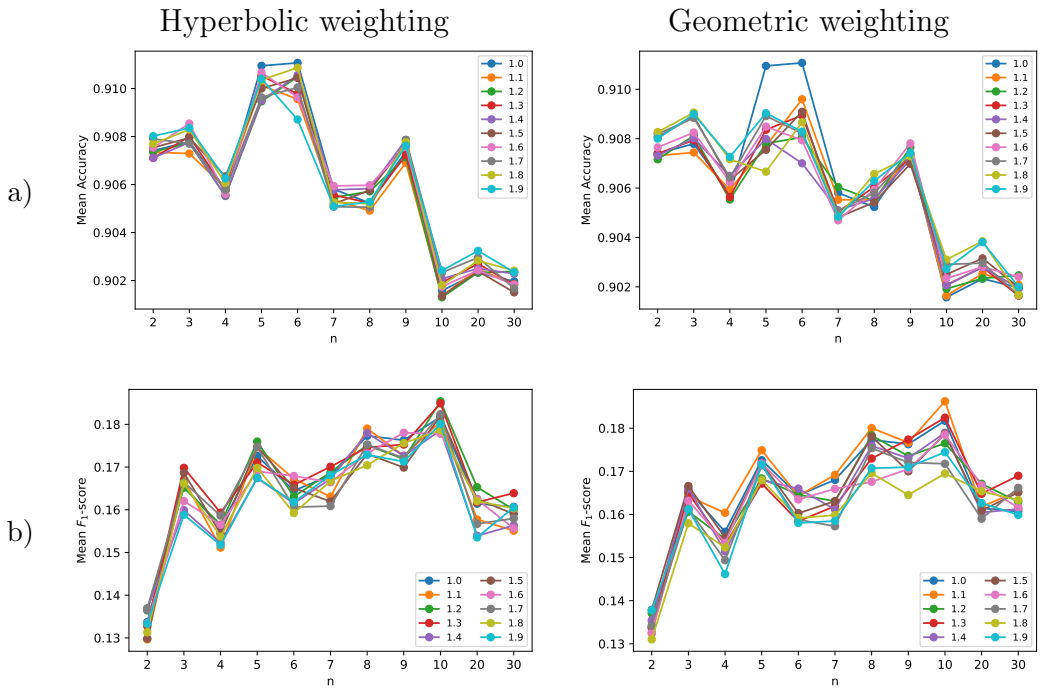

Figure 5.8: Comparison between the mean a) Accuracy and b) F_1-score in predicting the specific infected companies in the future $k = 10$ days for the the 2020 financial crisis, using the infections data from the previous n days when using hyperbolic (left column) and geometric (right column) weighting schemes.

5.3 Combining CPD and TiWML

In this section we combine the CPD method from Chapter 4, with the TiWML approach introduced in Section 5.2 as follows. Whenever a change point is detected during the past n days of a crisis, we discard the data before the change point has happened. However, while in Chapter 4 all the observations after a detected change point were weighted equally, in this section we instead apply the hyperbolic and geometric weighting schemes to the data after the change point and fit the model accordingly. If no change point is detected, we weight all past n observations. The step-by-step procedure is described in Algorithm 6.

Algorithm 6 Model predictions for the future k days: CPD + TiWML

1. Set $t = 1$; set threshold value $\mathcal{T}$; set weighting parameter ρ

2. Set $t_c = t$ (t_c is used to record and update the change-points times)

3. For $i \in \{t, ..., t + n\}$ calculate the 'anomaly' score a_i, using the ChangeFinder algorithm. If a_i is a local maxima and $a_i > \mathcal{T}$ (if the 'anomaly' score is over the threshold $\mathcal{T}$), then i is classified as a change-point. Set $t_c = i$.

4. Discard all data points before the change-point t_c and estimate the weighting coefficients $\tau_{t_c}, ..., \tau_{t+n-1}$, based on the chosen weighting scheme and the value of ρ, such that $\sum_{i=t_c}^{t+n-1} \tau_i = 1$.

5. Fit the model to the period from day t_c to day $t + n$ (window t) and maximizing the weighted likelihood functions $\mathcal{L}_{t_c,t+n}$ (see (5.1)) and $\mathcal{L}^p_{t_c,t+n}$ (see (5.3)) estimate $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$, respectively.

6. Run $N = 10{,}000$ SIR simulations on the network from day $t + n + 1$ to day $t + n + k$ (i.e., predict infections in the next k days) using the values of $\hat{\beta}_i$ $(1 \leq i \leq 4)$ and $\hat{p}$ estimated in Step 5.

(a) Predict the number of infected companies

- Count the number of newly infected (recovered, respectively) companies in each simulation and take the mean number of the newly infected (recovered, respectively) companies as a prediction for the number of newly infected (recovered, respectively) companies in the period from day $t + n + 1$ to day $t + n + k$ (prediction window t).
- Count the total number of infected companies at the end of each simulation and take the mean total number of infected companies as a prediction for the total number of infected companies on the final day with a prediction window t.

(b) Predict the infected companies' continents/sectors

- Collect a multiset (set allowing multiple instances of the same element) containing the continents/economic sectors of newly infected companies in each simulation and take it as a prediction for which continents/sectors are likely to be infected in prediction window t. For each simulation compare the predicted to the observed continents/sectors multiset using the Sørensen–Dice similarity coefficient (3.14). The mean over all N simulations is then used as an indicator of prediction quality for prediction window t.

(c) Predict the infected companies

- Collect a list of newly infected companies in each simulation and record the proportion of the N simulations in which each company was infected. Then, for every company, consider this proportion as the probability of infection during prediction window t of the crisis.

7. Update $t = t + 1$.

8. If $t > L - k - n$, where L is the number of crisis days, stop. Else go back to Step 2.

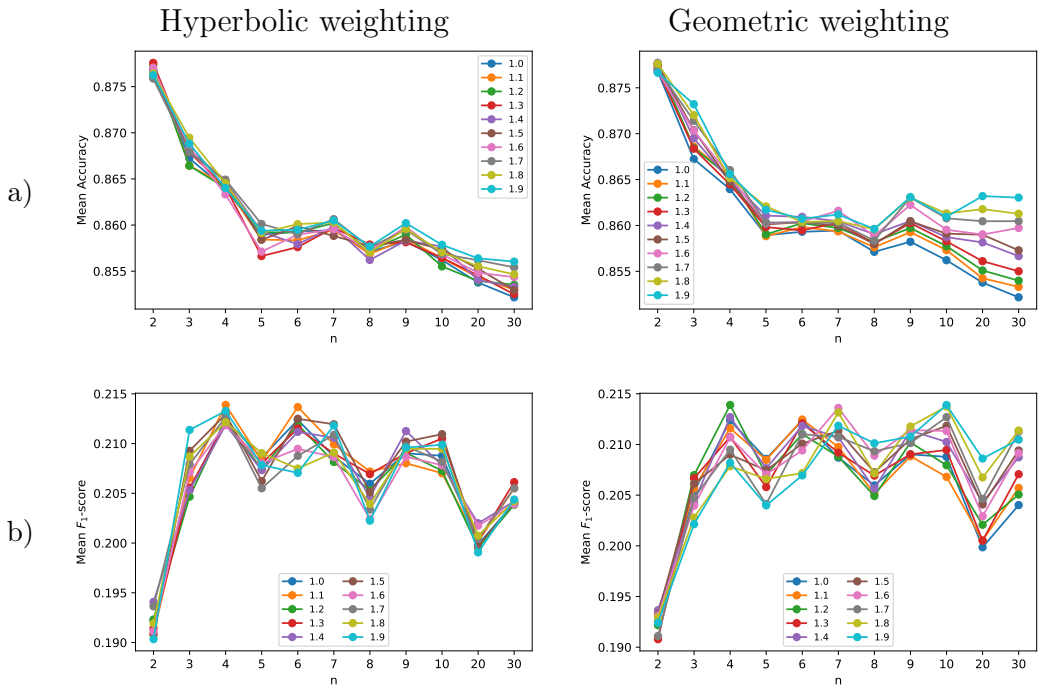

Figure 5.9: Comparison between the mean a) Accuracy and b) F_1-score in predicting the specific infected companies in the future $k = 20$ days for the the 2020 financial crisis, using the infections data from the previous n days when using hyperbolic (left column) and geometric (right column) weighting schemes.

5.3.1 Predicting the number of infected companies

We present the results in Figures 5.11, 5.12 and 5.13 for predicting the total number of infected companies (top row), the number of newly infected companies (mid row) and the number of newly recovered companies (bottom row) in the future 10, 20 and 30 days, respectively. We apply the CPD+TiWML method with the hyperbolic (left column) and geometric (right column) weighting schemes. It can be seen that for all values of k and under both weighting schemes, increasing the parameter ρ results in more accurate predictions across large window sizes $n \in \{10, 20, 30\}$, with most accurate predictions for larger n. Notably, for all combinations of n and k, the CPD+TiWML model with geometric weighting and parameter $\rho = 1.9$ yields the most accurate results in predicting the number of newly recovered com-

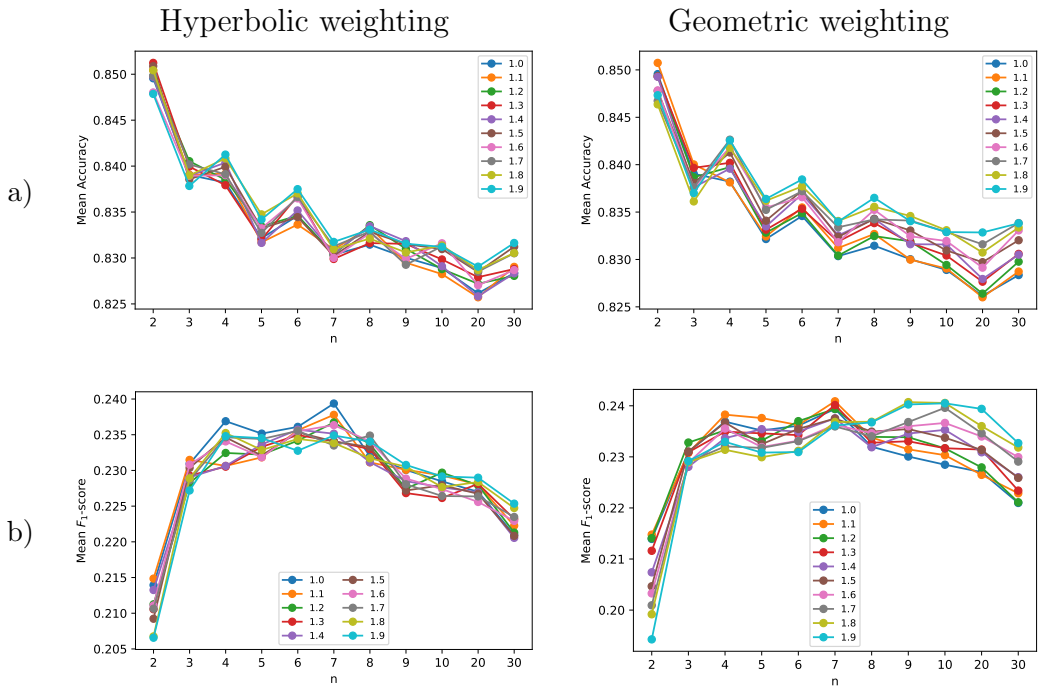

Figure 5.10: Comparison between the mean a) Accuracy and b) F_1-score in predicting the specific infected companies in the future $k = 30$ days for the the 2020 financial crisis, using the infections data from the previous n days when using hyperbolic (left column) and geometric (right column) weighting schemes.

panies. Specifically, the greatest improvement, compared to the non-weighted change point model, introduced in Chapter 4, is observed for $k = 30$ and $n = 30$, yielding more than 100% improvement. In addition, the greatest improvement in forecasting the total number of infected companies is achieved when $k = 30$ and $n = 30$, using geometric weighting with $\rho = 1.9$, leading to more than 50% improvement of the CPD+TiWML model, in comparison with the change point model from Chapter 4. In contrast, when predicting the number of newly infected companies for values of $n \in [2, 10]$, larger values of ρ lead to poorer predictions under both weighting schemes, compared to the non-weighted CPD model. Figures illustrating the relative percentage improvement achieved by incorporating the hyperbolic and geometric weighting schemes, compared to using only the change point model (introduced in Chapter 4) can be found in Appendix B.

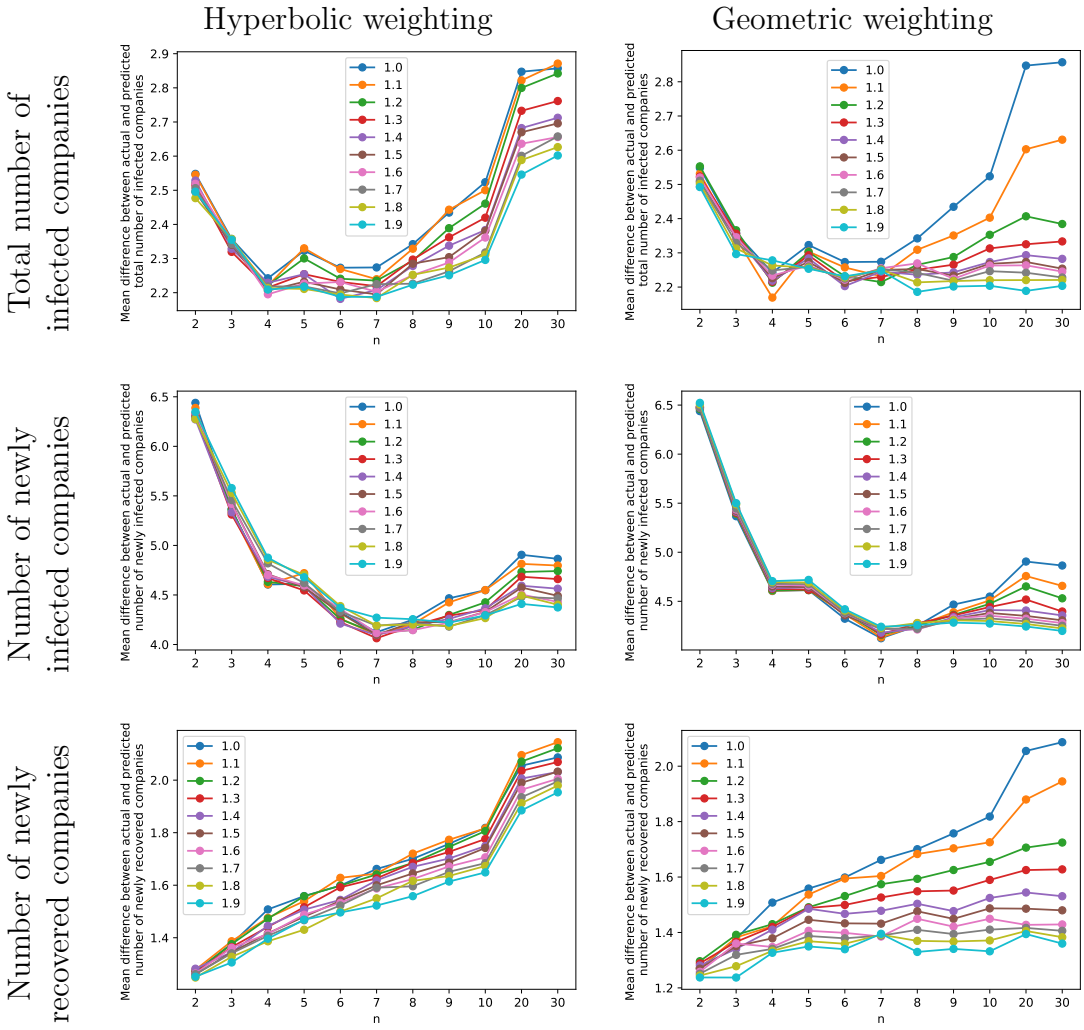

Figure 5.11: The plots show the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2020 financial crisis. The predictions are generated using the CPD+TiWML combination (see Algorithm 6), based on infection data from the previous n days and with a prediction horizon of $k = 10$ days. The left column shows the differences for the hyperbolic weighting scheme, while the right column shows the differences for the geometric weighting scheme, for various values of $\rho > 1$.

5.3.1.1 Sector- and continent-specific predictions

In this section we present a comparison of the model's accuracy in predicting the continents and the sectors of the infected companies in the future $k = 10$, $k = 20$ and $k = 30$ days. The results are shown in Figures 5.14, 5.15 and 5.16, respectively. It can be seen that while in the preceding section introducing

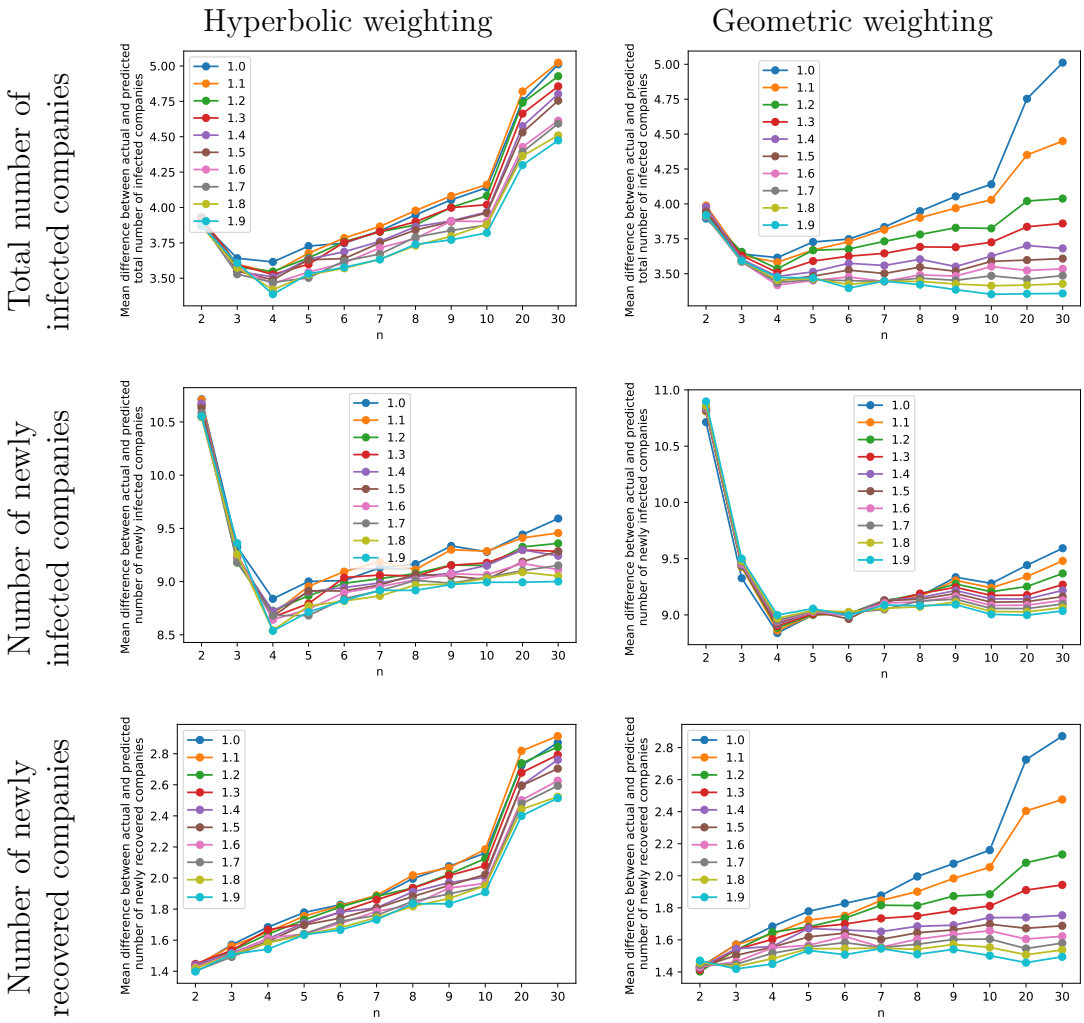

Figure 5.12: The plots show the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2020 financial crisis. The predictions are generated using the CPD+TiWML combination (see Algorithm 6), based on infection data from the previous n days and with a prediction horizon of $k = 20$ days. The left column shows the differences for the hyperbolic weighting scheme, while the right column shows the differences for the geometric weighting scheme, for various values of $\rho > 1$.

a weighting scheme in addition to the change point model improved predictions of the number of infected companies, here applying either hyperbolic or geometric weighting schemes leads to a lower mean Sørensen–Dice coefficient between the actual and predicted continents and sectors. Moreover, higher values of ρ consistently lead to poorer predictions across all values of k. In particular, for most values of n, the least accurate predictions are achieved

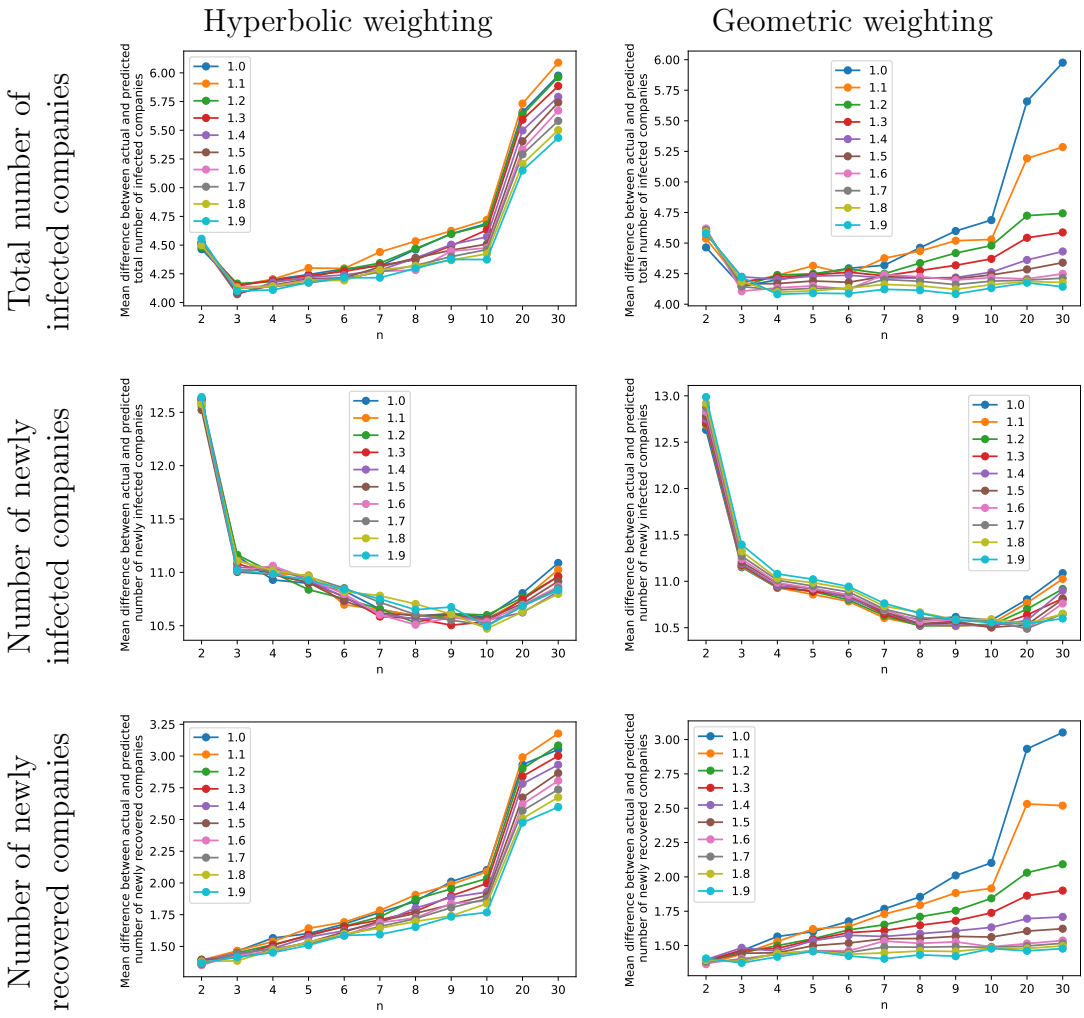

Figure 5.13: The plots show the absolute difference between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2020 financial crisis. The predictions are generated using the CPD+TiWML combination (see Algorithm 6), based on infection data from the previous n days and with a prediction horizon of $k = 30$ days. The left column shows the differences for the hyperbolic weighting scheme, while the right column shows the differences for the geometric weighting scheme, for various values of $\rho > 1$.

with geometric weighting and $\rho = 1.9$. Additionally, as n increases, prediction accuracy declines for both continents and sectors. Specifically, using the CPD+TiWML model with geometric weighting and large n results in more than 10% decline in mean Sørensen–Dice coefficient for both continents and sectors across all k. In addition, in Section 5.2.1.2 we observed that introducing either of the two weighting schemes, but without change point detection,

consistently led to more accurate predictions for $n \in \{20, 30\}$. In contrast, here, the larger the n, the poorer the predictions (see Appendix B for figures illustrating the relative percentage improvement in comparison to the change point model from Chapter 4). One possible explanation for larger n leading to poorer predictions is that when the model incorporates change point detection, data is already discarded and the weighting schemes are applied to a smaller number of days. Hence, the weighting scheme cannot successfully balance recent information with the broader historical data, leading to poorer predictions and loss of important data. Overall, the CPD+TiWML model works well for predicting the number of infected companies, but not for predicting the infected continents and sectors, suggesting that when it comes to predicting the spread across continents and sectors, historical data likely plays a more significant role.

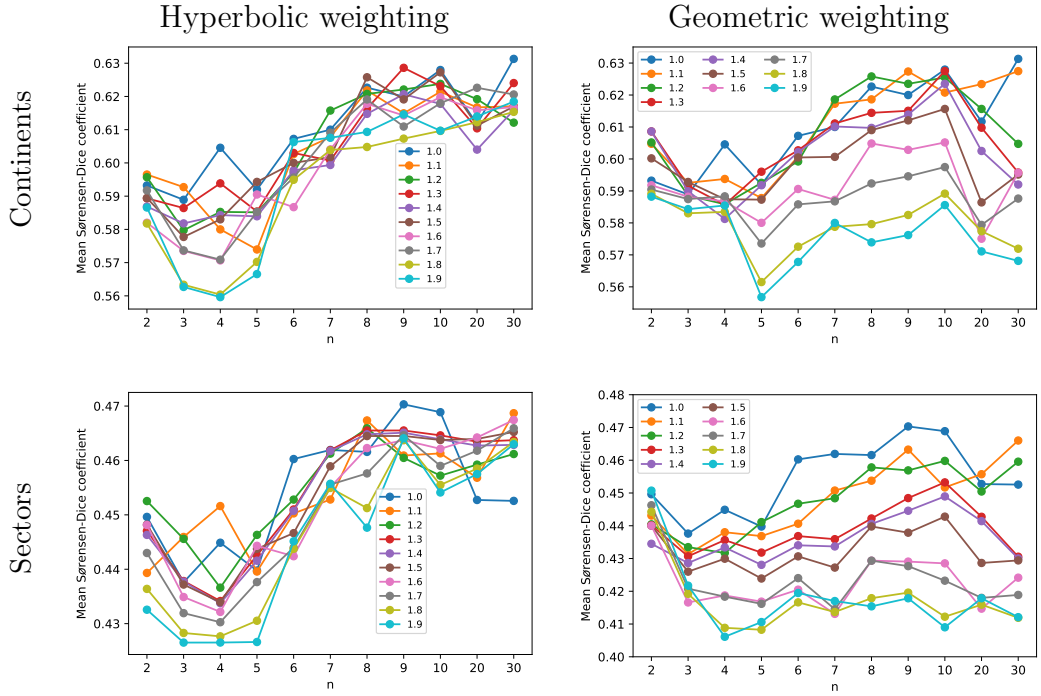

Figure 5.14: The plots show a comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 10$ days for the the 2020 financial crisis, using the infections data from the previous n days when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes.

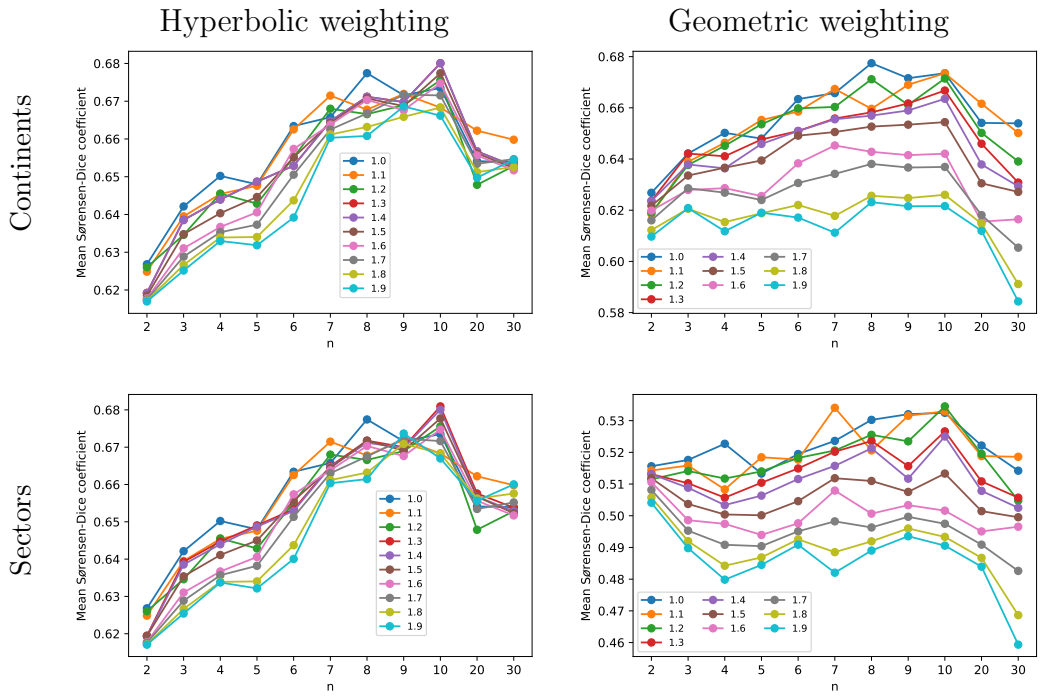

Figure 5.15: The plots show a comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 20$ days for the the 2020 financial crisis, using the infections data from the previous n days when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes.

5.3.1.2 Predicting the specific infected companies

In this section we evaluate the CPD+TiWML model's performance in predicting the specific infected companies. As in previous analyses, we calculate the mean Accuracy and F_1-scores by applying the CPD+TiWML with hyperbolic and geometric weighting schemes for different values of ρ. Figures 5.17, 5.18 and 5.19 present the mean a) Accuracy and b) F_1-score for different values of n and predicting the infected companies in the future $k = 10$, $k = 20$ and $k = 30$ days, respectively. The results suggest that, similarly to the TiWML model discussed earlier, introducing either the hyperbolic or geometric weighting schemes improves the Accuracy for most values of n, in comparison to the change point model from Chapter 4. Moreover, higher ρ yields higher Accuracy, implying that the CPD+TiWML model which emphasizes recent data

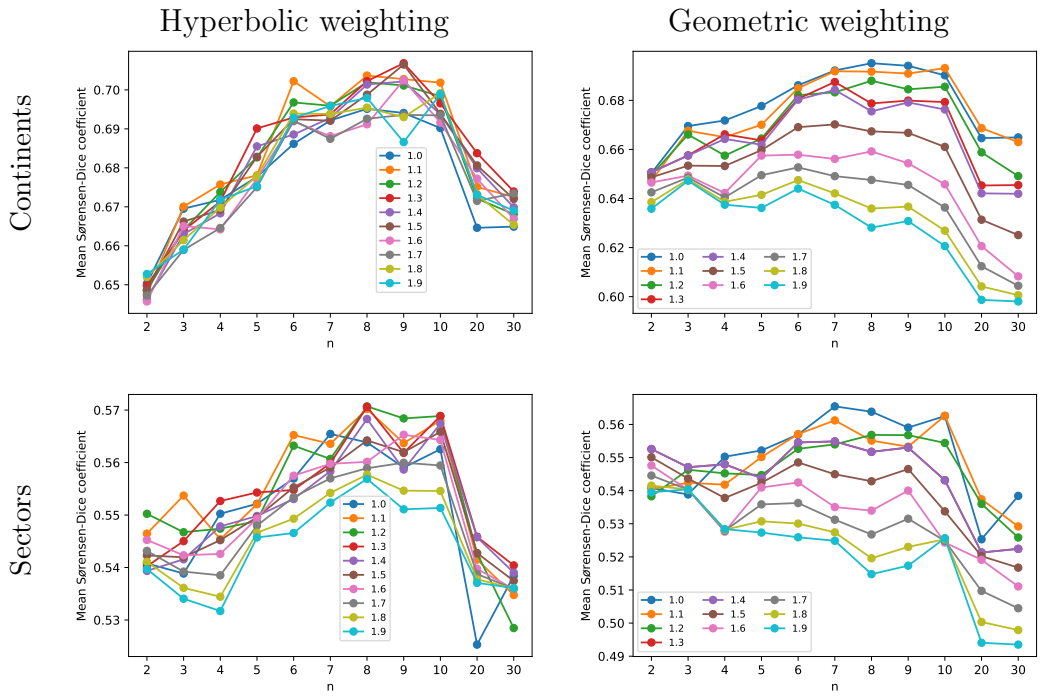

Figure 5.16: The plots show a comparison between the mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (top row) and sectors (bottom row) of newly infected companies in the future $k = 30$ days for the the 2020 financial crisis, using the infections data from the previous n days when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes.

more outperforms the change point model, which weights all the observations after a detected change point equally ($\rho = 1$). However, we should note that the highest improvement in Accuracy, compared to the non-weighted case, is only 0.6%, and is obtained with geometric weighting and $\rho = 1.9$ when $k = 30$ and $n = 20$.

Additionally, when $n \leq 20$ incorporating either weighting scheme in addition to the change point model typically leads to a lower F_1-score. On the other hand, $n \in \{20, 30\}$ combined with a large ρ leads to an improvement. Specifically, when $k = 10$ the highest improvement of nearly 4% in the F_1-score, compared to the non-weighted change point model from Chapter 4, is observed using hyperbolic weighting with $\rho = 1.9$ and $n = 30$. When $k = 20$ and $k = 30$ the highest improvements of 4% and 2%, respectively, are observed applying

geometric weighting with $\rho = 1.9$ and $n = 30$. Figures B.16 , B.17, and B.18 in Appendix B offer a visual representation of these percentage improvements.

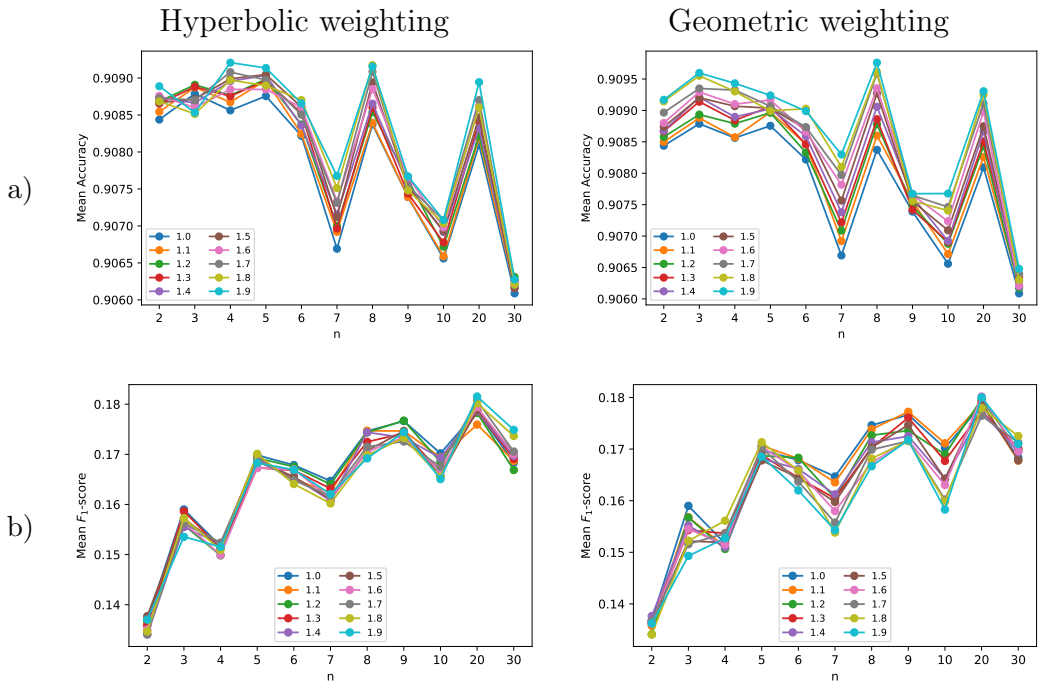

Figure 5.17: Comparison between the mean a) Accuracy and b) F_1-score in predicting the specific infected companies in the future $k = 10$ days for the the 2020 financial crisis, using the infections data from the previous n days when applying the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes.

5.4 Discussion and conclusion

In this chapter we introduced the TiWML and the CPD+TiWML models, aimed to enhance the model's predictive accuracy by applying different weighting schemes to the original model introduced in Chapter 3 and the change point model from Chapter 4, respectively. In particular, we employed two popular TiWML schemes, geometric and hyperbolic, in place of the conventional maximum likelihood method, where all observations are weighted equally. The models were evaluated based on the 2020 financial crisis simulation results. We observed the largest improvements when predicting the number of infected

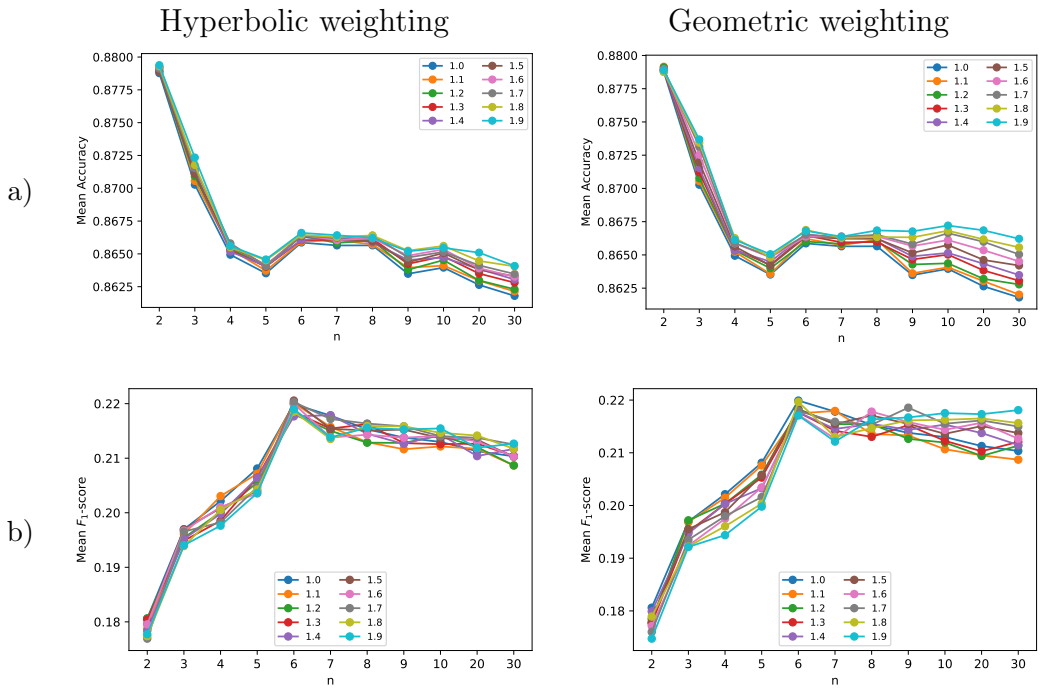

Figure 5.18: Comparison between the mean a) Accuracy and b) F_1-score in predicting the specific infected companies in the future $k = 20$ days for the the 2020 financial crisis, using the infections data from the previous n days when applying the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes.

companies. Specifically, in both the TiWML and CPD+TiWML models, the geometric weighting scheme consistently outperformed the hyperbolic weighting scheme, yielding substantially higher improvement in predictive accuracy, especially for higher values of the window length n.

While the TiWML and CPD+TiWML models showed clear improvements in many cases, there are several limitations and drawbacks that should be mentioned. Firstly, prioritising recent data might be beneficial in some cases, such as when predicting the number of infected companies, but in other cases it may overlook long-term trends or patterns that could be relevant, leading to poorer performance, as it was seen when applying the CPD+TiWML model for continent- and sector-specific predictions. In addition, the models' performance is heavily dependent on the choice of weighting scheme and the parameter ρ. While this chapter demonstrates the effectiveness of the

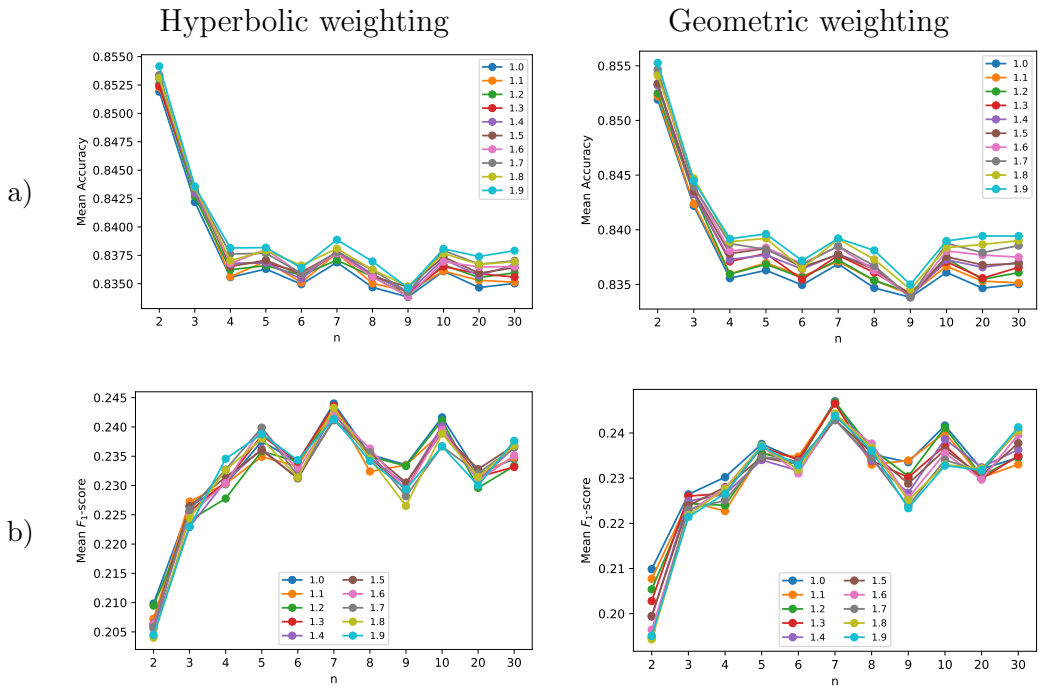

Figure 5.19: Comparison between the mean a) Accuracy and b) F_1-score in predicting the specific infected companies in the future $k = 30$ days for the the 2020 financial crisis, using the infections data from the previous n days when applying the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes.

observation-weighting approach using two popular weighting schemes, introducing an alternative weighting criterion could potentially lead to more accurate predictions. Hence, a future research direction is investigating other weighting strategies, such as adaptive or data-driven approaches that dynamically adjust weights based on changing patterns in the data. Additionally, combining different weighting schemes might further enhance model performance.

Chapter 6

The optimal model

In this chapter we compare all models introduced in the thesis: the original model from Chapter 3, the change point model from Chapter 4, the time-based weighted maximum likelihood (TiWML) model from Section 5.2 and the combination of the change point and the time-based weighted maximum likelihood (CPD+TiWML) model from Section 5.3. We assess each model for its performance across various prediction scenarios, such as forecasting the number of infected companies, the continents and economic sectors that will be most affected, and the specific companies most likely to be infected during future crisis days. The most effective model for each type of prediction is then identified based on the 2020 financial crisis simulation results.

6.1 Predicting the number of infected companies

In this section, for each combination of n and k, we present the model which yields most accurate predictions out of all models introduced in this thesis.

Figures 6.1, 6.2 and 6.3 show the most appropriate model choices for predicting a) the total number of infected companies, b) the number of newly infected companies and c) the number of newly recovered companies in the future $k = 10$, $k = 20$ and $k = 30$ days, respectively. The marker shapes represent the models as follows: circles for the TiWML model with hyperbolic weighting, crosses for the TiWML model with geometric weighting, triangles for CPD + TiWML model with hyperbolic weighting, squares for CPD + TiWML model with geometric weighting, a star for the change point model and a ring for the original model from Chapter 3. The marker colours represent different values of the parameter ρ. We remark that for $n = 1$ the results for all models are the same. As noted previously, a change point can be detected at least one day after its occurrence. Consequently, when $n = 1$, a detected change point aligns with the first day of the sliding window and hence no data will be discarded, equivalent to using the original model from Chapter 3. Similarly, having only one data observation, weighting is not feasible. Hence, for $n = 1$, we plot the predictions generated using the original model.

Figures 6.1a, 6.2a and 6.3a show that when predicting the total number of infected companies for all values of k and large values of n the most appropriate model choice is the CPD + TiWML model with geometric weighting and $\rho = 1.9$. For smaller values of n, however, there is not a consistent trend. In more detail, when $k = 10$, the most accurate predictions are achieved for $n \geq 4$, with best result yielded for $n = 4$ in combination with the CPD + TiWML model with geometric weighting and $\rho = 1.1$. When $k = 20$, the optimal choice is $n \in \{10, 20, 30\}$, combined with the CPD + TiWML model with geometric weighting and $\rho = 1.9$. Finally, when $k = 30$, the most accurate predictions are obtained when $n \geq 3$, with best model choice being $n = 3$ with CPD + TiWML model with hyperbolic weighting and $\rho = 1.4$.

Additionally, Figures 6.1b and 6.3b illustrate the optimal model choices for

predicting the number of newly infected companies in the future $k = 10$ and $k = 30$ days, respectively. For both values of k, when $n \in \{20, 30\}$, the optimal model choice is the CPD + TiWML model with geometric weighting, while for medium values, $n \in [6, 10]$, the optimal choice is the CPD + TiWML model with hyperbolic weighting. However, when $k = 20$ (see Figure 6.2b) for most values of n the optimal model choice is the TiWML model with geometric weighting scheme and $\rho = 1.9$. Overall, for all values of k, more accurate predictions are achieved when $n \geq 4$. Specifically, for $k = 10$, the most accurate predictions are achieved when $n = 7$ with the CPD + TiWML model with hyperbolic weighting and $\rho = 1.3$. For $k = 20$, the best result is yielded when $n = 4$, in combination with the TiWML model with hyperbolic weighting and $\rho = 1.8$. Finally, for $k = 30$, the most accurate predictions are achieved when $n = 20$ with the CPD + TiWML model with geometric weighting and $\rho = 1.6$.

Finally, we present the most appropriate model choices for predicting the number of newly recovered companies in the future $k = 10$ (Figure 6.1c), $k = 20$ (Figure 6.2c) and $k = 30$ (Figure 6.3c) days. Notably, for all values of k and $n \geq 3$, the best results are achieved with the CPD + TiWML model with geometric weighting and $\rho = 1.9$. In addition, for all k, larger values of n lead to less accurate predictions, with best results achieved when $n = 1$.

6.2 Sector- and continent-specific predictions

In this section we present the optimal model choices when predicting the infected companies' continents and sectors. Figure 6.4 displays the models which lead to highest mean Sørensen–Dice coefficient between predicted and actual continents (left column) and sectors (right column) of the newly in-

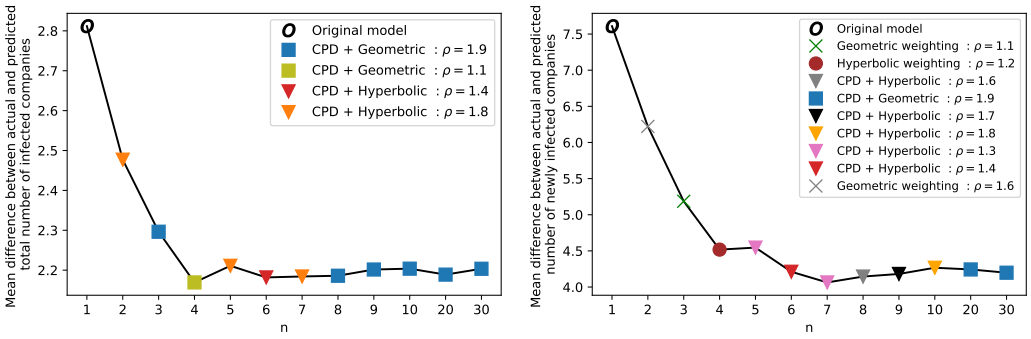

(a)

(b)

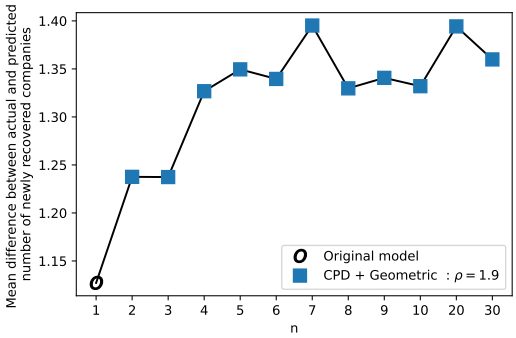

(c)

Figure 6.1: The figure presents the optimal model choices for predicting a) total number of infected, b) newly infected and c) newly recovered companies in the future $k = 10$ days for each value of n. The marker shapes represent the models as follows: circles for hyperbolic weighting, crosses for geometric weighting, triangles for CPD + hyperbolic weighting and squares for CPD + geometric weighting. The marker colours correspond to different ρ values.

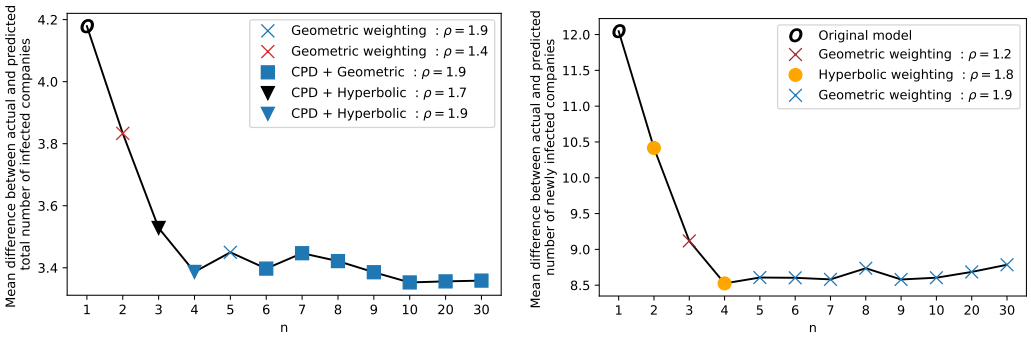

(a)

(b)

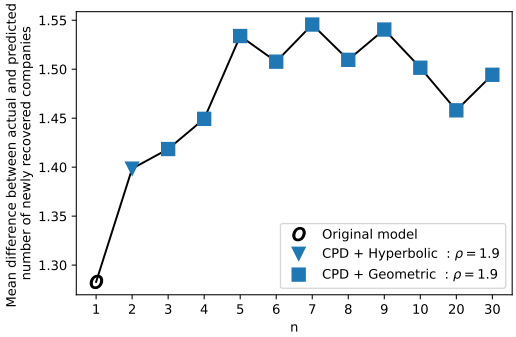

(c)

Figure 6.2: The figure presents the optimal model choices for predicting a) total number of infected, b) newly infected and c) newly recovered companies in the future $k = 20$ days for each value of n. The marker shapes represent the models as follows: circles for hyperbolic weighting, crosses for geometric weighting, triangles for CPD + hyperbolic weighting and squares for CPD + geometric weighting. The marker colours correspond to different ρ values.

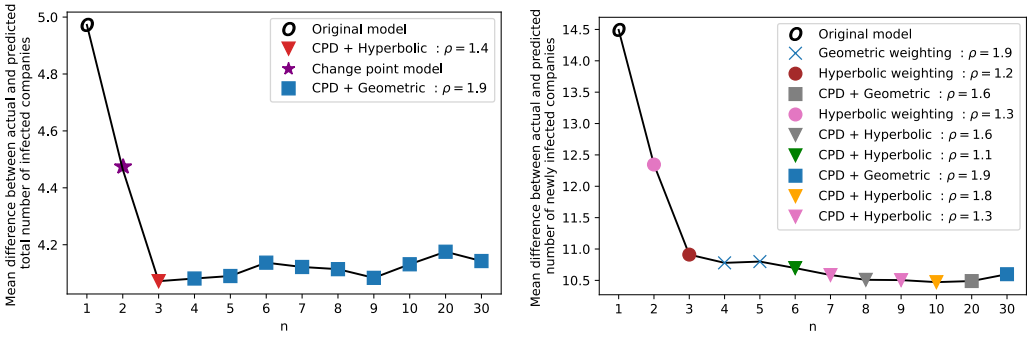

(a) (b)

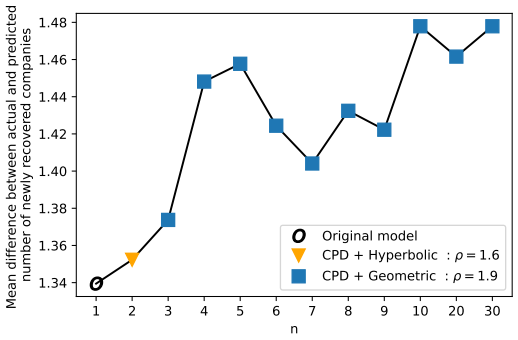

(c)

Figure 6.3: The figure presents the optimal model choices for predicting a) total number of infected, b) newly infected and c) newly recovered companies in the future $k = 30$ days for each value of n. The marker shapes represent the models as follows: circles for hyperbolic weighting, crosses for geometric weighting, triangles for CPD + hyperbolic weighting, squares for CPD + geometric weighting and a star for the change point model. The marker colours correspond to different ρ values.

fected companies in the future $k = 10$ (top row), $k = 20$ (middle row) and $k = 30$ (bottom row) days for the the 2020 financial crisis, using the infections data from the previous n days. It can be seen that, in comparison with predicting the number of infected companies, where the most appropriate models include weighting with higher values of ρ, indicating greater importance of more recent data, here in almost all cases the values of ρ range from 1.1 to 1.3. In addition, when predicting the number of infected companies in many cases the optimal model includes the geometric weighting scheme, usually in combination with the change point detection. In contrast, here the optimal combination in most cases consists of the hyperbolic weighting and the change point detection. The findings suggest that incorporating a wider range of historical data after a change point, rather than focusing solely on recent trends, is more effective for geographic- and sector-specific predictions. In addition, for $k = 10$ the best result for both continents and sectors is obtained with $n = 30$, with the optimal model choices being the change point model and the TiWML model with hyperbolic weighting and $\rho = 1.1$, respectively. In contrast, when $k = 20$ or $k = 30$, values of $n \in \{20, 30\}$ lead to a decrease in the mean Sørensen–Dice coefficient. For $k = 20$ the most accurate predictions are obtained with $n = 10$, with optimal models being the CPD+TiWML with hyperbolic weighting and $\rho = 1.4$ and $\rho = 1.3$, respectively, for predicting the continents and sectors. In addition, for $k = 30$ the highest mean Sørensen–Dice coefficients are obtained when $n = 9$, combined with CPD+TiWML model with hyperbolic weighting and $\rho = 1.3$ and $\rho = 1.2$, when predicting continents and sectors, respectively. For all values of k, the least accurate predictions are obtained for $n = 1$, highlighting that a single day's worth of data is insufficient for reliable predictions. Overall, the consistent success of the CPD+TiWML model with hyperbolic weighting and a small ρ indicates that while incorporating change point detection and discarding the data prior to it is beneficial for sector- and continent-specific predictions, much of the data

after the change point remains valuable. Rather than heavily prioritising the most recent data, a balanced weighting of post-change point data contributes to more accurate predictions in these contexts.

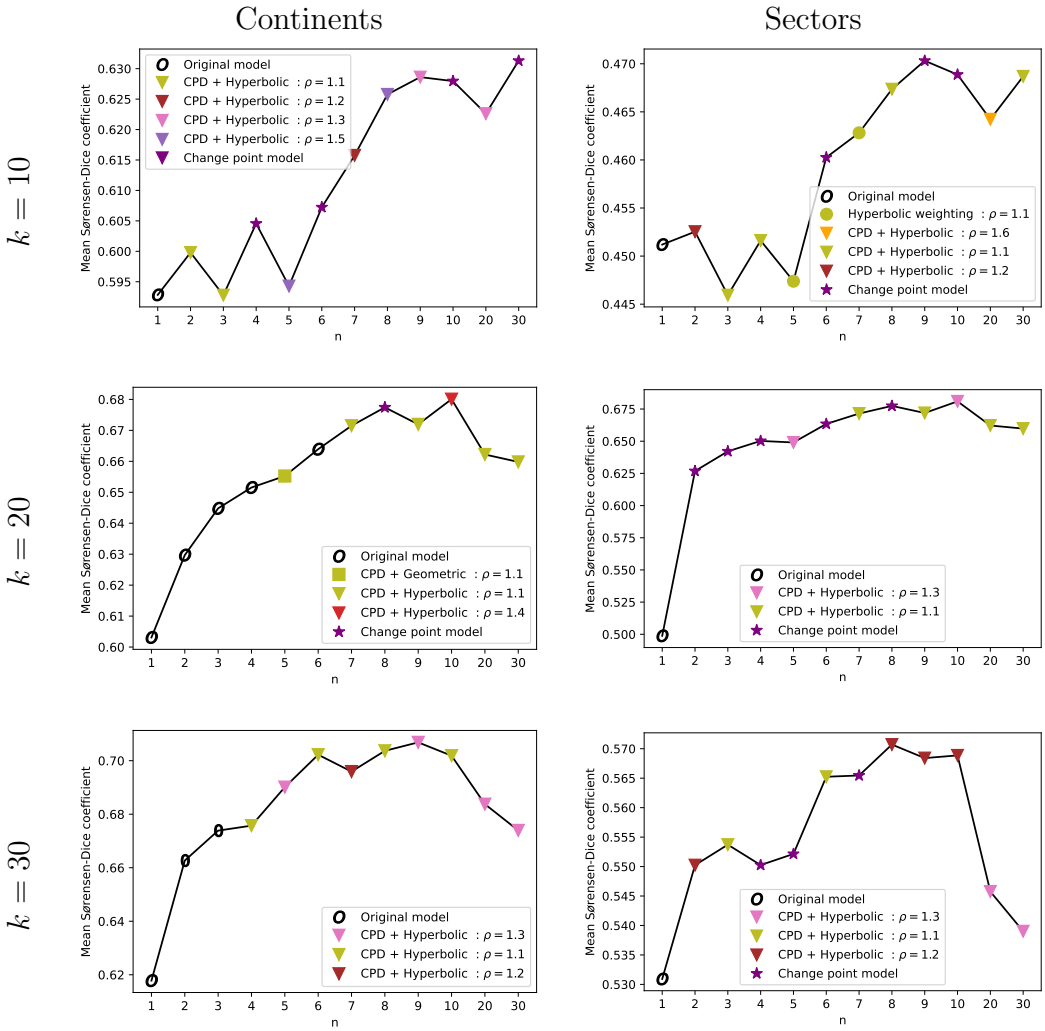

Figure 6.4: The plots show the model resulting in the highest mean Sørensen–Dice coefficient, averaged over all prediction windows, between predicted and actual continents (left column) and sectors (right column) of the newly infected companies in the future $k = 10$ (top row), $k = 20$ (middle row) and $k = 30$ (bottom row) days for the the 2020 financial crisis, using the infections data from the previous n days. The marker shapes represent the models as follows: circles for hyperbolic weighting, crosses for geometric weighting, triangles for CPD + hyperbolic weighting, squares for CPD + geometric weighting, a star for the change point model and a ring for the original model. The marker colours correspond to different ρ values.

6.3 Predicting the specific infected companies

We now present the models which perform best in predicting the specific infected companies in the future $k = 10$, $k = 20$ and $k = 30$ crisis days. Figure 6.5 shows the optimal model choice based on the Accuracy (left column) and F_1-score (right column) for the prediction horizons of $k = 10$ (top row), $k = 20$ (middle row) and $k = 30$ (bottom row) crisis days. For each studied value of n we show the model yielding the highest Accuracy and F_1-score. It can be seen that across all values of k and most values of n the highest Accuracy is achieved with the CPD+TiWML model with geometric weighting and $\rho = 1.9$. In contrast, when comparing the models based on F_1-score there is no obvious pattern. However, notably for $n \in \{20, 30\}$, each optimal model, based on both Accuracy and F_1-score, uses parameter $\rho = 1.9$, while for smaller n, lower values of ρ are preferred. Overall, across all values of k the highest Accuracy is consistently achieved by the original model from Chapter 3 with $n = 1$, with decreasing Accuracy as n increases. On the other hand, the same model yields lowest F_1-scores, with higher values of n generally improving the F_1-score. The results show that incorporating more historical data might give important information to the model, hence improving the F_1-score, while using only recent data including discarding the data prior to a detected change point, a large ρ and a small n enhances Accuracy.

6.4 Discussion and conclusion

This chapter presented the most effective models, based on the 2020 financial crisis simulation results, for predicting the number of infected companies, as well as for continents-, sectors- and company-specific predictions, among all models introduced in this thesis. Overall, the results indicate that in many

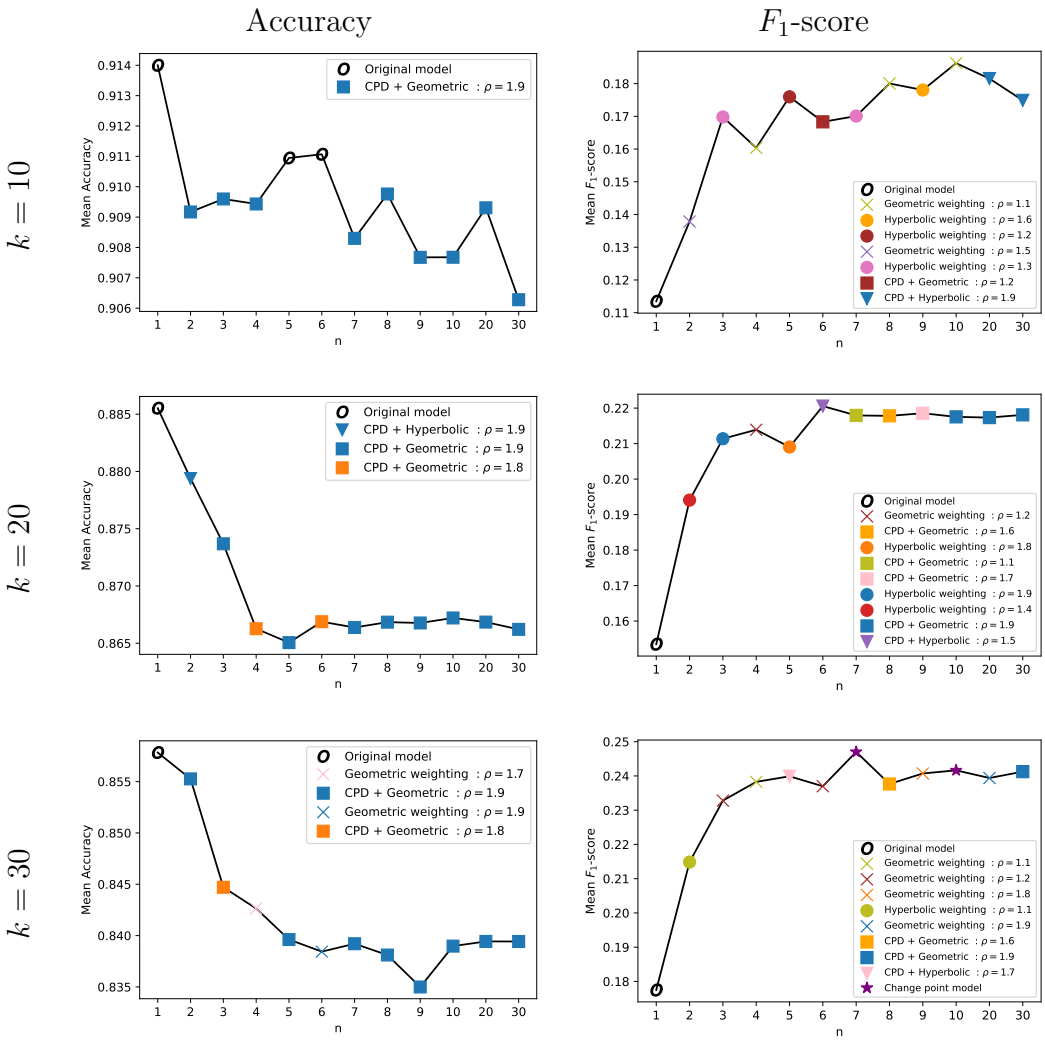

Figure 6.5: The plots show the model resulting in highest mean Accuracy (left column) and F_1-score (right column) when predicting the specific infected companies in the future $k = 10$ (top row), $k = 20$ (middle row) and $k = 30$ (bottom row) days for the the 2020 financial crisis, using the infections data from the previous n days. The marker shapes represent the models as follows: circles for hyperbolic weighting, crosses for geometric weighting, triangles for CPD + hyperbolic weighting, squares for CPD + geometric weighting, a star for the change point model and a ring for the original model. The marker colours correspond to different ρ values.

cases the optimal model choice is dependent on k (the prediction horizon) and n (the amount of historical data used), as well as on the metric of interest.

One of the key takeaways is that there is no single model that performs best in all scenarios. Hence, the model choice should be tailored based on the predictions of interest. This implies that different prediction tasks, whether pre-

dicting the number of infected companies or sector- , continent- or company-specific predictions, may require distinct approaches to achieve more accurate forecasts. In addition, it is important to select models that strike a balance between various performance metrics to achieve optimal results. In this context, a direction for future research will be to develop adaptive models which can adjust based on the specific nature of the task or achieving a trade-off between different prediction objectives.

Chapter 7

The impact of network and node characteristics on risk propagation

7.1 Introduction

Studying the effect of network structure and node characteristics on risk propagation has been a hot research topic in the past few years [65]. Centrality measures and the role of highly connected nodes in a network has become a key research area, as they provide insights into how risk spreads [68]. In this chapter we explore these topics, studying different node characteristics and their influence on crisis propagation. Specifically, in Section 7.2 we introduce the notion of contagion distance and contagion centrality, closeness centrality and betweenness centrality and study their correlation with the probability of a company becoming infected. The results show that there is indeed a positive correlation between centrality and the probability of infection, indicating that more central nodes are more likely to be infected. Then, in Section 7.3, we

analyse the relationship between various company's characteristics, including credit scores from the Bureau van Dijk dataset, and the length of time a company remained infected during the 2020 financial crisis. The findings suggest that a credit score of a company can be indicative of how long it would be infected during the crisis, but a previous infection is not. Finally, Section 7.4 presents several simulation studies exploring various counterfactual scenarios, such as different lockdown strategies during the 2020 financial crisis, starting a crisis with the most central nodes and changing the number of initially infected companies. The simulations highlight the crucial role that initial conditions and intervention strategies play in financial crises propagation.

7.2 The impact of node centrality on infection probability

The interconnectedness of companies within complex financial networks can amplify the impact of a crisis, as highly connected nodes can spread the risk to a broader range of companies [193]. In this context, node centrality has emerged as one of the key concepts in the analysis of complex financial networks, providing insights into which entities are of key importance for the propagation of risk [27]. In this section we explore the correlation between a company's resilience to financial shocks and different node centrality measures. Specifically, we focus on the notions of contagion distance and contagion centrality (7.2.1), closeness centrality (7.2.2) and betweenness centrality (7.2.3), and analyse their correlation with the probability of a company becoming infected. The determination of the most central agents is particularly challenging in systems with multiple types of interactions, resulting in interconnected multilayer networks, such as our network structure, introduced in

Chapter 3. In this section, we first introduce a mathematical framework for calculating centrality within multiplex networks. We then rank the nodes based on their centrality, thereby identifying the most influential companies. By understanding these centrality measures, we gain insights into how the position and influence of a company within the network affect its ability to withstand and recover from disruptions. The findings suggest that nodes with higher contagion and closeness centrality are more likely to be infected, while betweenness centrality shows a weaker correlation with infection probability.

7.2.1 Contagion distance and contagion centrality

In this section we present two measures introduced for single-layered networks in [1], contagion distance and contagion centrality. In particular, in Section 7.2.1.1 we introduce the notion of contagion distance and we extend its application to multiplex networks. We then study the contagion distance between pairs of companies grouped by continent and sector, with results revealing that companies within the same continent and sector have lower contagion distance between each other, implying a higher probability of transmitting infection to one another. Then, in Section 7.2.1 we define contagion centrality, indicating how close a node is to the rest of the network in terms of contagion distances, and study its correlation with infection probability. The results show a moderate positive correlation for both the 2008 and the 2020 financial crises, indicating that nodes with higher contagion centrality are in general more likely to become infected during the crisis.

7.2.1.1 Contagion distance

The contagion distance $d_c(i,j)$ between two nodes i and j in a network is the distance on the path $\gamma_{i,j}$ that finds a tradeoff between these two factors: minimizing the length $H(\gamma_{i,j})$ of the path $\gamma_{i,j}$ and maximizing the log probability of shock transmission along the path $\gamma_{i,j}$, that is:

$$\min_{\gamma_{i,j}} \left[H(\gamma_{i,j}) - \sum_{(i_{c-1},i_c)\in\mathcal{E}_{\gamma_{i,j}}} \log w_{i_{c-1},i_c} \right]. \tag{7.1}$$

Here w_{i_{c-1},i_c} is the weight on the edge (i_{c-1},i_c) (i.e., the probability of transmission on the edge (i_{c-1},i_c)), $H(\gamma_{i,j})$ is the number of edges on the path $\gamma_{i,j}$ and $\mathcal{E}_{\gamma_{i,j}}$ is the set of edges on the path $\gamma_{i,j}$. In more detail, the path $\gamma_{i,j}$ is an ordered sequence of $H(\gamma_{i,j})+1$ distinct nodes $\mathcal{V}_{\gamma_{i,j}} = \{i = i_0, \ldots, i_{H(\gamma_{i,j})} = j\}$, connected by $H(\gamma_{i,j})$ edges between them $\mathcal{E}_{\gamma_{i,j}} = \{(i = i_0, i_1), \ldots, (i_{H(\gamma_{i,j})-1}, i_{H(\gamma_{i,j})} = j)\}$, where each edge links a pair of consecutive nodes in the sequence.

The reasoning behind the definition of contagion distance is that if node i_{c-1} transmits to node i_c through an edge (i_{c-1},i_c) with probability w_{i_{c-1},i_c}, then the log probability that an infection transmits through the whole path $\gamma_{i,j}$ is the log of the product of the probabilities on all edges on the path $\gamma_{i,j}$, that is:

$$\log\left(\prod_{(i_{c-1},i_c)\in\mathcal{E}_{\gamma_{i,j}}} w_{i_{c-1},i_c} \right) = \sum_{(i_{c-1},i_c)\in\mathcal{E}_{\gamma_{i,j}}} \log w_{i_{c-1},i_c}. \tag{7.2}$$

The notion of contagion distance can be easily generalised for multiplex networks with l layers by representing them as single-layered as follows. For each node i in layer α of the multiplex network, create a new node $i^{[\alpha]}$ in the single-layered network. This transformation is done for every node and layer combination. Hence, the set of of nodes in the single-layered network is $V' = \{i^{[\alpha]} \mid i \in V, \alpha \in \{1,\ldots,l\}\}$, where V is the set of nodes in the original

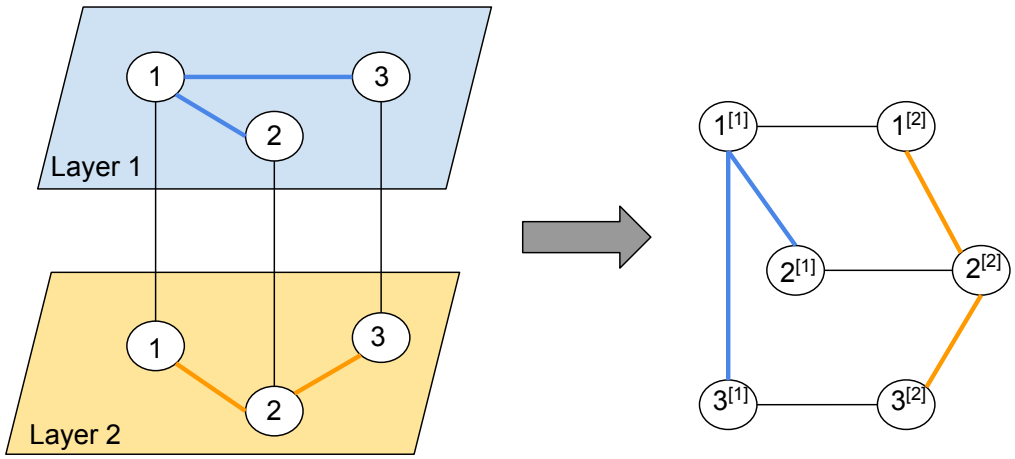

Figure 7.1: Representation of a multiplex network as a single layered network.

multiplex network. Then, for each edge (i, j) in layer α, create an edge between $i^{[\alpha]}$ and $j^{[\alpha]}$ in the single-layered network. Additionally, for inter-layer connections, if there is an edge between i in layer α and j in layer β, create an edge between $i^{[\alpha]}$ and $j^{[\beta]}$ in the single-layered network. An example is shown in Figure 7.1.

The contagion distance $d_c(i^{[\alpha]}, j^{[\beta]})$ between each pair of nodes $i^{[\alpha]}$ and $j^{[\beta]}$ for $1 \leq \alpha, \beta \leq l$ is calculated using equation (7.1). The contagion distance $d_c(i, j)$ between nodes i and j in the multiplex network can then be found by taking the minimum of the contagion distances between all pairs of nodes $i^{[\alpha]}$ and $j^{[\beta]}$ in the single-layered network, that is, $d_c(i, j) = \min\{d_c(i^{[\alpha]}, j^{[\beta]}) | 1 \leq \alpha, \beta \leq l\}$.

Contagion distances are calculated based on the underlying multiplex network structure, introduced in Chapter 3. We first represent the multiplex network as a single-layered one using the procedure described above. We then apply Dijkstra's shortest path algorithm [74] to find a shortest path from each node $i^{[\alpha]}$ to every other node $j^{[\beta]}$, where $1 \leq \alpha, \beta \leq 4$, by minimising the contagion distance as per equation (7.1). The edge weights are determined using the daily transmission probabilities, as introduced in Chapter 3. That is, we fit the model to the period from day t to day $t+1$ (window length $n = 1$) to get

estimates of $\hat{\beta}_i(t)$ $(1 \leq i \leq 4)$ for each day $1 \leq t \leq L-1$ where L is the length of the crisis in days. Therefore, for each layer $1 \leq \alpha \leq 4$, we get daily weights, as defined in (3.4). In other words, we end up with a dynamic network with edge weights changing over time.

We calculate the contagion distance between each pair of companies for each day of the 2008 and 2020 financial crises. Figures 7.2a and 7.2b show heatmaps of the mean contagion distance between each pair of companies, averaged over all sliding windows, and grouped by sector, in the 2008 and 2020 financial crises, respectively. In the 2020 financial crisis the shortest contagion distances (highest probability of contagion) occur between companies in the same sector, in the 2008 financial crisis there is no clear pattern to be seen. This indicates that during the 2020 financial crisis, economic sector played a more significant role in contagion spread than during the 2008 financial crisis. Furthermore, in the 2008 crisis, the greatest contagion distances are between companies in the Chemicals sector, suggesting that these companies were the least likely to spread the infection to each other.

Similar heatmaps, but for the companies grouped by continent, are presented in Figures 7.3b and 7.3b. It can be seen that in both financial crises the shortest contagion distances are between companies in the same continent, while the highest average contagion distances are between pairs of companies, for which one of the companies is in Asia. This observation suggests that companies on the same continent are more likely to transmit infection between each other, whereas companies in Asia were least likely to spread infection to or be infected by companies in other continents.

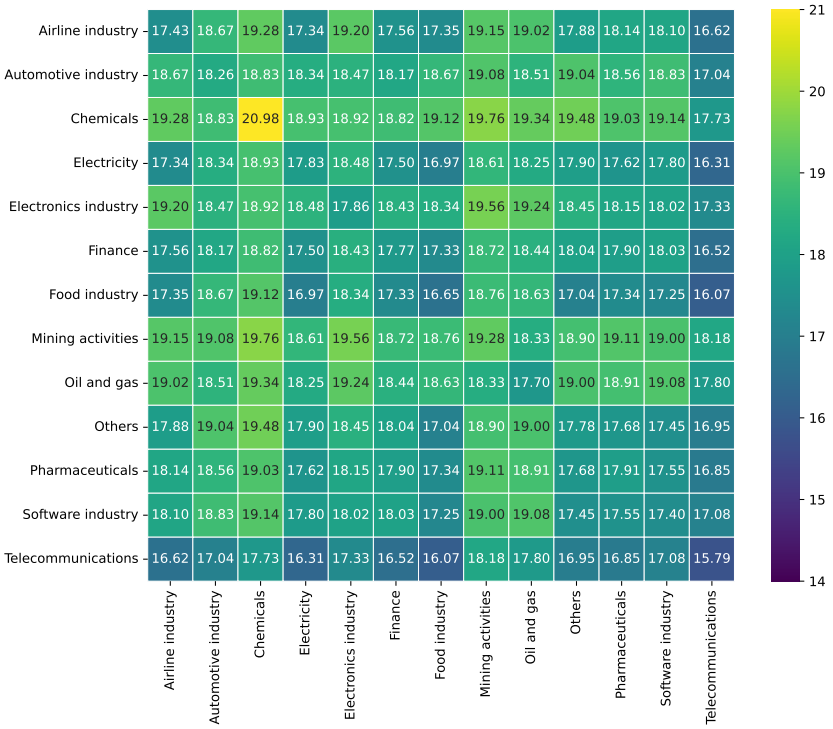

(a)

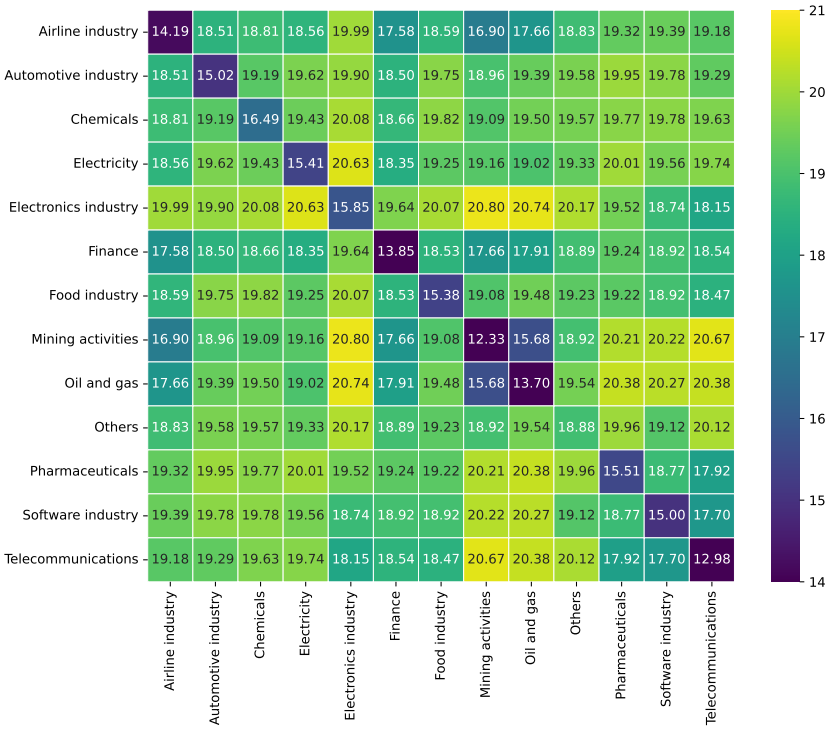

(b)

Figure 7.2: Heatmaps showing the mean contagion distance, averaged over all crisis days, in the (a) 2008 and (b) 2020 financial networks, grouped by sectors. Each cell in the heatmap represents the average contagion distance over all crisis days and all pairs of distinct nodes where one node is from the sector indicated by the row and the other is from the sector indicated by the column.

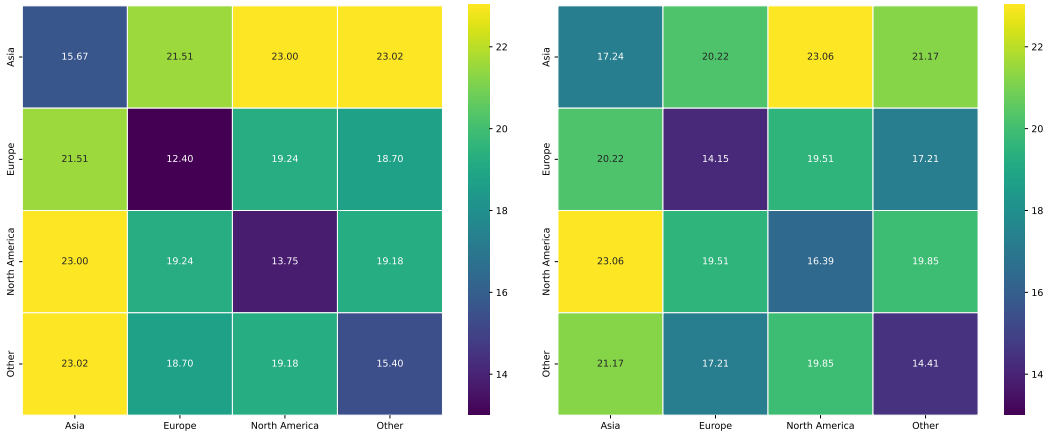

(a) (b)

Figure 7.3: Heatmaps showing the mean contagion distance, averaged over all crisis days, in the (a) 2008 and (b) 2020 financial networks, grouped by continents. Each cell in the heatmap represents the average contagion distance over all crisis days and all pairs of distinct nodes where one node is from the continent indicated by the row and the other is from the continent indicated by the column.

7.2.1.2 Contagion centrality

The contagion centrality of node i, as defined in [1] is:

$$CC_i = \frac{1}{\sqrt{\mu_i^2 + \sigma_i^2}}, \tag{7.3}$$

where

$$\mu_i = \frac{\sum_{j \neq i} d_c(i,j)}{N-1} \text{ and } \sigma_i = \sqrt{\frac{\sum_{j \neq i} (d_c(i,j) - \mu_i)^2}{N-2}}. \tag{7.4}$$

Here N is the number of nodes in the network, μ_i and σ_i are the sample mean and sample standard deviation of contagion distances $d_c(i,j)$ of node i to all other nodes j. The contagion centrality measure defined in equation (7.3) indicates the closeness of that node to the rest of the network in terms of contagion distances. A more central node has higher contagion centrality value and is highly effective at spreading the shock while also being more vulnerable to shocks itself.

For each company we estimate the probability of it becoming infected in each

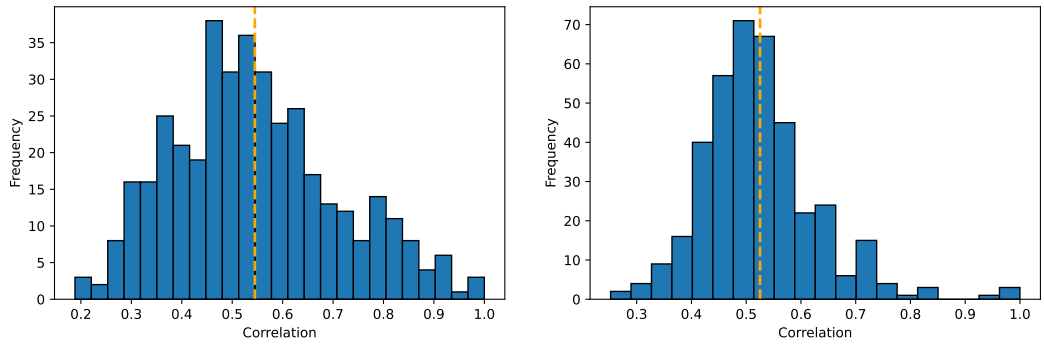

(a) (b)

Figure 7.4: Histograms showing the Pearson correlation between probability of a company becoming infected and its contagion centrality in the a) 2008 and b) 2020 financial crises. The dashed vertical line marks the mean correlation across all companies.

crisis day, using the procedure described in Section 3.4.2.2.1 and we also calculate its contagion centrality. We then, for each company, compute the Pearson correlation [176] between the estimated probability of infection and its contagion centrality. The results for both crises are shown in Figure 7.4. The dashed vertical line illustrates the mean correlation across all companies. It can be seen that in both crises there is a moderate positive correlation between the probability of a company being infected and its contagion centrality, with most values lying in the interval between 0.4 and 0.6.

7.2.2 Closeness centrality

The closeness centrality of node i, as defined in Salter-Townshend et al. [194], is given by

$$C(i) = \frac{N-1}{\sum_{j=1}^{N} d(i,j)}, \tag{7.5}$$

where N is the number of nodes and $d(i,j)$ denotes the geodesic distance (the length of the shortest path) between nodes i and j. Specifically, the shortest path between nodes i and j is the one with the minimum sum of edge weights among all paths starting from i and ending at j. Since, $d(i,j)$ is finite if and

only if nodes i and j are connected (for two vertices that are not connected in a graph, the geodesic distance is defined as infinite), the closeness centrality measure is meaningful only for connected networks.

We transform the weights in our network using the following transformation $w_{i,j}^{[\alpha]} = 1 - \log w_{i,j}^{[\alpha]}$, where $w_{i,j}^{[\alpha]}$ denotes the weight of edge (i, j) in layer α, $(i^{[\alpha]}, j^{[\alpha]})$ (see Section 3.3.2 for more details). This transformation ensures that the weight $w_{i,j}^{[\alpha]}$ reflects the inverse relationship between infection probability and distance, where a higher probability of infection between two nodes corresponds to a shorter distance between them. With this transformation, the contagion distance measure $d_c(i, j)$ (see equation (7.1)), between nodes i and j, is equivalent to the geodesic distance $d(i, j)$ between nodes i and j. Therefore, the closeness centrality measure $C(i)$ can be written as $C(i) = 1/\mu_i$, where μ_i was introduced in (7.4). Hence, the closeness centrality of node i can be considered as the inverse of the average contagion distance from node i to all other nodes. The higher the average contagion distance, the less central the node is within the network.

Figure 7.5 shows a histogram illustrating the Pearson correlation between company's closeness centrality and its probability of becoming infected. The results show that in both crises, there is a strong positive correlation, with average correlation around 0.57 in the 2008 financial crisis and 0.69 in the 2020 financial crisis. These findings suggest that closeness centrality plays a role in susceptibility to infection. In particular, more central nodes seem to be more likely to get infected.

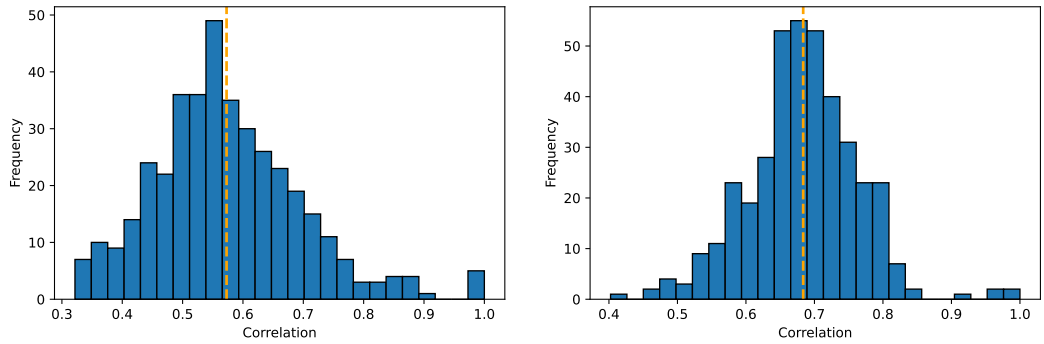

(a) (b)

Figure 7.5: Histograms showing the Pearson correlation between probability of a company becoming infected and its closeness centrality in the a) 2008 and b) 2020 financial crises. The dashed vertical line marks the mean correlation across all companies.

7.2.3 Betweenness centrality

The betweenness centrality of node i shows how important the node is for connecting other nodes in the network, and is given by:

$$C_B(i) = \sum_{j \neq k} \frac{\sigma_{jk}(i)}{\sigma_{jk}} \frac{1}{\binom{N}{2}} \tag{7.6}$$

where σ_{jk} is the number of shortest paths between j and k and $\sigma_{jk}(i)$ is the number of shortest paths between j and k that contain i [194]. We study if the betweenness centrality of a node is correlated with its probability of becoming infected. Figure 7.6 illustrates a histogram of the Pearson correlation between company's betweenness centrality and its probability of becoming infected. It can be seen that, compared to Figures 7.4 and 7.5, the correlation between betweenness centrality and infection probability is more dispersed, with a mean value near 0.3 in both crises. Moreover, unlike the correlations between infection probability and the contagion and closeness centrality measures, in Figure 7.6 there is also some negative correlation, indicating that nodes with high betweenness centrality do not always correspond to higher probability of infection.

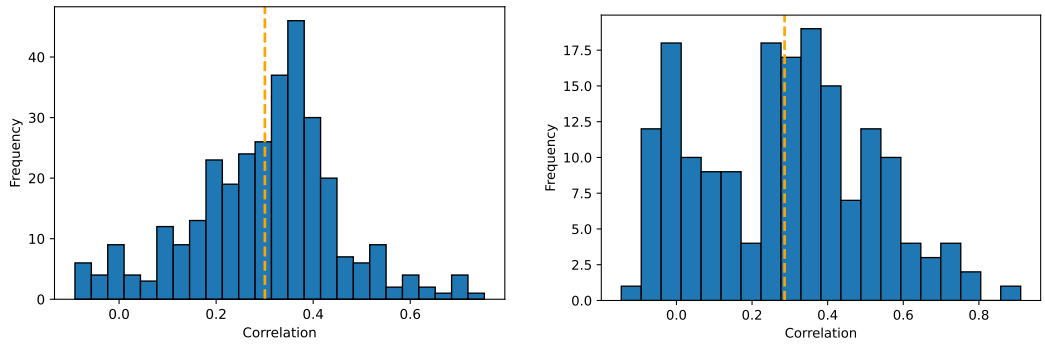

(a) (b)

Figure 7.6: Histograms showing the Pearson correlation between probability of a company becoming infected and its betweenness centrality in the a) 2008 and b) 2020 financial crises. The dashed vertical line marks the mean correlation across all companies.

7.3 Company's characteristics and recovery time during financial crises

In this section we examine potential indicators that might be related to the duration of a company's infection during a crisis. We first study whether the duration a company remained infected during the 2008 financial crisis is related to how long it stayed infected during the 2020 financial crisis. For each company that was infected in both crises we compare the time (in days) from the first day of infection to recovery. Figure 7.7 shows the number of days from infection to recovery in the 2020 financial crisis (x-axis) and 2008 financial crisis (y-axis) for each company that was infected in both crises. It can be seen that there is no obvious correlation between the recovery period in 2008 and 2020 financial crisis, as indicated by a Pearson correlation coefficient of -0.1408. The weak negative correlation shows that the duration of recovery in one crisis does not reliably predict the duration of recovery in the other crisis. This suggests that the factors influencing a company's ability to rebound from a financial crisis in 2008 may not be the same as those affecting its recovery in the 2020 financial crisis.

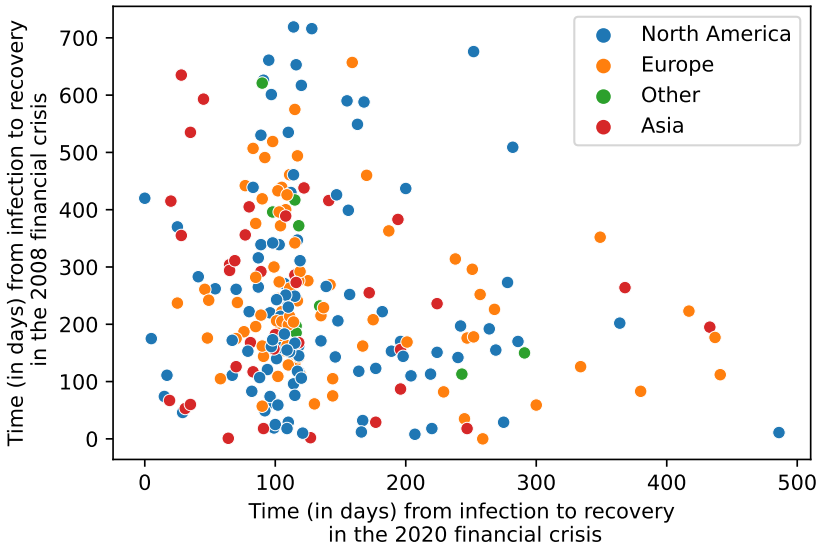

(a)

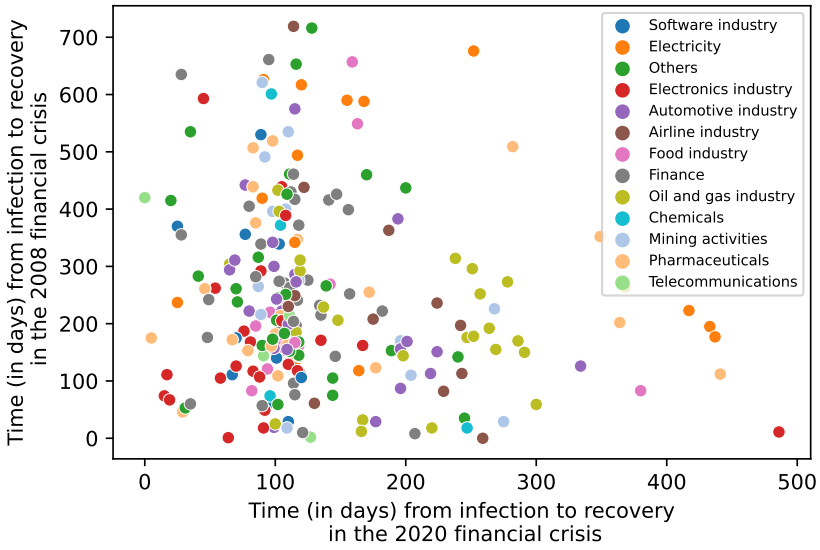

(b)

Figure 7.7: The number of days from infection to recovery in the 2020 financial crisis (x-axis) and 2008 financial crisis (y-axis) for each company that was infected in both 2008 and 2020 financial crisis. The companies are coloured by a) continent and b) sector.

We now study the relationship between different companies' characteristics given in the Bureau van Dijk dataset and see if any of them can be used as a predictor of a company's healthiness during a specific crisis. As the Bureau van Dijk dataset goes as back only as far as 2019, the following analysis considers only the 2020 financial crisis. Specifically, we study if each specific company's performance before the start of the crisis is indicative of how long it is likely to be infected during the crisis. We consider two different credit score measures: the MORE (multi objective rating evaluation) credit score and the VPI (Vadis predictive indicator) propensity to become bankrupt score. The MORE credit score is an assessment of the credit-worthiness of a company. It grades companies based on how well they can meet their financial commitments. The VPI score measures the likelihood of a company to declare bankruptcy in the next 18 months. The indicator is calculated for nearly 15 million companies worldwide and regularly extended. Both credit scores are qualitative. The MORE credit score ranges from AAA ('healthy' companies) to D ('risky' companies), whereas the VPI bankruptcy score ranges from 1 (least likely to declare bankruptcy) to 5 (highest risk of bankruptcy). Figure 7.8 shows boxplots of the distribution of the recovery period per company associated with the different scores. The green line indicates the mean recovery time per group, while the orange lines mark the median of each group. It can be seen in both plots that companies which have better MORE and VPI scores are on average more likely to recover quickly than companies with worse such scores. Therefore, a company's performance before the 2020 financial crisis is indicative of its financial stability and the speed of its recovery.

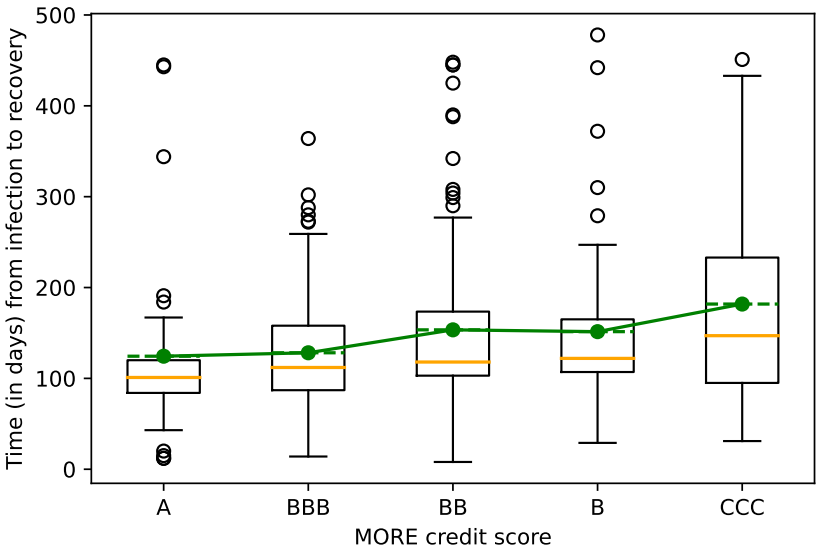

(a)

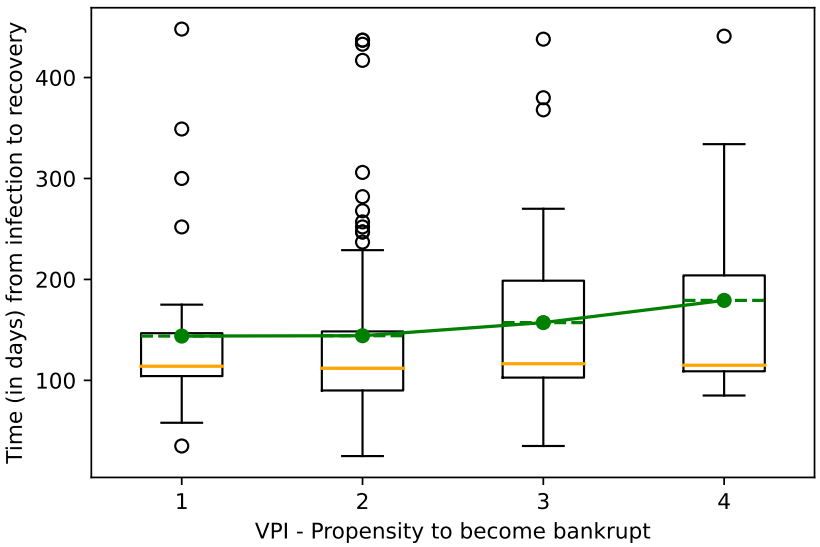

(b)

Figure 7.8: Boxplots showing the number of days from infection to recovery for companies with different (a) MORE credit score and (b) VPI- Propensity to become bankrupt score. The green line illustrates the mean, while the orange line illustrates the median.

7.4 Counterfactual scenarios: Factors influencing crisis spread

In this section we study different factors that can potentially influence crisis spread. In particular, in Section 7.4.1 we study how node centrality can influence crisis propagation. We then, in Section 7.4.2, explore how the number of initially infected companies impacts the spread of a crisis. Finally, Section 7.4.3 presents a simulation study exploring how the timing of lockdowns during the 2020 financial crisis could affect its spread. The results indicate that both company's centrality and the number of initially infected companies have a notable impact on the epidemic spread. In addition, implementing timely lockdowns is crucial for reducing infection spread and enabling faster recovery.

7.4.1 Impact of node centrality on crisis spread

In this section we present a simulation study exploring how starting a financial crisis from the most (least) influential nodes influences its spread. Specifically, the crisis begins with the same number of initially infected companies as originally (five in the 2008 financial crisis and six in the 2020 financial crisis), but these are chosen to be the most (least) influential ones based on the three centrality measures introduced in Section 7.2. For this simulation study we use the estimated daily probabilities of infection and recovery as per Chapter 3. The step-by-step procedure is described in Algorithm 7.

Based on $10,000$ SIR simulations, Figure 7.9 shows the distribution of the total number of companies that are infected at any day during the simulated crisis for the (a) 2008 crisis and (b) 2020 crisis, respectively. The orange histograms correspond to the scenario in which the simulations begin with

the least central companies initially infected, while for the blue histograms the crisis is initiated from the most central companies. The blue and orange vertical lines indicate the mean total number of infected companies, averaged across all simulations. The green solid lines show the total number of infected companies in each financial crisis that is actually observed in the data set. The figure demonstrates, as expected, that starting from the least (most) central companies results in a lower (higher) total number of infected companies in both crises. In particular, in the 2008 financial crisis, initiating the crisis with the least central nodes results in nearly 25 fewer infected companies on average compared to starting with the most central nodes. In the 2020 financial crisis this difference increases to 50. This finding suggests that the centrality of the initially infected companies plays a crucial role in the spread of financial crises. Note that the sets of the five/six (the number of initially infected companies in the 2008/2020 financial crisis) companies with highest or lowest contagion, closeness and betweenness centrality are identical. Hence, the results are similar regardless of the centrality measure used, so we display only one set of results.

Algorithm 7 SIR simulations

1. Set $t = 1$, $N = 0$ (Start from the first crisis day and set $N = 0$ to be the number of simulations).
2. Fit the model to the period from day t to day $t+1$ (window length $n = 1$) and estimate $\hat{\beta}_i(t)$ $(1 \leq i \leq 4)$ and $\hat{p}(t)$ for day t.
3.. If $t < L$, update $t = t + 1$ and go back to Step 2. Else, go to Step 4.
4. Choose the initially infected companies.
5. Run an SIR simulation on the network from day 1 to day L (the whole crisis period) using the values of $\hat{\beta}_i(t)$ $(1 \leq i \leq 4)$ and $\hat{p}(t)$ for each day $1 \leq t \leq L$, estimated in Step 2.
6. Update $N = N + 1$. If $N < 10,000$, go back to Step 4. Else, stop.

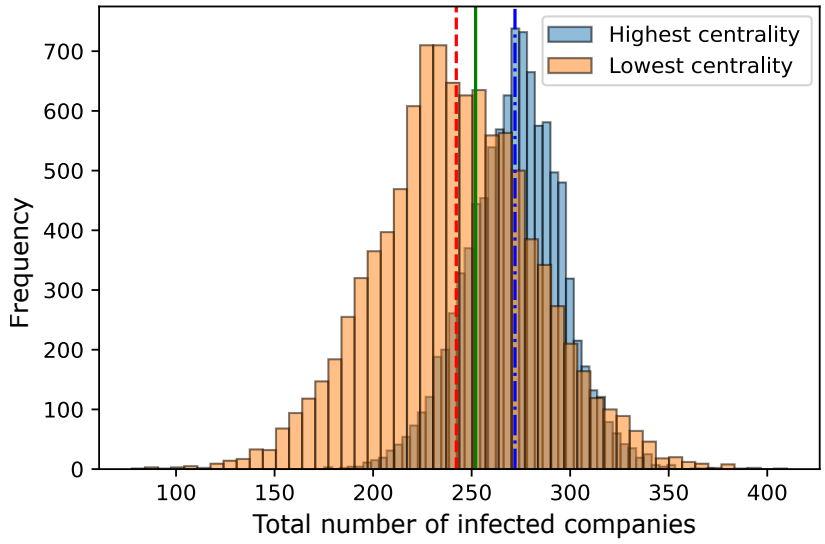

(a)

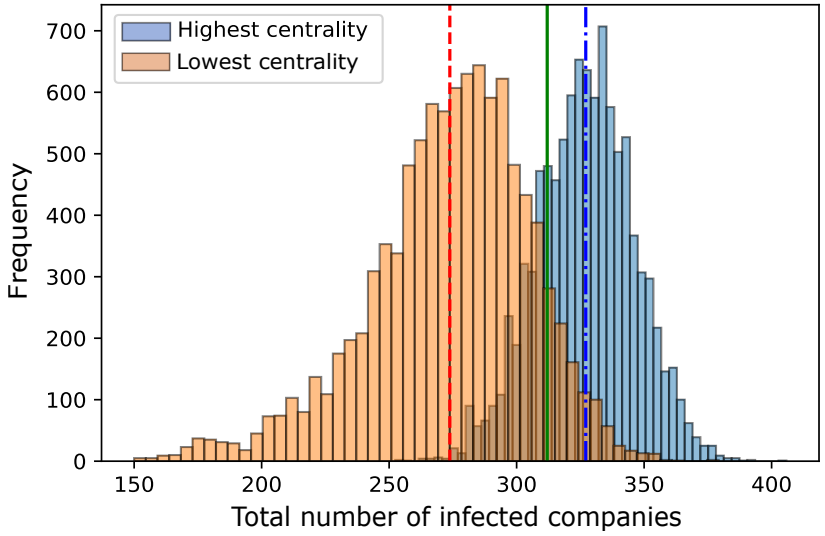

(b)

Figure 7.9: Histograms showing the distribution of the total number of infected companies out of 10,000 SIR simulations (as per Algorithm 7) in the (a) 2008 and (b) 2020 financial crises, where the simulations have started with the (a) five and (b) six most central initially infected companies (in blue) and the least central ones (in orange). The dashed vertical lines, in blue and orange, respectively, illustrate the mean averaged over all simulations. The green solid lines illustrate the actual total number of infected companies in each of the crises.

7.4.2 Impact of the number of initially infected companies on crisis spread

In this section we study how the number of initially infected companies can influence the spread of the crisis. Using Algorithm 7, we simulate an SIR epidemic, initiating each of the 2008 and 2020 financial crises with a different number of infected companies. We first choose the number of initially infected companies and then, for each SIR simulation, we choose uniformly at random a subset of all the nodes of the required size to be initially infected. Figures 7.10a and 7.10b illustrate the distribution of the total number of infected companies, over $10,000$ simulations, in the 2008 and 2020 financial crises, respectively. The dashed horizontal line indicates the actual total number of infected companies in each of the two crises, while the star's x-coordinate denotes the initial number of infected companies in each crisis. It can be seen that starting the crisis with lower number of initially infected companies leads to a higher variability in the results, including some cases where the epidemic dies out early. Also, in each of the crisis, higher number of initially infected companies leads to higher average total number of infected companies, with highest jump when comparing scenarios with one and five initial infections. This jump might indicate the so-called 'threshold effect', which entails that after a particular number of companies is infected, the crisis spreads rapidly and becomes difficult to manage. Therefore, the results indicate that understanding the thresholds that lead to rapid contagion is important for minimizing the overall impact of financial crises. Some important results on the threshold behaviour of the stochastic SIR model, outside the scope of this thesis, can be found in [119].

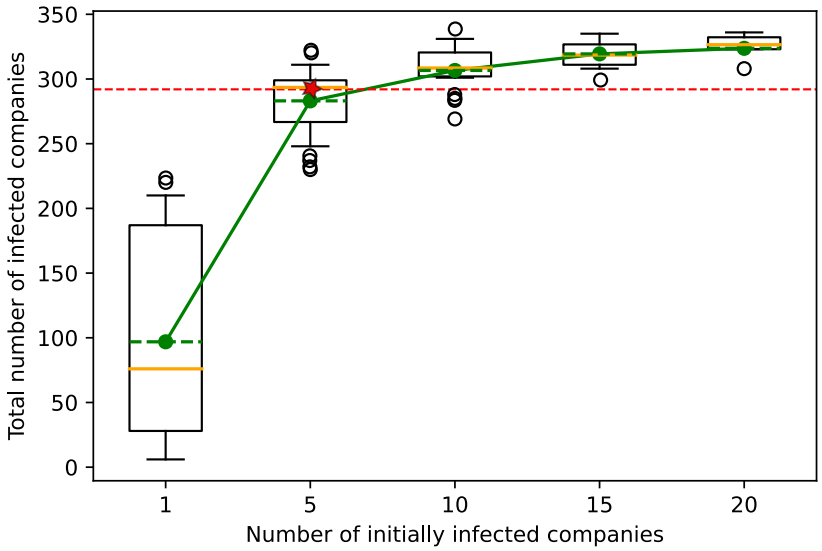

(a)

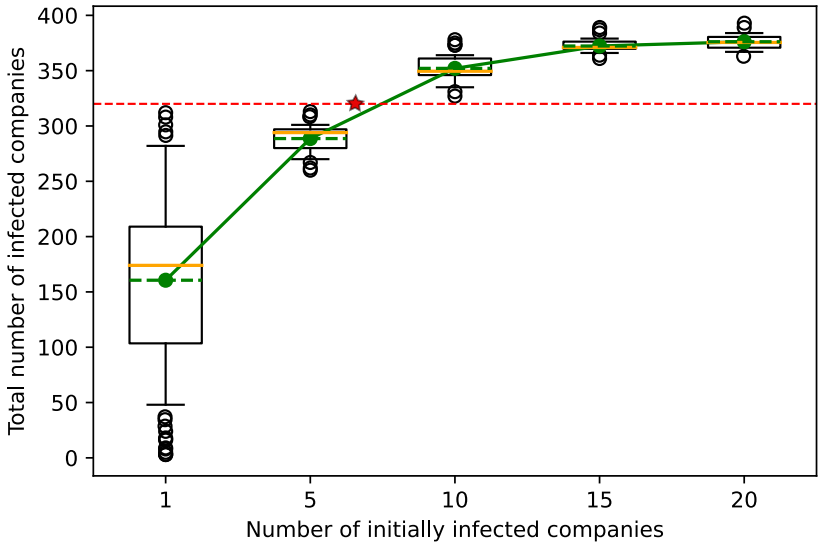

(b)

Figure 7.10: Boxplots showing the distribution of the total number of infected companies out of 10,000 SIR simulations (as per Algorithm 7) in the (a) 2008 and (b) 2020 financial crises, where the simulations have started with different number of initially infected companies. The dashed red horizontal lines indicate the total number of infected companies in each of the two crises and the star in the each plot indicates the initial number of infected companies in each crisis, i.e. 5 initially infected companies in the 2008 financial crisis and 6 in the 2020 financial crisis. The dashed green line in each separate boxplot indicates the mean, while the solid orange line indicates the median out of all 10,000 simulations.

7.4.3 Impact of lockdown timing on the 2020 financial crisis spread

In this section we simulate the possibility of a lockdown starting or ending earlier (later). Around significant events such as lockdowns the model's parameters change significantly, as illustrated in Figure 7.11, which shows the estimated parameters for each crisis day. It can be seen that when the COVID-19 was declared a global pandemic and lockdowns were imposed, the transmission probabilities $\beta_i, 1 \leq i \leq 4$, jump significantly. In addition, after governments ease lockdown restrictions the daily recovery probability p increases substantially. However, this increase is followed by a sharp drop corresponding to a surge in COVID-19 cases.

To simulate lockdown starting earlier or later, we adjust the timing of the parameter change in the model. For example, β_1 jumped from 1.3775×10^{-5} in the day before lockdown ($t = 25$) to 2.35243×10^{-1} in the day in which lockdown was imposed ($t = 26$). Then, to simulate lockdown starting one day earlier we shift this change to day 25. Figure 7.12 shows the output of 100 simulations, where the lockdown has been simulated to a) start one day earlier, b) start one day later, c) end 10 days earlier and d) end 10 days later. The results indicate that the timing of lockdown can influence the spread of infections among companies during the 2020 financial crisis. Starting the lockdown even one day earlier can result in nearly 40 more infected companies (see Figure 7.12a), while delaying it by one day leads to approximately 10 fewer infected companies (see Figure 7.12a). In addition, the timing of ending the lockdown has minimal impact on the number of infected companies but affects the speed of recovery (see Figures 7.12c and 7.12d). Consequently, implementing timely lockdowns and carefully considering their duration can be critical for minimizing the spread of infections and facilitating quicker

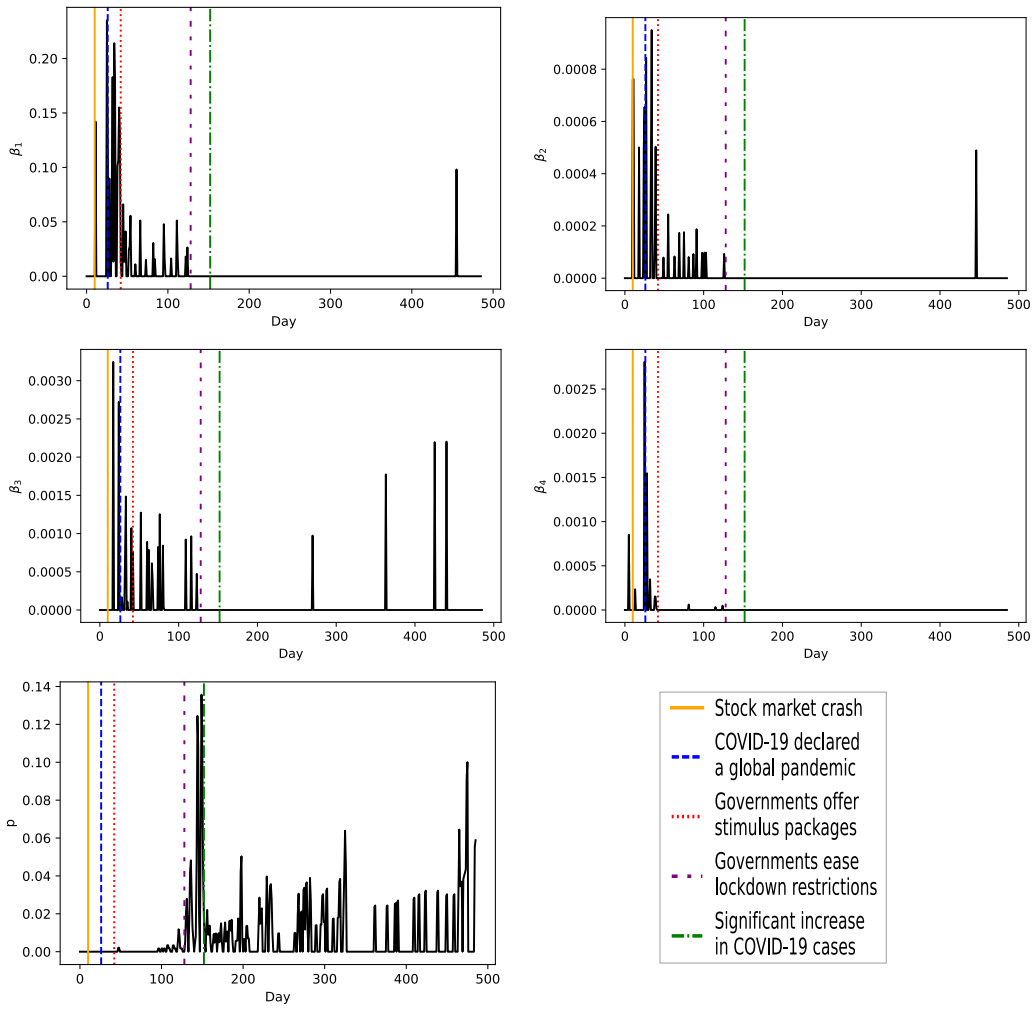

Figure 7.11: The figure illustrates the daily ($n = 1$) estimated model's parameters as per the procedure described in 3.4.2.2.1. The most significant events during the 2020 financial crisis are denoted by the coloured vertical lines.

recovery for companies. However, it is important to note that this study does not account for the human impact, such as the socioeconomic effects on employees and communities, which can be crucial for companies health. For example, if a lockdown is not imposed on time, more people might become infected by COVID-19. As a result, customer spending could decrease as people avoid shopping, dining out and traveling. In addition, the workforce might be significantly reduced due to illness, causing disruptions in production and supply chains and healthcare costs for companies could rise. However, our simulations are based on the assumption that if, for example, lockdown starts later, the transmission rates before the lockdown remain unchanged, just as if the lockdown had started earlier. However, delaying a lockdown in reality could lead to increased economic difficulties for companies as more employees become sick and fewer customers are available to purchase goods and services.

7.5 Discussion and conclusion

In this chapter we explored the relationship between node centrality measures and the probability of infection within the 2008 and 2020 financial networks. The work illustrated that contagion and closeness centrality are positively correlated with the probability of infection, indicating that companies positioned closer to others or with higher potential to spread infection are more likely to become infected themselves. However, betweenness centrality does not exhibit a strong correlation with infection probability, suggesting that being a bridge between other nodes does not necessarily increase a company's likelihood of becoming infected in a crisis. These findings can be potentially valuable for identifying highly vulnerable companies and taking proactive measures to mitigate the spread of financial contagion. For example, companies with high

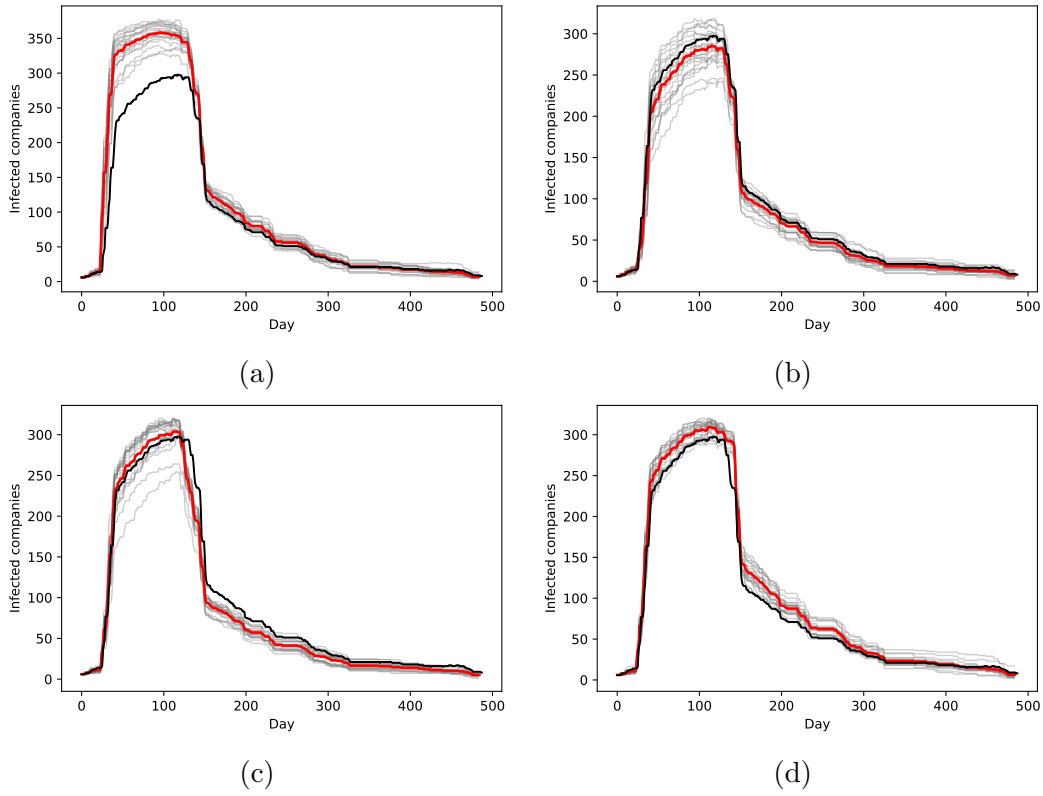

Figure 7.12: Each grey curve illustrate a single SIR simulation output out of 100 SIR simulations, while the red curve shows the average result from all simulations. In each of the plots lockdown is simulated to start a) one day earlier and b) one day later, and to end c) 10 days earlier and d) 10 days later. The black lines show the observed number of infected companies as determined in 3.4.1, providing a reference for comparison with the model predictions.

contagion or closeness centrality might be prioritised for early interventions, such as increased monitoring or support.

We then studied the connections between credit scores and the duration of infection, highlighting how a company's financial health impacts its recovery. This information might be valuable in future modelling choices, as different recovery probabilities might be chosen depending on the credit score of a company. In addition, strategies to support the most vulnerable companies might be developed to enhance the resilience of financial systems.

We finally performed various simulation studies analysing multiple scenarios including different lockdown strategies, starting a crisis with the most central nodes and changing the number of initially infected companies. The results showed that starting a crisis at more central nodes can lead to a higher crisis spread, again indicating that early interventions might be preferable for companies with higher centrality. In addition, the significant jump between the total number of infected companies, when starting the crisis with five infected companies, compared with a single infected company, suggests that early interventions are important for mitigating the epidemic spread. Simulating what-if scenarios, such as those performed in this chapter, can be informative for the potential impact of different decision-makers interventions. However, simulations often require simplifying assumptions (such as our synthetic lockdown scenarios) to make the model manageable, sometimes leading to invalid and unreliable results. In addition, simulations are typically based on existing data, so they can only predict outcomes within the scope of those data and unprecedented events may not be accurately modelled. In this context, it is important to node that the simulation studies in this chapter are intended to illustrate how such models are of value, rather than to provide realistic predictions.

Chapter 8

Discussion and conclusion

8.1 Thesis summary

In this thesis, we have presented and compared different models for financial crisis propagation. We employed different methods, such as network theory, extreme value theory and time series analysis, to create and analyse various financial contagion propagation models. In this section we summarise the contributions of this thesis to the understanding of financial contagion.

Firstly, Chapter 3 used stock price, geographical and economic sector data for a set of 398 companies to construct a multiplex network. We proposed a novel framework for modelling financial contagion using a discrete-time SIR epidemic model on the network. The findings suggested that each of the layers in the multiplex network structure provides important information for enhancing prediction power and capturing the interdependencies among companies driving financial contagion dynamics.

Secondly, in Chapter 4 we used the ChangeFinder algorithm for detecting significant changes in past crisis data and incorporating them into our model

accordingly. We achieved this by continuously monitoring the data for anomalies, and as soon as a change point was identified, we discarded the prior data and used only data subsequent to the change point to inform model predictions. The results indicated that incorporating change points into the model improved prediction accuracy, compared to the model from Chapter 3.

Chapter 5 further enhanced the model's predictive accuracy by applying weighting schemes to the models introduced in Chapters 3 and 4. In particular, in Section 5.2 we used the geometric and hyperbolic weighting schemes to weight the observed data, where more recent data has more relative importance compared to older data. The resulting model was applied to the 2020 financial crisis, with results showing that introducing both weighting schemes improves the model's accuracy, compared to the original model of Chapter 3. Then, in Section 5.3 we combined the change point model, introduced in Chapter 4 with the two weighting schemes, leading to an improved predictive accuracy.

In Chapter 6 we compared all proposed models in this thesis, including the original model (Chapter 3), the change point model (Chapter 4), the time weighted maximum likelihood (TiWML) model (Section 5.2) and the combination of the change point and the time weighted maximum likelihood (CPD+TiWML) model (Section 5.3). Each model was assessed for its performance across various prediction scenarios, such as forecasting the number of infected companies, the continents and economic sectors that will be most affected, and the specific companies most likely to be infected during future crisis days. We identified the most effective model for each type of prediction, offering guidance on which model to use based on the specific forecasting needs.

Finally, Chapter 7 studied the impact on different network characteristics, such as node centrality, on financial crisis spread. The results indicated that

node centrality is positively correlated with the likelihood of infection. We also presented different counterfactual simulation scenarios, including starting a crisis with the most central nodes or altering the number of initially infected companies. The findings suggested that the initial conditions of the crisis play crucial role for its spread.

Overall, the work presented in this thesis offers valuable insights into contagion modelling. While the main body of the thesis focuses on detailed analysis and specific findings, the models presented could also be useful in real-life applications. The models in the thesis are based on the data of only 398 companies, offering an informative snapshot but not fully capturing the diversity of the real-world network of companies. For these models to be applied effectively, they should be tested on larger and more diverse datasets. Hence, we recommend conducting additional studies that include thousands of companies to validate and refine the models. Although incorporating a significantly larger number of nodes into the network could lead to a higher computational burden, recent advances in statistical procedures for modelling processes on large networks have made it feasible to manage large networks effectively. For example, graph reduction techniques [108] and clustering-based approaches [63] have been developed to handle the challenges of scaling up networks.

Additionally, the thesis highlights an important insight: no single model consistently outperforms others across all prediction scenarios. Hence, a real-life applications would require establishing criteria for when to use different models. Alternatively, implementing a system that automatically selects or combines models based on the task at hand can optimize performance across different situations.

We finally demonstrated the usefulness of such models in counterfactual scenarios. In real-life scenarios such studies can be used for evaluating different

strategic decisions. For example, companies can use these models in supply chain management, simulating how changes to suppliers or disruptions in the supply chain might have played out. In addition, companies might simulate different possible strategies, such as merging or breaking connections with other businesses or putting new regulations into place. However, the models presented in the thesis are currently too simple to be applied in a real-world scenarios. Hence, we recommend adding more variables to the model, such as external economic factors, industry-specific information, supply chain dynamics and market trends.

8.2 Research gaps

While there has been considerable progress in modelling financial contagion on networks, there are still several research gaps that need to be addressed.

1. **Lack of standardised methodology**: Despite the rapid advancements in complex network theory in recent years, there remains no standardised methodology for modelling stochastic processes, including financial contagion, on networks. Different studies use different approaches for building network structures, as well as different contagion measures, making it difficult to compare results and draw meaningful conclusions. For example, different definitions of contagion are found in the financial literature [122]. While [11] defined contagion to be a simple transmission of shocks between countries, Bekaert et al. [33] defined it as unexplained increases in residual correlations. In addition, many papers have proposed quantitative measures of contagion [33, 77, 121]. The differing methodologies and focus areas highlight the challenges in establishing a unified framework. The lack of standardisation in constructing

network models, as well as in defining and measuring contagion leads to significant variability in findings, making it difficult to compare studies directly and to draw universally applicable insights.

2. **Lack of data**: One of the biggest challenges in modelling financial contagion on networks is the lack of comprehensive data on financial institutions and their interconnections. Most of the existing academic research on financial contagion relies on low frequency data such as monthly, quarterly or yearly data instead of high frequency, such as minute-by-minute or intra-day data [122]. Nevertheless, high-frequency data is essential for understanding how financial shocks propagate. For instance, using high-frequency data allows for the early detection of financial shocks and contagion events, as well as identification of the rapid spread of financial contagion and capture of short-term interactions and reactions within the market. However, such data is often difficult to obtain. While there are several data sources available, such as https://finance.yahoo.com/ and https://www.kaggle.com/, they are often incomplete or inaccurate. In addition, most of the financial data available pertains to financial markets of developed countries rather than emerging or developing countries [122]. Hence, the majority of existing studies have focused on modelling financial contagion among developed countries. To improve our understanding, future research should incorporate data from emerging and developing markets, leading to more comprehensive and globally applicable financial models.

3. **Empirical validation**: While there are numerous studies on financial contagion on networks, there is a need for more empirical validation. For example, many theoretical models suggest how shocks in one part of the financial network can propagate and impact other parts, but there is a lack of real-world data to confirm these mechanisms. In addition,

many of the mentioned studies in this literature review study counterfactual scenarios, where the authors simulate different scenarios on the network and analyse the results. However, usually these scenarios are not empirically validated.

4. **The 'human factor'**: The 'human factor' in financial network studies refers to the behaviours, decisions and psychological biases of individuals and institutions participating in financial markets. Humans are prone to various biases, such as overconfidence, herd behaviour and loss aversion, which can lead to irrational market movements and contagion effects that are difficult to predict with traditional models. For example, each company is directed by a group of people, who are responsible to make decisions related to the company. Individual and institutional decision-making processes are often influenced by factors such as experience, risk tolerance and expectations about future market movements. These subjective elements can introduce difficulties in predicting future market movements. In addition, political movements and public figures influence financial markets and institutions. For example, on August 7, 2018, Elon Musk tweeted,'Am considering taking Tesla private at $420. Funding secured.', which caused Tesla's stock price to spike initially due to the potential buyout at a significant premium. However, the stock price later experienced volatility as doubts emerged about the feasibility of the plan and regulatory scrutiny followed. Another popular example involving Cristiano Ronaldo occurred in July 2021 when he removed two bottles of Coca-Cola from his sight during a press conference for the Euro 2020 football tournament. Following this action, Coca-Cola's stock price declined significantly, reportedly losing billions of dollars in market value. Overall, 'the human factor' adds a layer of complexity to financial network modelling, making it challenging to develop models

that can accurately predict future market behavior. Incorporating behavioral economics can help address these challenges and improve the predictive power of financial contagion studies.

8.3 Future work

The modelling framework from Chapter 3 provides several directions that might form a basis for future research. The first natural extension to the model is considering specific transmission and recovery rates for each continent and sector. Even though this might lead to more accurate results, it is important to recognise that it will inevitably lead to higher computational complexity and may result in overparameterization. Secondly, as already mentioned in Chapter 3, the framework is universal in nature and hence could be applied to analyse any contagious financial crisis, not only in a network of companies but also in a network of countries and economic sectors. Lastly, as previously mentioned, our model operates under several basic assumptions, such as the division into susceptible, infected and recovered companies. Possible extension will be adding more complexity to the model, such as introducing additional stages of infection or recovery. In addition, the underlying network structure used throughout this thesis is considered to be static. Future work might involve using a dynamic network structure. For example, in real-life situations companies might break or develop financial ties, such as loans, investments or supply-chain contracts. Moreover, companies might merge, resulting in the integration of two or mode nodes into a single one, or new companies might emerge, leading to additional nodes. Therefore, developing a dynamic network structure could better capture the evolving nature of interactions within financial systems [159].

The change point detection model, introduced in Chapter 4 also provides a natural direction for future research. In this thesis the change point detection was performed on the time series of infections. However, an alternative choice might be performing change point detection on the time series of the estimated model's parameters. In this approach, detecting a change point would indicate a significant shift in a specific model parameter. This method would also give the opportunity to incorporate different change points for the different layers, as opposed to the current approach where detecting a change point in the infection time series leads to reestimating all parameters.

Chapter 5 provides a natural extension for combining the change point detection model and the TiWML approach. While the TiWML approach is a popular tool in financial time series forecasting, the existence of outlier observations (which are frequent in many scientific areas) may provide a valuable information about underlying market dynamics. In financial time series outliers sometimes represent significant events or market shifts, and can offer crucial insights that may improve the performance of forecasting models. In other cases, the outliers are just noisy data points, which can disproportionately influence the results of statistical models. Hence, instead of using the hyperbolic or geometric weighting schemes, one could use a weighting based on the outlier scores, introduced in Chapter 4. In particular, one could decide to give lower weight to outliers, hence minimizing their influence, or give higher weight when outliers are believed to provide valuable information about significant changes or events.

Chapter 6 demonstrates that there is no single model that excels in all scenarios, whether predicting the number of infected companies or making company-specific predictions. Hence, each prediction task may require a distinct model to achieve accurate forecasts. For example, models that perform well in predicting the numbers of companies are not necessarily effective at predicting

individual company outcomes and vice versa. Therefore, it is important to select models that strike a balance between various performance metrics. Given these challenges, a straightforward direction for future research will be developing adaptive models which can adjust based on the specific nature of the task, achieving a trade-off between different prediction objectives. In addition, future evaluation of model predictions could be enhanced by incorporating additional (probabilistic) distance measures, commonly referred to as scoring rules. The scoring rules provide measures for the evaluation of probabilistic forecasts by assigning a numerical score, based on the event being predicted [103]. Unlike point-based measures such as the Accuracy and F_1-score used in this thesis, scoring rules assess the full probabilistic forecast, providing insights on both model's predictive accuracy and uncertainty. Furthermore, as crises evolve, the optimal model may change, requiring dynamic adaptation in the model selection process. Identifying how to efficiently choose an optimal model in response to changing conditions presents a significant challenge. Hence, future research should involve developing adaptive models that can be adjusted in real-time.

Finally, Chapter 7 provides examples of how the model could be used to study different counterfactual scenarios. An important area of further work is using such simulations to study possible strategies for minimizing crisis spread. In addition, as discussed in this chapter, some credit scores are correlated with the duration a company stays infected during a crisis. This provides a natural extension of the model, where recovery probabilities could be dependent on a company's credit score. Similarly, subject to data availability, additional company-specific attributes can be incorporated, such as size or operating revenue, to model how different characteristics affect susceptibility or recovery rates.

Bibliography

[1] K. Abduraimova. Contagion and tail risk in complex financial networks. *Journal of Banking & Finance*, page 106560, 2022.

[2] D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.

[3] A. AghaKouchak, S. Sellars, and S. Sorooshian. Methods of tail dependence estimation. *Extremes in a changing climate*, pages 163–179, 2013.

[4] S. Ahmed and O. Farooq. The effect of 9/11 on the stock market volatility dynamics: Empirical evidence from a front line state. *International Research Journal of Finance and Economics*, (16):71–83, 2008.

[5] T. Ahnert and C.-P. Georg. Information contagion and systemic risk. *Journal of Financial Stability*, 35:159–171, 2018.

[6] A. Akhmanova and M. O. Steinmetz. Tracking the ends: a dynamic protein network controls the fate of microtubule tips. *Nature reviews Molecular cell biology*, 9(4):309–322, 2008.

[7] R. Albert. Network inference, analysis, and modeling in systems biology. *The Plant Cell*, 19(11):3327–3338, 2007.

[8] R. Albert, H. Jeong, and A.-L. Barabási. Diameter of the world-wide web. *Nature*, 401(6749):130–131, 1999.

[9] C. Alexander. *Market risk analysis, pricing, hedging and trading financial instruments*. John Wiley & Sons, 2008.

[10] F. Allen and A. Babus. Networks in finance. *The network challenge: strategy, profit, and risk in an interlinked world*, 367, 2009.

[11] F. Allen and D. Gale. Financial contagion. *Journal of political economy*, 108(1):1–33, 2000.

[12] S. Aminikhanghahi and D. J. Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017.

[13] K. Anand, P. Gai, S. Kapadia, S. Brennan, and M. Willison. A network model of financial system resilience. *Journal of Economic Behavior & Organization*, 85:219–235, 2013.

[14] I. L. Arancibia-Cárcamo, M. C. Ford, L. Cossell, K. Ishida, K. Tohyama, and D. Attwell. Node of Ranvier length as a potential regulator of myelinated axon conduction speed. *Elife*, 6:e23329, 2017.

[15] A. Arnon, J. Ricco, and K. Smetters. Epidemiological and economic effects of lockdown. *Brookings Papers on Economic Activity*, 2020(3):61–108, 2020.

[16] B. E. Aslanertik, S. Erdem, and G. K. Gümüş. Extreme value theory in finance: A way to forecast unexpected circumstances. In *Risk Management, Strategic Thinking and Leadership in the Financial Services Industry*, pages 177–190. Springer, 2017.

[17] S. A. Asongu. The 2011 Japanese earthquake, tsunami and nuclear crisis: evidence of contagion from international financial markets. *Journal of Financial Economic Policy*, 4(4):340–353, 2012.

[18] T. Aste, T. Di Matteo, and S. Hyde. Complex networks on hyperbolic surfaces. *Physica A: Statistical Mechanics and its Applications*, 346(1-2):20–26, 2005.

[19] A. G. Awan. Analysis of the impact of 2008 financial crisis on the economic, political and health systems and societies of advanced countries. *Global Journal of Management and Social Sciences*, 1(1):1–16, 2015.

[20] S. Bansal, J. Read, B. Pourbohloul, and L. A. Meyers. The dynamic nature of contact networks in infectious disease epidemiology. *Journal of biological dynamics*, 4(5):478–489, 2010.

[21] A.-L. Barabási. Network science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1987):20120375, 2013.

[22] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[23] M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli. Pathways towards instability in financial networks. *Nature communications*, 8(1):14416, 2017.

[24] L. Bargigli, G. Di Iasio, L. Infante, F. Lillo, and F. Pierobon. The multiplex structure of interbank networks. *Quantitative Finance*, 15(4):673–691, 2015.

[25] A. Barja, A. Martínez, A. Arenas, P. Fleurquin, J. Nin, J. J. Ramasco, and E. Tomás. Assessing the risk of default propagation in interconnected sectoral financial networks. *EPJ Data Science*, 8(1):32, 2019.

[26] R. Barrell and E. P. Davis. The evolution of the financial crisis of 2007—8. *National Institute economic review*, 206:5–14, 2008.

[27] P. Bartesaghi, M. Benzi, G. P. Clemente, R. Grassi, and E. Estrada. Risk-dependent centrality in economic and financial networks. *SIAM Journal on Financial Mathematics*, 11(2):526–565, 2020.

[28] D. S. Bassett and E. Bullmore. Small-world brain networks. *The neuroscientist*, 12(6):512–523, 2006.

[29] D. S. Bassett and O. Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353–364, 2017.

[30] F. Battiston, V. Nicosia, and V. Latora. Structural measures for multiplex networks. *Physical Review E*, 89(3):032804, 2014.

[31] S. Battiston, G. Caldarelli, and A. Garas. *Multiplex and multilevel networks*. Oxford University Press, 2018.

[32] J. Beirlant, Y. Goegebeur, J. Segers, and J. L. Teugels. *Statistics of extremes: theory and applications*. John Wiley & Sons, 2006.

[33] G. Bekaert, M. Ehrmann, M. Fratzscher, and A. Mehl. The global crisis and equity market contagion. *The Journal of Finance*, 69(6):2597–2649, 2014.

[34] A. Belke, I. Dubova, and T. Osowski. Policy uncertainty and international financial markets: the case of Brexit. *Applied Economics*, 50(34-35):3752–3770, 2018.

[35] Y. Bensalah. Steps in applying extreme value theory to finance: A review. Technical report, Bank of Canada, 2000.

[36] T. Berge, J.-S. Lubuma, G. Moremedi, N. Morris, and R. Kondera-Shava. A simple mathematical model for Ebola in Africa. *Journal of biological dynamics*, 11(1):42–74, 2017.

[37] N. Berger, S. Schulze-Schwering, E. Long, and S. Spinler. Risk management of supply chain disruptions: An epidemic modeling approach. *European Journal of Operational Research*, 304(3):1036–1051, 2023.

[38] P. D. Z. Bermudez and Z. Mendes. Extreme value theory in medical sciences: Modeling total high cholesterol levels. *Journal of statistical theory and practice*, 6(3):468–491, 2012.

[39] B. S. Bernanke. Asia and the global financial crisis. *Conference on Asia and the Global Financial Crisis*, page 11, 2009.

[40] A. Bhan and E. Mjolsness. Static and dynamic models of biological networks. *Complexity*, 11(6):57–63, 2006.

[41] G. Bonanno, G. Caldarelli, F. Lillo, and R. N. Mantegna. Topology of correlation-based minimal spanning trees in real and model markets. *Physical Review E*, 68(4):046130, 2003.

[42] G. Bonanno, F. Lillo, and R. N. Mantegna. High-frequency cross-correlation in a set of stocks. *Quantitative Finance*, 1(1):96–104, 2001.

[43] G. Bonanno, N. Vandewalle, and R. N. Mantegna. Taxonomy of stock market indices. *Physical review E*, 62(6):R7615, 2000.

[44] S. P. Borgatti, A. Mehra, D. J. Brass, and G. Labianca. Network analysis in the social sciences. *Science*, 323(5916):892–895, 2009.

[45] M. Bozhidarova, F. Ball, Y. Van Gennip, R. D. O'Dea, and G. Stupfler. Describing financial crisis propagation through epidemic modelling on multiplex networks. *Proceedings of the Royal Society A*, 480(2287):20230787, 2024.

[46] F. Brauer. Mathematical epidemiology: Past, present, and future. *Infectious Disease Modelling*, 2(2):113–127, 2017.

[47] D. Brounen and J. Derwall. The impact of terrorist attacks on international stock markets. *European Financial Management*, 16(4):585–598, 2010.

[48] J.-J. Cai, J. H. J. Einmahl, L. de Haan, and C. Zhou. Estimation of the marginal expected shortfall: the mean when a related variable is extreme. *Journal of the Royal Statistical Society: Series B*, 77(2):417–442, 2015.

[49] S.-M. Cai, W. Chen, D.-B. Liu, M. Tang, and X. Chen. Complex network analysis of brain functional connectivity under a multi-step cognitive task. *Physica A: Statistical Mechanics and Its Applications*, 466:663–671, 2017.

[50] G. Caldarelli. *Complex networks*. EOLSS Publications, 2010.

[51] L. Cantini, E. Medico, S. Fortunato, and M. Caselle. Detection of gene communities in multi-networks reveals cancer drivers. *Scientific reports*, 5(1):1–10, 2015.

[52] D. A. Carter and B. J. Simkins. The market's reaction to unexpected, catastrophic events: the case of airline stock returns and the September 11th attacks. *The Quarterly Review of Economics and Finance*, 44(4):539–558, 2004.

[53] X. Chang and J. Li. Business performance prediction in location-based social commerce. *Expert Systems with Applications*, 126:112–123, 2019.

[54] Q. Chen, D. E. Giles, and H. Feng. The extreme-value dependence between the Chinese and other international stock markets. *Applied Financial Economics*, 22(14):1147–1160, 2012.

[55] J.-J. Cheng, Y. Liu, B. Shen, and W.-G. Yuan. An epidemic model of rumor diffusion in online social networks. *The European Physical Journal B*, 86(1):29, 2013.

[56] Y. Chernin and Y. Lahav. "The people demand social justice" a case study on the impact of protests on financial markets. *Accounting, Economics and Law*, 4(2):99–121, 2014.

[57] T. Chevapatrakul and K.-H. Tee. The effects of news events on market contagion: evidence from the 2007–2009 financial crisis. *Research in international Business and Finance*, 32:83–105, 2014.

[58] D.-T. Chu, S.-M. Vu Ngoc, H. Vu Thi, Y.-V. Nguyen Thi, T.-T. Ho, V.-T. Hoang, V. Singh, and J. A. Al-Tawfiq. COVID-19 in Southeast Asia: current status and perspectives. *Bioengineered*, 13(2):3797–3809, 2022.

[59] M. S. Claessens, M. A. Kose, M. L. Laeven, and M. F. Valencia. *Financial crises: Causes, consequences, and policy responses*. International Monetary Fund, 2014.

[60] S. Claessens and K. Forbes. *International financial contagion*. Springer Science & Business Media, 2013.

[61] S. Coles, J. Heffernan, and J. Tawn. Dependence measures for extreme value analyses. *Extremes*, 2(4):339–365, 1999.

[62] E. Cozzo, M. Kivelä, M. De Domenico, A. Solé-Ribalta, A. Arenas, S. Gómez, M. A. Porter, and Y. Moreno. Structure of triadic relations in multiplex networks. *New Journal of Physics*, 17(7):073029, 2015.

[63] H. Crane and M. Xu. Root and community inference on the latent growth process of a network. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(4):825–865, 2024.

[64] J. J. Crofts, M. Forrester, and R. D. O'Dea. Structure-function clustering in multiplex brain networks. *EPL (Europhysics Letters)*, 116(1):18003, 2016.

[65] B. Cui, L. Meilong, and J. Zhu. Review of the network risk propagation research. *Aeron Aero Open Access J*, 3(2):66–74, 2019.

[66] D. M. Cutler and L. H. Summers. The COVID-19 pandemic and the $16 trillion virus. *Jama*, 324(15):1495–1496, 2020.

[67] J. Danielsson and C. G. De Vries. Value-at-risk and extreme returns. *Annales d'Economie et de Statistique*, pages 239–270, 2000.

[68] G. F. De Arruda, A. L. Barbieri, P. M. Rodriguez, F. A. Rodrigues, Y. Moreno, and L. d. F. Costa. Role of centrality for the identification of influential spreaders in complex networks. *Physical Review E*, 90(3):032812, 2014.

[69] M. De Domenico, A. Solé-Ribalta, E. Omodei, S. Gómez, and A. Arenas. Centrality in interconnected multilayer networks. *Physica D: Nonlinear Phenomena*, 2013.

[70] R. M. del Rio-Chanona, Y. Korniyenko, M. Patnam, and M. A. Porter. The multiplex nature of global financial contagions. *Applied Network Science*, 5(1):1–23, 2020.

[71] N. Demiris, T. Kypraios, and L. V. Smith. On the epidemic of financial crises. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, pages 697–723, 2014.

[72] D. W. Diamond and P. H. Dybvig. Bank runs, deposit insurance, and liquidity. *Journal of political economy*, 91(3):401–419, 1983.

[73] F. X. Diebold, T. Schuermann, and J. D. Stroughair. Pitfalls and opportunities in the use of extreme value theory in risk management. In *Decision technologies for computational finance*, pages 3–12. Springer, 1998.

[74] E. W. Dijkstra. A note on two problems in connexion with graphs. In *Edsger Wybe Dijkstra: his life, work, and legacy*, pages 287–290. 2022.

[75] P. Domingos and M. Richardson. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 57–66, 2001.

[76] R. Dornbusch, Y. C. Park, and S. Claessens. Contagion: How it spreads and how it can be stopped. *World Bank Research Observer*, 15(2):177–197, 2000.

[77] M. Dungey, R. Fry, B. González-Hermosillo, and V. L. Martin. Empirical modelling of contagion: a review of methodologies. *Quantitative finance*, 5(1):9–24, 2005.

[78] J. H. Einmahl, A. Kiriliouk, and J. Segers. A continuous updating weighted least squares estimator of tail dependence in high dimensions. *Extremes*, 21(2):205–233, 2018.

[79] M. Elliott, B. Golub, and M. O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–3153, 2014.

[80] P. Embrechts, S. I. Resnick, and G. Samorodnitsky. Extreme value theory as a risk management tool. *North American Actuarial Journal*, 3(2):30–41, 1999.

[81] S. Engelke and A. S. Hitz. Graphical models for extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4):871–932, 2020.

[82] S. Engelke and J. Ivanovs. Sparse structures for multivariate extremes. *Annual Review of Statistics and its Application*, 8, 2020.

[83] M. Eryiğit and R. Eryiğit. Network structure of cross-correlations among the world market indices. *Physica A: Statistical Mechanics and its Applications*, 388(17):3551–3562, 2009.

[84] S. Eubank, H. Guclu, V. Anil Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184, 2004.

[85] L. Euler. Solutio problematis ad geometriam situs pertinentis. *Commentarii academiae scientiarum Petropolitanae*, pages 128–140, 1741.

[86] P. M. Fairfield, S. Ramnath, and T. L. Yohn. Does industry-level analysis improve profitability and growth forecasts? *Journal of Accounting Research*, 47(1):147–178, 2009.

[87] B. Fakhry, B. Aktan, O. Masood, M. Tvaronavičienė, and S. Celik. The impact of a recent natural disaster on the Japanese financial markets: Empirical evidence. *Journal of Competitiveness*, 10(2):56, 2018.

[88] V. Fanelli and L. Maddalena. A nonlinear dynamic model for credit risk contagion. *Mathematics and Computers in Simulation*, 174:45–58, 2020.

[89] C. Ferguson. *Inside job: The financiers who pulled off the heist of the century*. Simon and Schuster, 2012.

[90] P. Fiedor. Networks in financial markets based on the mutual information rate. *Physical Review E*, 89(5):052801, 2014.

[91] E. C. Gabrick, E. Sayari, D. L. Souza, F. S. Borges, J. Trobia, E. K. Lenzi, and A. M. Batista. Fractal and fractional SIS model for syphilis data. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(9), 2023.

[92] P. Gai and S. Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2120):2401–2423, 2010.

[93] A. Garas, P. Argyrakis, C. Rozenblat, M. Tomassini, and S. Havlin. Worldwide spreading of economic crisis. *New journal of Physics*, 12(11):113043, 2010.

[94] D. Gembris, J. G. Taylor, and D. Suter. Trends and random fluctuations in athletics. *Nature*, 417(6888):506–506, 2002.

[95] S. Gerlach and F. Smets. Contagious speculative attacks. *European Journal of Political Economy*, 11(1):45–63, 1995.

[96] I. Gigauri et al. Influence of COVID-19 crisis on human resource management and companies' response: the expert study. *International Journal of Management Science and Business Administration*, 6(6):15–24, 2020.

[97] E. N. Gilbert. Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.

[98] F. Gilbert, P. Simonetto, F. Zaidi, F. Jourdan, and R. Bourqui. Communities and hierarchical structures in dynamic social networks: analysis and visualization. *Social Network Analysis and Mining*, 1(2):83–95, 2011.

[99] M. Gilli et al. An application of extreme value theory for measuring financial risk. *Computational Economics*, 27(2):207–228, 2006.

[100] N. Gissibl and C. Klüppelberg. Max-linear models on directed acyclic graphs. *Bernoulli*, 24(4A):2693–2720, 2018.

[101] P. Giudici and G. Polinesi. Crypto price discovery through correlation networks. *Annals of Operations Research*, 299(1):443–457, 2021.

[102] R. Glick and A. K. Rose. Contagion and trade: why are currency crises regional? *Journal of international Money and Finance*, 18(4):603–617, 1999.

[103] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

[104] M. S. Granovetter. The strength of weak ties. *American journal of sociology*, 78(6):1360–1380, 1973.

[105] J. Hallikas, V.-M. Virolainen, and M. Tuominen. Understanding risk and uncertainty in supplier networks–a transaction cost approach. *International Journal of Production Research*, 40(15):3519–3531, 2002.

[106] Z. Hammoud and F. Kramer. Multilayer networks: aspects, implementations, and application in biomedicine. *Big Data Analytics*, 5(1):1–18, 2020.

[107] H. Hasan, N. Salam, and M. B. Adam. Modelling extreme temperature in Malaysia using generalized extreme value distribution. *International Journal of Mathematical and Computational Sciences*, 7(6):983–989, 2013.

[108] M. Hashemi, S. Gong, J. Ni, W. Fan, B. A. Prakash, and W. Jin. A comprehensive survey on graph reduction: Sparsification, coarsening, and condensation. *arXiv preprint arXiv:2402.03358*, 2024.

[109] P. Hiram Guzzi, F. Petrizzelli, and T. Mazza. Disease spreading modeling and analysis: A survey. *Briefings in Bioinformatics*, 23(4):bbac230, 2022.

[110] M. Hmimida and R. Kanawati. Community detection in multiplex net-

works: A seed-centric approach. *Networks & Heterogeneous Media*, 10(1):71, 2015.

[111] D. Hodson and L. Quaglia. European perspectives on the global financial crisis: Introduction. *JCMS: Journal of Common Market Studies*, 47(5):939–953, 2009.

[112] M. Hohlmeier and C. Fahrholz. The impact of Brexit on financial markets—taking stock. *International Journal of Financial Studies*, 6(3):65, 2018.

[113] C. Huaihu and Z. Jianming. Research on banking crisis contagion dynamics based on the complex network of system engineering. *Systems Engineering Procedia*, 5:156–161, 2012.

[114] J.-K. Hwang. Spillover effects of the 2008 financial crisis in Latin America stock markets. *International Advances in Economic Research*, 20:311–324, 2014.

[115] K. Imakubo, Y. Soejima, et al. The transaction network in Japan's interbank money markets. *Monetary and Economic Studies*, 28:107–150, 2010.

[116] T. Iwata, K. Nakamura, Y. Tokusashi, and H. Matsutani. Accelerating online change-point detection algorithm using 10 GbE FPGA NIC. In *Euro-Par 2018: Parallel Processing Workshops: Euro-Par 2018 International Workshops, Turin, Italy, August 27-28, 2018, Revised Selected Papers 24*, pages 506–517. Springer, 2019.

[117] M. O. Jackson and A. Pernoud. Distorted investment incentives, regulation, and equilibrium multiplicity in a model of financial networks. *Regulation, and Equilibrium Multiplicity in a Model of Financial Networks (March 1, 2019)*, 2019.

[118] K. A. Jacobsen, M. G. Burch, J. H. Tien, and G. A. Rempała. The large graph limit of a stochastic epidemic model on a dynamic multilayer network. *Journal of biological dynamics*, 12(1):746–788, 2018.

[119] C. Ji and D. Jiang. Threshold behaviour of a stochastic SIR model. *Applied Mathematical Modelling*, 38(21-22):5067–5079, 2014.

[120] J. Joe and M. Noble. A Triangle-free, 4-chromatic Q3 Euclidean distance graph scavenger hunt! *Discrete Applied Mathematics*, 339:216–226, 2023.

[121] G. A. Karolyi. Does international financial contagion really exist? *International Finance*, 6(2):179–199, 2003.

[122] H. Kaur. Insights into financial contagion: A bibliometric study. *FIIB Business Review*, 2023.

[123] S. Kawashima and F. Takeda. The effect of the Fukushima nuclear accident on stock prices of electric power utilities in Japan. *Energy Economics*, 34(6):2029–2038, 2012.

[124] M. J. Keeling and K. T. Eames. Networks and epidemic models. *Journal of the royal society interface*, 2(4):295–307, 2005.

[125] D. Y. Kenett and S. Havlin. Network science: a useful tool in economics and finance. *Mind & Society*, 14(2):155–167, 2015.

[126] W. O. Kermack and A. G. McKendrick. Contributions to the mathematical theory of epidemics. II. The problem of endemicity. *Proceedings of the Royal Society of London. Series A, containing papers of a mathematical and physical character*, 138(834):55–83, 1932.

[127] I. Z. Kiss, L. Berthouze, T. J. Taylor, and P. L. Simon. Modelling approaches for simple dynamic networks and applications to disease

transmission models. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 468(2141):1332–1355, 2012.

[128] I. Z. Kiss, J. C. Miller, P. L. Simon, et al. *Mathematics of epidemics on networks*, volume 598. Springer, 2017.

[129] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter. Multilayer networks. *Journal of complex networks*, 2(3):203–271, 2014.

[130] P. R. Kleindorfer, Y. J. R. Wind, and R. E. Gunther. *The network challenge: strategy, profit, and risk in an interlinked world*. Prentice Hall Professional, 2009.

[131] E. D. Kolaczyk and G. Csárdi. *Statistical analysis of network data with R*, volume 65. Springer, 2014.

[132] O. Kostylenko and O. Nazarenko. Modelling and identification of macroeconomic system dynamic interindustry balance. *Mechanism of Economic Regulation*, 1(63):76–86, 2014.

[133] G. B. Koumou. Diversification and portfolio theory: a review. *Financial Markets and Portfolio Management*, 34(3):267–312, 2020.

[134] W. Kristjanpoller, A. Fadic, and M. C. Minutolo. Volatility forecast using hybrid neural network models. *Expert Systems with Applications*, 41(5):2437–2442, 2014.

[135] F. Kyriakopoulos, S. Thurner, C. Puhr, and S. W. Schmitz. Network and eigenvalue analysis of financial transaction networks. *The European Physical Journal B*, 71:523–531, 2009.

[136] P. R. Lane and G. M. Milesi-Ferretti. The cross-country incidence of the global crisis. *IMF Economic Review*, 59(1):77–110, 2011.

[137] T. H. Le, H. X. Do, D. K. Nguyen, and A. Sensoy. Covid-19 pandemic and tail-dependency networks of financial assets. *Finance research letters*, 38:101800, 2021.

[138] A. W. Ledford and J. A. Tawn. Modelling dependence within joint tail regions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2):475–499, 1997.

[139] A. W. Ledford and J. A. Tawn. Diagnostics for dependence within time series extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):521–543, 2003.

[140] M. Li, X. Wang, K. Gao, and S. Zhang. A survey on information diffusion in online social networks: Models and methods. *Information*, 8(4):118, 2017.

[141] L. Liu, Z. Shao, J. Lv, F. Xu, S. Ren, Q. Jin, J. Yang, W. Ma, H. Xie, D. Zhang, et al. Identification of early warning signals at the critical transition point of colorectal cancer based on dynamic network analysis. *Frontiers in bioengineering and biotechnology*, 8:530, 2020.

[142] Y. Liu, Z. Li, H. Xiong, X. Gao, and J. Wu. Understanding of internal clustering validation measures. In *2010 IEEE international conference on data mining*, pages 911–916. IEEE, 2010.

[143] M. López, A. Peinado, and A. Ortiz. An extensive validation of a SIR epidemic model to study the propagation of jamming attacks against iot wireless networks. *Computer Networks*, 165:106945, 2019.

[144] J. D. Loyal and Y. Chen. Statistical network analysis: A review with applications to the coronavirus disease 2019 pandemic. *International Statistical Review*, 88(2):419–440, 2020.

[145] U. Y. Madaki and A. A. Daya. The spread of contagious disease (measles) using SIR model: A case study of Yobe State specialist hospital, Damaturu. *ASEANA Science and Education Journal*, 1(1):7–16, 2021.

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

[147] H. Marfatia. The financial market's ability to forecast economic growth: information from sectoral movements. *Journal of Economic Studies*, 50(7):1467–1484, 2023.

[148] G. P. Massara, T. Di Matteo, and T. Aste. Network filtering for big data: Triangulated maximally filtered graph. *Journal of complex Networks*, 5(2):161–178, 2016.

[149] R. M. May, S. A. Levin, and G. Sugihara. Ecology for bankers. *Nature*, 451(7181):893–895, 2008.

[150] S. Mittnik and M. S. Paolella. Conditional density and value-at-risk prediction of Asian currency exchange rates. *Journal of Forecasting*, 19(4):313–333, 2000.

[151] J. Moody, D. McFarland, and S. Bender-deMoll. Dynamic network visualization. *American journal of sociology*, 110(4):1206–1241, 2005.

[152] J. L. Moreno. *Who shall survive?: A new approach to the problem of human interrelations.* Nervous and Mental Disease Publishing Co., 1934.

[153] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *Science*, 328(5980):876–878, 2010.

[154] M. C. Münnix, T. Shimada, R. Schäfer, F. Leyvraz, T. H. Seligman, T. Guhr, and H. E. Stanley. Identifying states of a financial market. *Scientific reports*, 2(1):644, 2012.

[155] A. Murphy. An analysis of the financial crisis of 2008: causes and solutions. *SSRN Electronic Journal*, 2008.

[156] L. Musango, A. Nundoochan, and J. M. Kirigia. The discounted money value of human life losses associated with COVID-19 in Mauritius. *Frontiers in public health*, page 766, 2020.

[157] N. Musmeci, T. Aste, and T. Di Matteo. Relation between financial market structure and the real economy: comparison between clustering methods. *PloS one*, 10(3):e0116201, 2015.

[158] A. Nagurney. Networks in finance. In *Handbook on Information Technology in Finance*, pages 383–419. Springer, 2008.

[159] A. Nagurney and S. Siokos. *Financial networks: Statics and dynamics*. Springer Science & Business Media, 2012.

[160] M. S. Neto, M. J. do Carmo, E. M. S. Ribeiro, and W. V. G. da Cruz. Corporate ownership network in the automobile industry: Owners, shareholders and passive investment funds. *Research in Globalization*, 2:100016, 2020.

[161] T. M. Newcomb. *The acquaintance process as a prototype of human interaction.* Holt, Rinehart & Winston, 1961.

[162] M. E. Newman. Models of the small world. *Journal of Statistical Physics*, 101(3):819–841, 2000.

[163] M. E. Newman. Mixing patterns in networks. *Physical review E*, 67(2):026126, 2003.

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

[165] M. E. Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006.

[166] T. W. Ng, G. Turinici, and A. Danchin. A double epidemic model for the SARS propagation. *BMC Infectious Diseases*, 3:1–16, 2003.

[167] J. Nikkinen, M. M. Omran, P. Sahlström, and J. Äijö. Stock returns and volatility following the September 11 attacks: Evidence from 53 equity markets. *International Review of Financial Analysis*, 17(1):27–46, 2008.

[168] T. Nishizeki and N. Chiba. *Planar graphs: Theory and algorithms*. Elsevier, 1988.

[169] M. Norouzi, D. J. Fleet, and R. R. Salakhutdinov. Hamming distance metric learning. *Advances in neural information processing systems*, 25, 2012.

[170] D. Osthus, K. S. Hickmann, P. C. Caragea, D. Higdon, and S. Y. Del Valle. Forecasting seasonal influenza with a state-space SIR model. *The annals of applied statistics*, 11(1):202, 2017.

[171] A. Pak, O. A. Adegboye, A. I. Adekunle, K. M. Rahman, E. S. McBryde, and D. P. Eisen. Economic consequences of the COVID-19 outbreak: the need for epidemic preparedness. *Frontiers in public health*, 8:241, 2020.

[172] G. Palla, I. Derényi, I. Farkas, and T. Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.

[173] B. Ø. Palsson. *Systems biology: simulation of dynamic network states*. Cambridge University Press, 2011.

[174] M. S. Paolella and S. C. Steude. Risk prediction: a DWARF-like approach. *Journal of Risk Model Validation*, 2(1):25–44, 2008.

[175] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001.

[176] K. Pearson. VII. Note on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58(347-352):240–242, 1895.

[177] L. Peng. Estimation of the coefficient of tail dependence in bivariate extremes. *Statistics & Probability Letters*, 43(4):399–409, 1999.

[178] G. Peralta and R. Crisóstomo. Financial contagion with spillover effects: a multiplex network approach. ESRB Working Paper Series 32, European Systemic Risk Board, 2016.

[179] S. Pilosof, M. A. Porter, M. Pascual, and S. Kéfi. The multilayer nature of ecological networks. *Nature Ecology & Evolution*, 1(4):1–9, 2017.

[180] E. R. Pinto, E. G. Nepomuceno, and A. S. Campanharo. Impact of network topology on the spread of infectious diseases. *TEMA (São Carlos)*, 21:95–115, 2020.

[181] T. Plümper and E. Neumayer. Lockdown policies and the dynamics of the first wave of the Sars-Cov-2 pandemic in Europe. *Journal of European Public Policy*, 29(3):321–341, 2022.

[182] S. Poledna, J. L. Molina-Borboa, S. Martínez-Jaramillo, M. Van Der Leij, and S. Thurner. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *Journal of Financial Stability*, 20:70–81, 2015.

[183] S.-H. Poon, M. Rockinger, and J. Tawn. Modelling extreme-value dependence in international stock markets. *Statistica Sinica*, pages 929–953, 2003.

[184] M. Pósfai and A.-L. Barabási. *Network science*. Citeseer, 2016.

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

[186] L. Qiu, C. Gu, Q. Xiao, H. Yang, and G. Wu. State network approach to characteristics of financial crises. *Physica A: Statistical Mechanics and its Applications*, 492:1120–1128, 2018.

[187] F. Quesnay. Tableau économique avec son explication, ou Extrait des économies royales de Sully. *Paris 1758*, 1758.

[188] A. Ramos and A. Ledford. Regular score tests of independence in multivariate extreme values. *Extremes*, 8(1):5–26, 2005.

[189] Y. S. Rao, A. K. Rauta, H. Saini, and T. C. Panda. Mathematical model for cyber attack in computer network. *International Journal of Business Data Communications and Networking (IJBDCN)*, 13(1):58–65, 2017.

[190] K. Rasul, C. Seward, I. Schuster, and R. Vollgraf. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*, pages 8857–8868. PMLR, 2021.

[191] S. Rezaei and J. Behnamian. Competitive planning of partnership supply networks focusing on sustainable multi-agent transportation and virtual alliance: A matheuristic approach. *Journal of Cleaner Production*, 333:130073, 2022.

[192] S. J. Roberts. Extreme value statistics for novelty detection in biomedical data processing. *IEE Proceedings-Science, Measurement and Technology*, 147(6):363–367, 2000.

[193] M. Z. Salim, D. Ramdhan, and K. Daly. Centrality measures of financial system interconnectedness: A multiple crises study. *Heliyon*, 9(4), 2023.

[194] M. Salter-Townshend, A. White, I. Gollini, and T. B. Murphy. Review of statistical network analysis: models, algorithms, and software. *Statistical Analysis and Data Mining*, 5(4):243–264, 2012.

[195] S. Schmidl, P. Wenig, and T. Papenbrock. Anomaly detection in time series: a comprehensive evaluation. *Proceedings of the VLDB Endowment*, 15(9):1779–1797, 2022.

[196] R. Schmidt. Tail dependence. In *Statistical tools for finance and insurance*, pages 65–91. Springer, 2005.

[197] P. Schober, C. Boer, and L. A. Schwarte. Correlation coefficients: appropriate use and interpretation. *Anesthesia & analgesia*, 126(5):1763–1768, 2018.

[198] G. W. Schwert. Stock volatility in the new millennium: how wacky is Nasdaq? *Journal of Monetary Economics*, 49(1):3–26, 2002.

[199] J. M. Shine and R. A. Poldrack. Principles of dynamic network reconfiguration across diverse brain states. *NeuroImage*, 180:396–405, 2018.

[200] M. Sibuya. Bivariate extreme statistics, I. *Annals of the Institute of Statistical Mathematics*, 11(3):195–210, 1960.

[201] A. K. Singh, D. E. Allen, and R. J. Powell. Tail dependence analysis of stock markets using extreme value theory. *Applied Economics*, 49(45):4588–4599, 2017.

[202] M. Sklar. Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris*, 8:229–231, 1959.

[203] W.-M. Song, T. Di Matteo, and T. Aste. Hierarchical information clustering by means of topologically embedded graphs. *PloS one*, 7(3):e31929, 2012.

[204] M. Starnini, M. Boguñá, and M. Á. Serrano. The interconnected wealth of nations: Shock propagation on global trade-investment multiplex networks. *Scientific reports*, 9(1):1–10, 2019.

[205] U. Stelzl, U. Worm, M. Lalowski, C. Haenig, F. H. Brembeck, H. Goehler, M. Stroedicke, M. Zenkner, A. Schoenherr, S. Koeppen, et al. A human protein-protein interaction network: a resource for annotating the proteome. *Cell*, 122(6):957–968, 2005.

[206] S. C. Steude. Weighted maximum likelihood for risk prediction. *NCCR FINRISK*, (689), 2011.

[207] S. T. Straetmans, W. F. Verschoor, and C. C. Wolff. Extreme US stock market fluctuations in the wake of 9/11. *Journal of Applied Econometrics*, 23(1):17–42, 2008.

[208] J.-i. Takeuchi and K. Yamanishi. A unifying framework for detecting outliers and change points from time series. *IEEE transactions on Knowledge and Data Engineering*, 18(4):482–492, 2006.

[209] P. Tasca, S. Battiston, and A. Deghi. Portfolio diversification and systemic risk in interbank networks. *Journal of Economic dynamics and control*, 82:96–124, 2017.

[210] I. W. Taylor, R. Linding, D. Warde-Farley, Y. Liu, C. Pesquita, D. Faria, S. Bull, T. Pawson, Q. Morris, and J. L. Wrana. Dynamic modularity

in protein interaction networks predicts breast cancer outcome. *Nature biotechnology*, 27(2):199–204, 2009.

[211] J. B. Taylor. The financial crisis and the policy responses: An empirical analysis of what went wrong. Technical report, National Bureau of Economic Research, 2009.

[212] E. Terazi and S. Şenel. The effects of the global financial crisis on the central and eastern European Union countries. *International Journal of Business and Social Science*, 2(17), 2011.

[213] A. V. Thakor. The financial crisis of 2007–2009: Why did it happen and what did we learn? *The Review of Corporate Finance Studies*, 4(2):155–205, 2015.

[214] M. Thomas, M. Lemaitre, M. L. Wilson, C. Viboud, Y. Yordanov, H. Wackernagel, and F. Carrat. Applications of extreme value theory in public health. *PloS one*, 11(7):e0159312, 2016.

[215] A. K. Tiwari, E. J. A. Abakah, O. S. Yaya, and K. O. Appiah. Tail risk dependence, co-movement and predictability between green bond and green stocks. *Applied Economics*, 55(2):201–222, 2023.

[216] M. Toivanen. Contagion in the interbank network: An epidemiological approach. *Bank of Finland Research Discussion Paper*, (19), 2013.

[217] V. Torre and T. A. Poggio. On edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (2):147–163, 1986.

[218] E. Towler, B. Rajagopalan, E. Gilleland, R. S. Summers, D. Yates, and R. W. Katz. Modeling hydrologic and water quality extremes in a changing climate: A statistical approach based on extreme value theory. *Water Resources Research*, 46(11), 2010.

[219] N. Tran, J. Buck, and C. Kluppelberg. Causal inference for extremes on river networks. In *EGU General Assembly Conference Abstracts*, pages EGU21–341, 2021.

[220] M. Tumminello, T. Aste, T. Di Matteo, and R. N. Mantegna. A tool for filtering information in complex systems. *Proceedings of the National Academy of Sciences*, 102(30):10421–10426, 2005.

[221] M. Tumminello, F. Lillo, and R. N. Mantegna. Correlation, hierarchies, and networks in financial markets. *Journal of economic behavior & organization*, 75(1):40–58, 2010.

[222] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proceedings of the 26th annual international conference on machine learning*, pages 1073–1080, 2009.

[223] G.-J. Wang, C. Xie, Z.-Q. Jiang, and H. E. Stanley. Extreme risk spillover effects in world gold markets and the global financial crisis. *International Review of Economics & Finance*, 46:55–77, 2016.

[224] S. Wasserman, K. Faust, et al. *Social network analysis: Methods and applications*. Cambridge university press, 1994.

[225] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world'networks. *Nature*, 393(6684):440–442, 1998.

[226] F. Wen, X. Yang, and W.-X. Zhou. Tail dependence networks of global stock markets. *International Journal of Finance & Economics*, 24(1):558–567, 2019.

[227] M. Wolf. *The Shifts and the Shocks: What we've Learned–and have Still to Learn–from the Financial Crisis*. Penguin UK, 2014.

[228] N. Xuan, V. Julien, S. Wales, and J. Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research*, 11(95), 2010.

[229] X. Yan, L. G. Jeub, A. Flammini, F. Radicchi, and S. Fortunato. Weight thresholding on complex networks. *Physical Review E*, 98(4):042304, 2018.

[230] C. C. Yang and T. D. Ng. Terrorism and crime related weblog social network: Link, content analysis and information visualization. In *2007 IEEE Intelligence and Security Informatics*, pages 55–58. IEEE, 2007.

[231] J. Yin and S. Sun. Incomplete multi-view clustering with cosine similarity. *Pattern Recognition*, 123:108371, 2022.

[232] M. Yuanyuan, Z. Xintian, and L. Linguan. Susceptible-Infected-Removed (SIR) model of crisis spreading in the correlated network of listed companies and their main stockholders. *Journal of Management Sciences in China*, 16(7):80–94, 2013.

[233] D. Zaring. The Government's economic response to the Covid crisis. *Rev. Banking & Fin. L.*, 40:315, 2020.

[234] J. Zhang, J. Fei, X. Song, and J. Feng. An improved Louvain algorithm for community detection. *Mathematical Problems in Engineering*, 2021(1):1485592, 2021.

[235] P. J. Zhao. A social network model of the COVID-19 pandemic. *MedRxiv*, pages 2020–03, 2020.

Appendices

Appendix A

Additional details

A.1 Modularity

Modularity, as defined in [165], is 'up to a multiplicative constant, the number of edges falling within groups minus the expected number in an equivalent network with edges placed at random.' Mathematically, the modularity can be expressed as:

$$Q = \frac{1}{4m} \sum_{i=1}^{n} \sum_{j=1}^{n} \left[A_{i,j} - \frac{k_i k_j}{2m} \right] \delta(i,j), \tag{A.1}$$

where n is the number of nodes in a network, $A_{i,j}$ is the i,j-th entry of the adjacency matrix A, k_i and k_j are the degrees of nodes i and j, respectively, m is the total number of edges in the network, and $\delta(i,j) = 1$ if i and j are in the same cluster and 0 otherwise. The intuition behind the formula in equation (A.1) is that the expected number of edges between i and j, given random placement of edges, is $k_i k_j/2$. Hence, the modularity is the sum $A_{i,j} - k_i k_j/2$ over all pairs of nodes. The factor of $1/4m$ is included for compatibility with previous definitions of modularity.

A.2 Louvain community algorithm

The Louvain algorithm is a hierarchical clustering method, which extracts non-overlapping communities from networks. Its core principle involves iteratively placing nodes into different groups in order to maximize the modularity of the community partition, ultimately achieving the optimal community division. Algorithm 8 shows the step-by-step procedure of the Louvain algorithm. More details can be found in [234].

Algorithm 8 Louvain community algorithm

1. Place each node i into its own separate community c_i.
2. For each node i, find all the communities connected to it and calculate the change of modularity after moving the node to each neighboring community. Place node i in the community which gives highest gain in modularity score. If no increase is possible, then node i stays in its original community.
3. Merge each community in step 2 into a single node. Edges connecting nodes between the different communities are likewise reduced to a single weighted edge.
4. If after a complete iteration the modularity of the graph has not increased, terminate the algorithm. Else, return to step 2.

A.3 Adjusted mutual information

Let S be a set of n elements and $U = \{U_1, U_2, \ldots, U_m\}$ and $V = \{V_1, V_2, \ldots, V_p\}$ be partitions of S with m and p clusters, respectively. The partitions are assumed to be so-called 'hard clusters', i.e. $U_i \cap U_j = \emptyset = V_i \cap V_j$ for all $i \neq j$, and complete, i.e. $\cup_{i=1}^{m} U_i = \cup_{j=1}^{p} V_j = S$. The mutual information (MI) between U and V is defined as:

$$MI(U,V) = \sum_{i=1}^{m} \sum_{j=1}^{p} P_{U,V}(i,j) \log \left(\frac{P_{U,V}(i,j)}{P_U(i) P_V(j)} \right), \tag{A.2}$$

where $P_U(i) = |U_i|/n$, $P_V(i) = |V_i|/n$ and $P_{U,V}(i,j) = |U_i \cap V_j|/n$. The entropy, associated with clustering U, is then defined as:

$$H(U) = -\sum_{i=1}^{m} P_U(i) \log P_U(i). \tag{A.3}$$

The baseline value of mutual information between two random clusterings is not constant and tends to increase when both partitions contain a greater number of clusters (while the total number of elements n remains fixed). The expected MI can be computed by adopting a hypergeometric model of randomness, as shown in [222], and takes the following form:

$$E[MI(U,V)] = \sum_{i=1}^{m}\sum_{j=1}^{p}\sum_{n_{ij}=\max(a_i+b_j-n,0)}^{\min(a_i,b_j)} \frac{n_{i,j}}{n} \log\left(\frac{n \cdot n_{ij}}{a_i b_j}\right) \times \frac{a_i! b_j! (n-a_i)!(n-b_j)!}{n! n_{ij}! (a_i - n_{ij})!(b_j - n_{ij})!(n - a_i - b_j + n_{ij})!}, \tag{A.4}$$

where $a_i = \sum_{j=1}^{p} n_{i,j}$ and $b_j = \sum_{i=1}^{m} n_{i,j}$. The AMI is then defined as:

$$AMI(U,V) = \frac{MI(U,V) - E[MI(U,V)]}{\max(H(U), H(V)) - E[MI(U,V)]}. \tag{A.5}$$

While the MI measures the information shared between two clusterings without considering the possibility of agreement by chance, the AMI adjusts MI by subtracting the expected mutual information between random clusterings, thus taking into account the baseline level of agreement that can be expected by chance alone. By using AMI, we can obtain a measure that not only captures the shared information between clusterings but also considers the level of agreement that is likely to occur randomly, which makes AMI a more suitable measure when evaluating and comparing clusterings in practice.

Appendix B

Additional figures

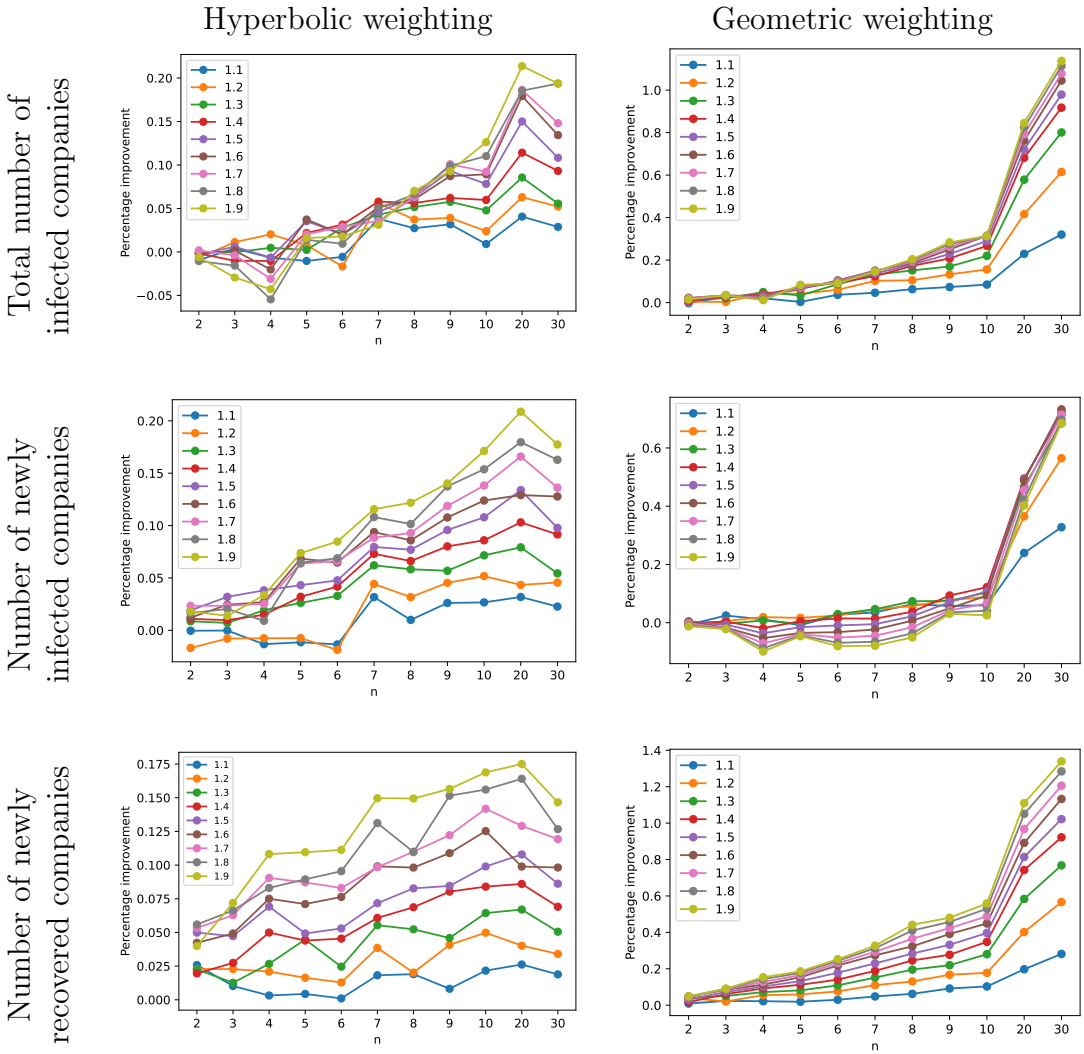

Figure B.1: The plots indicate the percentage improvement when using the hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results when $\rho = 1$ (equal weighting for all observations). These improvements are measured in predicting the total number of infected companies (top row), newly infected companies (middle row), and newly recovered companies (bottom row) during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 10$ days.

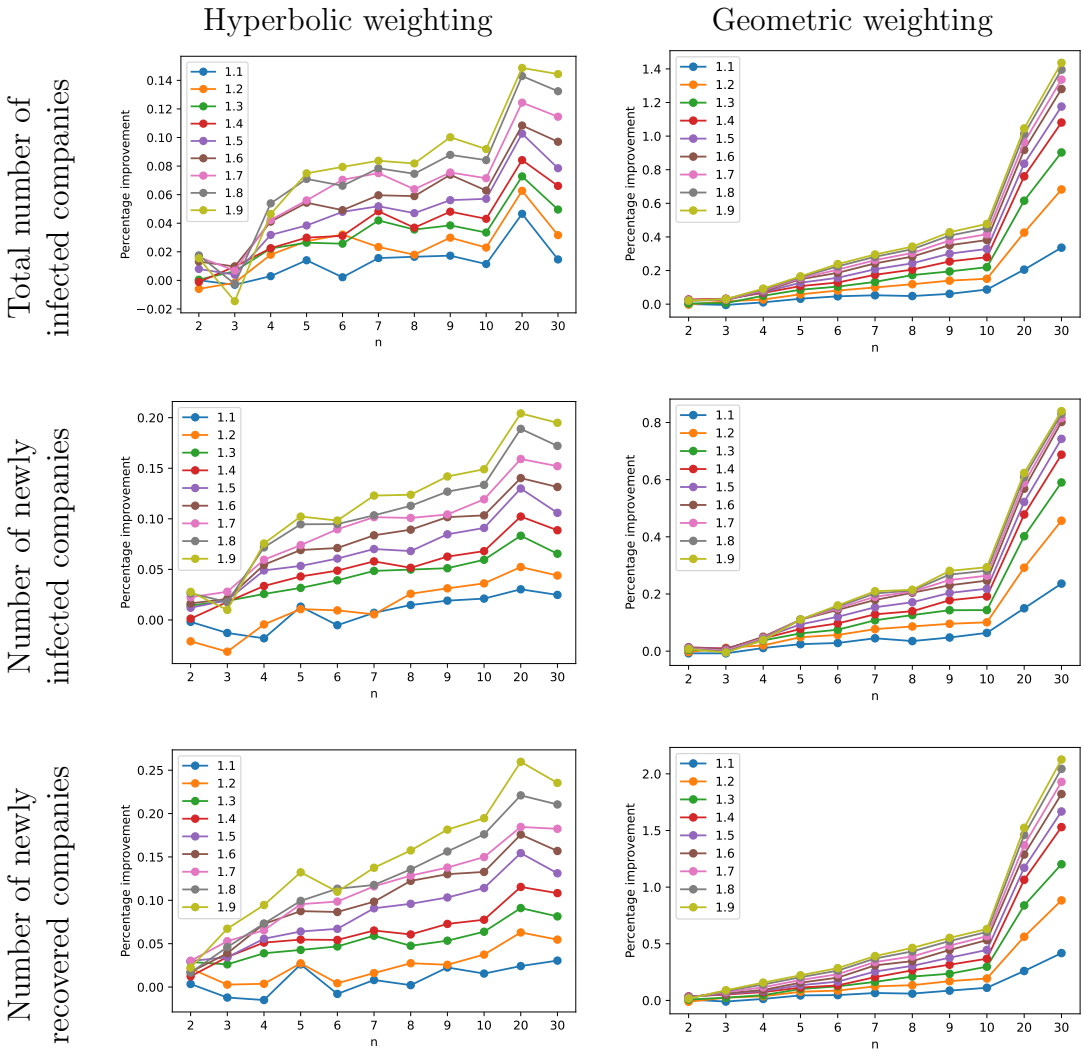

Figure B.2: The plots indicate the percentage improvement when using the hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results when $\rho = 1$ (equal weighting for all observations). These improvements are measured in predicting the total number of infected companies (top row), newly infected companies (middle row), and newly recovered companies (bottom row) during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 20$ days.

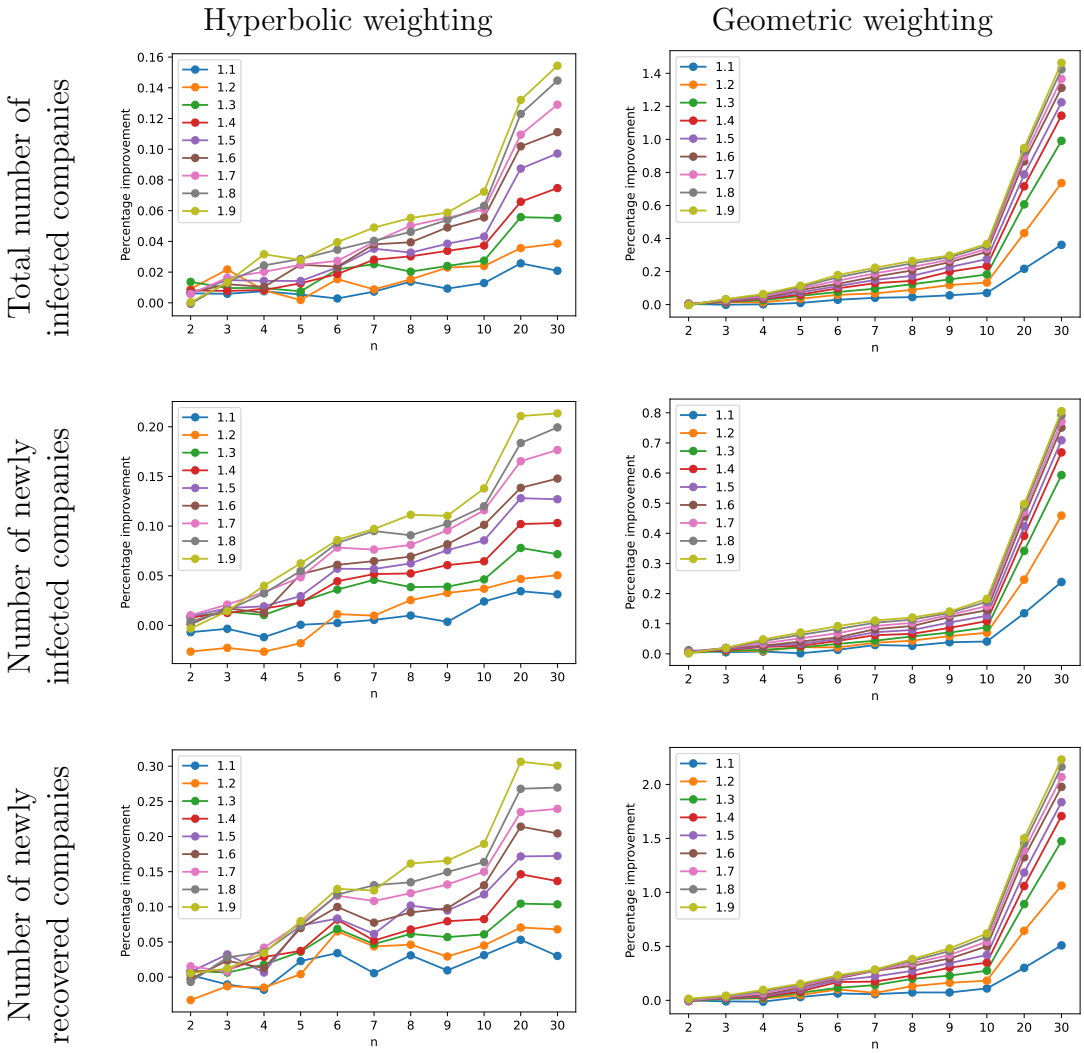

Figure B.3: The plots indicate the percentage improvement when using the hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results when $\rho = 1$ (equal weighting for all observations). These improvements are measured in predicting the total number of infected companies (top row), newly infected companies (middle row), and newly recovered companies (bottom row) during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 30$ days.

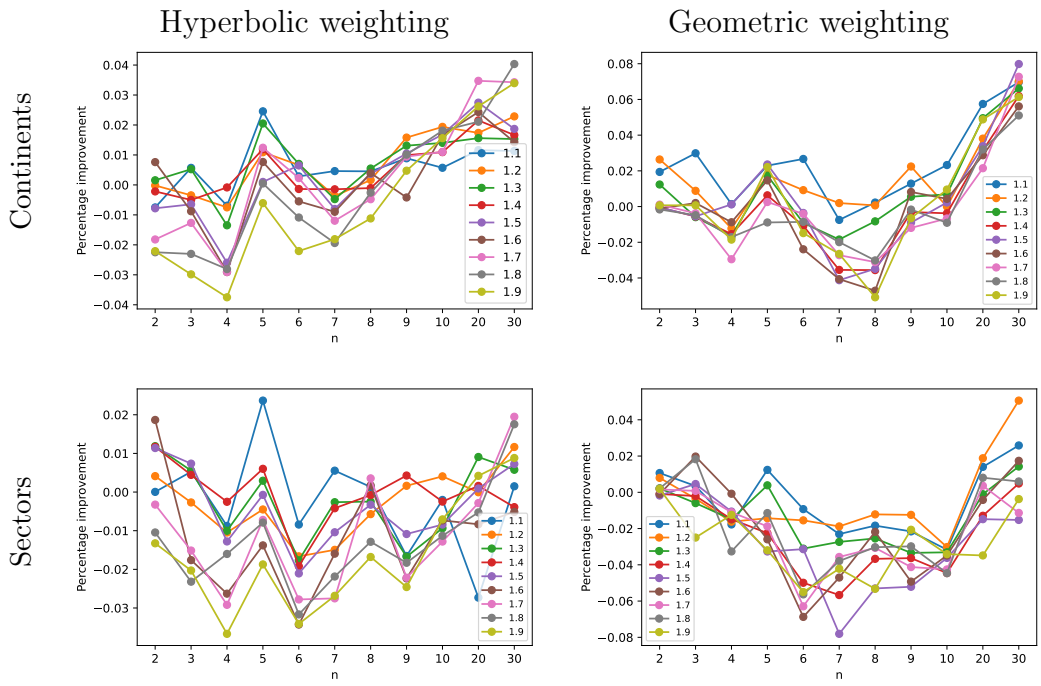

Figure B.4: The plots show the percentage improvement in the Sørensen–Dice coefficient when using the hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured in predicting the continents (top row) and the sectors (bottom row) of the newly infected companies during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 10$ days.

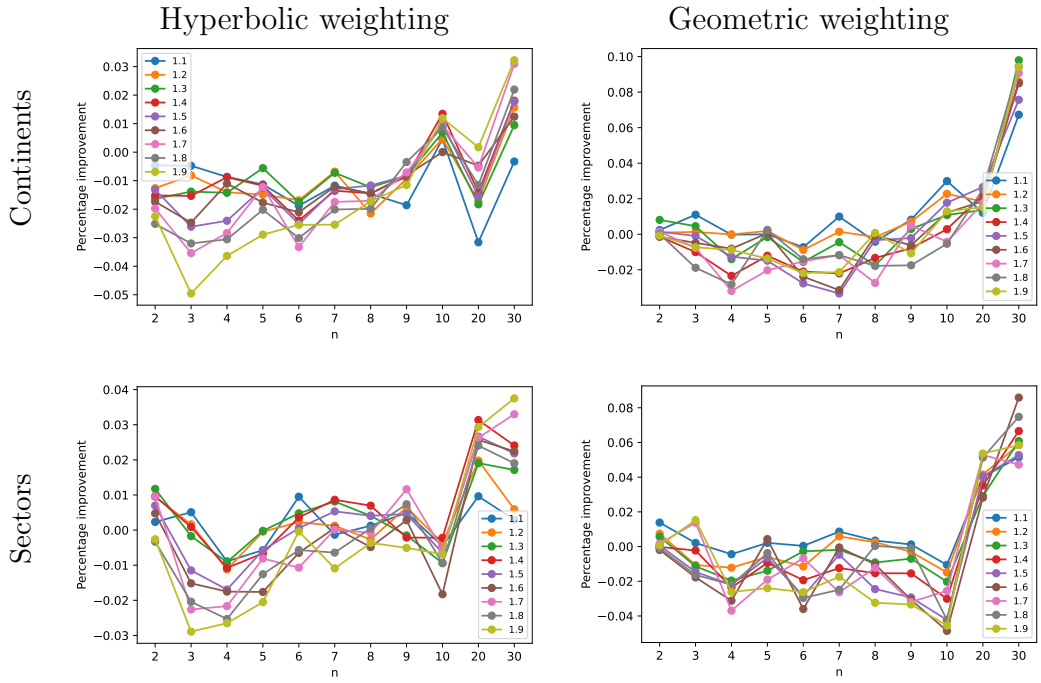

Figure B.5: The plots show the percentage improvement in the Sørensen–Dice coefficient when using the hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured in predicting the continents (top row) and the sectors (bottom row) of the newly infected companies during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 20$ days.

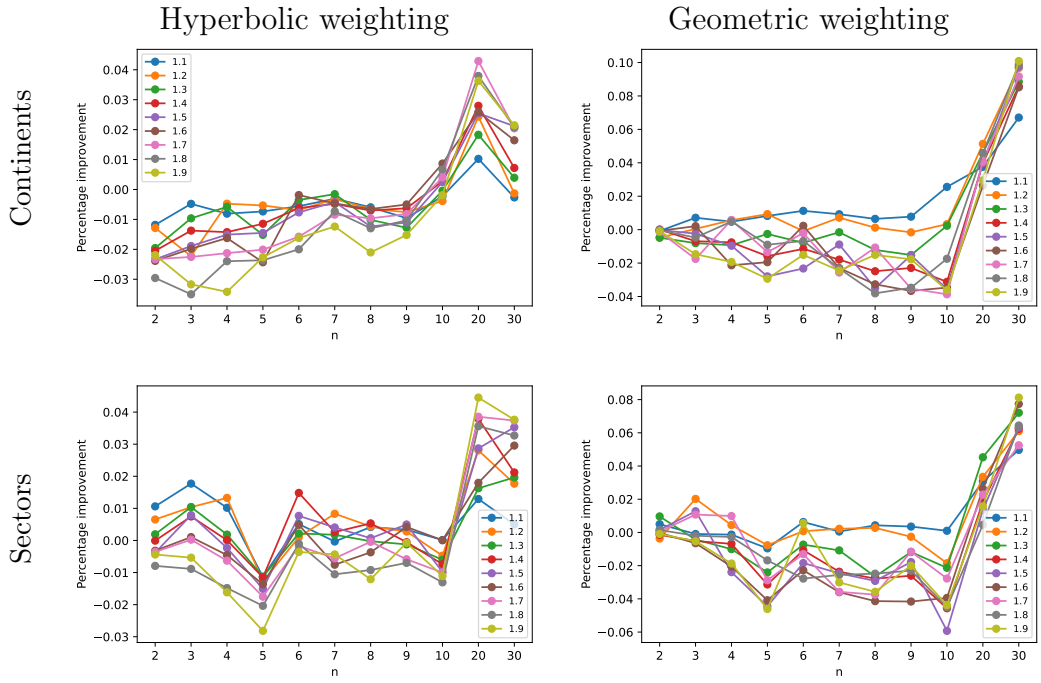

Figure B.6: The plots show the percentage improvement in the Sørensen–Dice coefficient when using the hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured in predicting the continents (top row) and the sectors (bottom row) of the newly infected companies during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 30$ days.

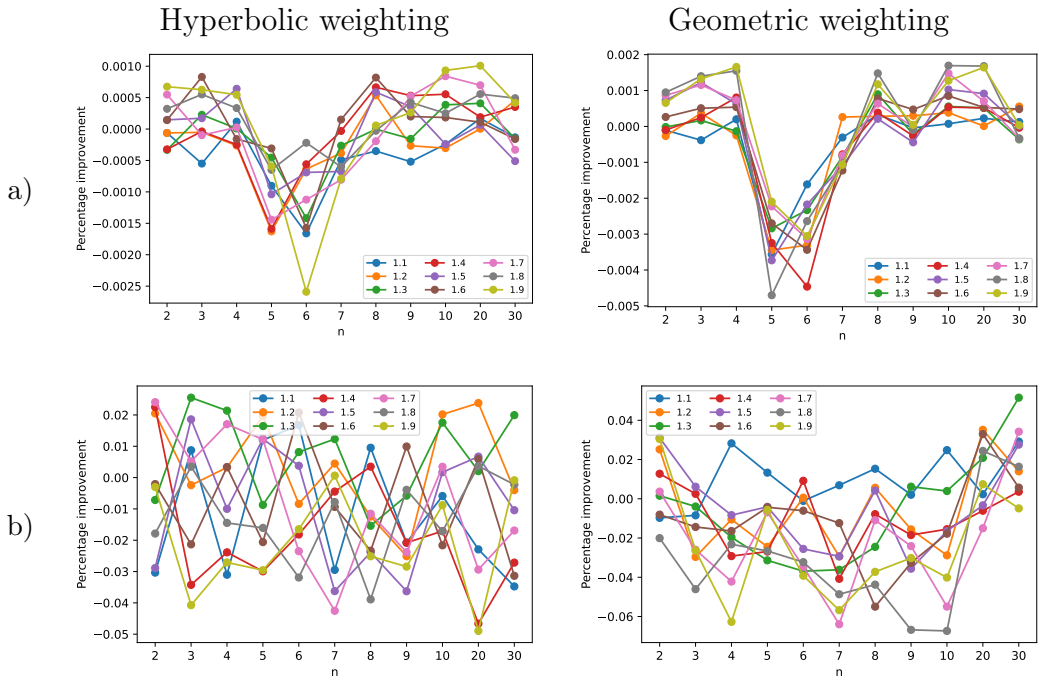

Figure B.7: The plots show the percentage improvement in the a) Accuracy and b) F_1-score when using the TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured for the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 10$ days.

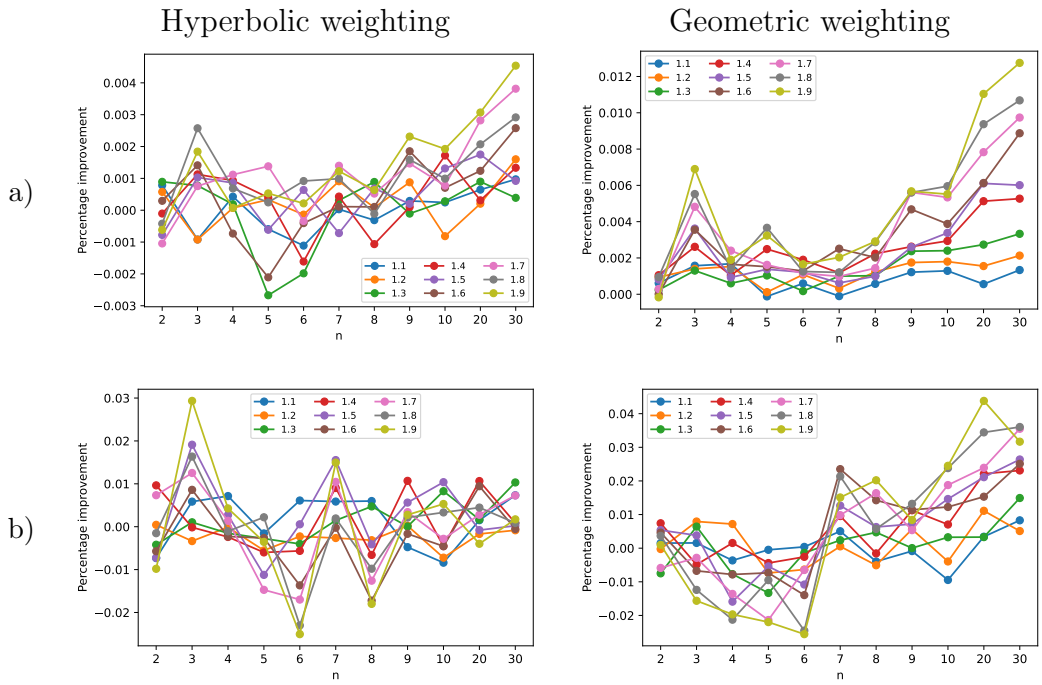

Figure B.8: The plots show the percentage improvement in the a) Accuracy and b) F_1-score when using the TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured for the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 20$ days.

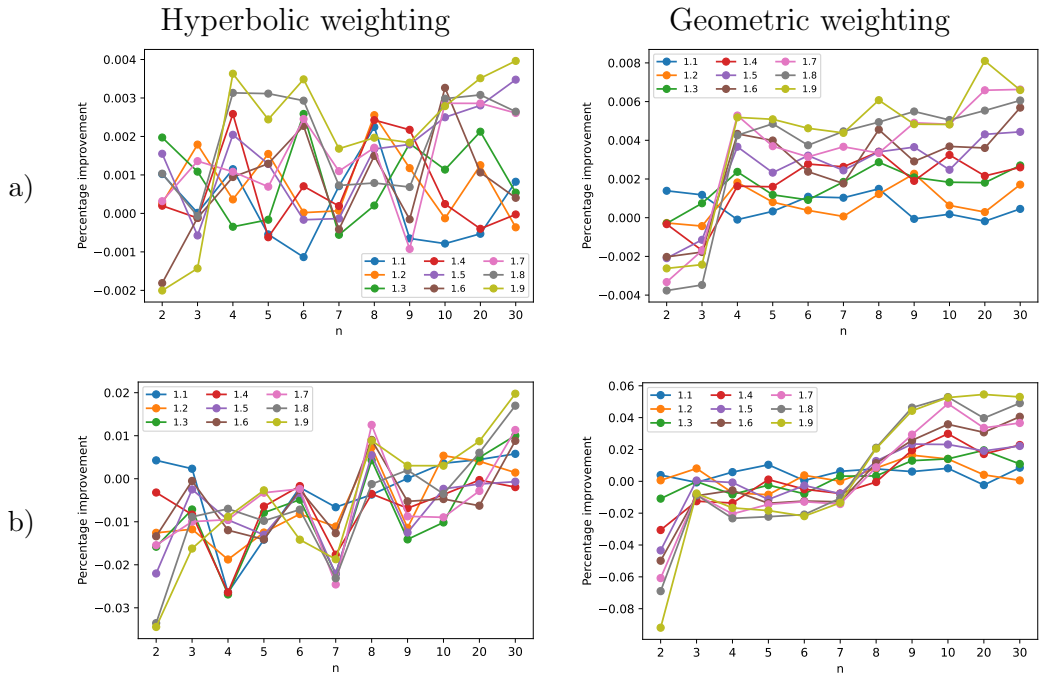

Figure B.9: The plots show the percentage improvement in the a) Accuracy and b) F_1-score when using the TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured for the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 30$ days.

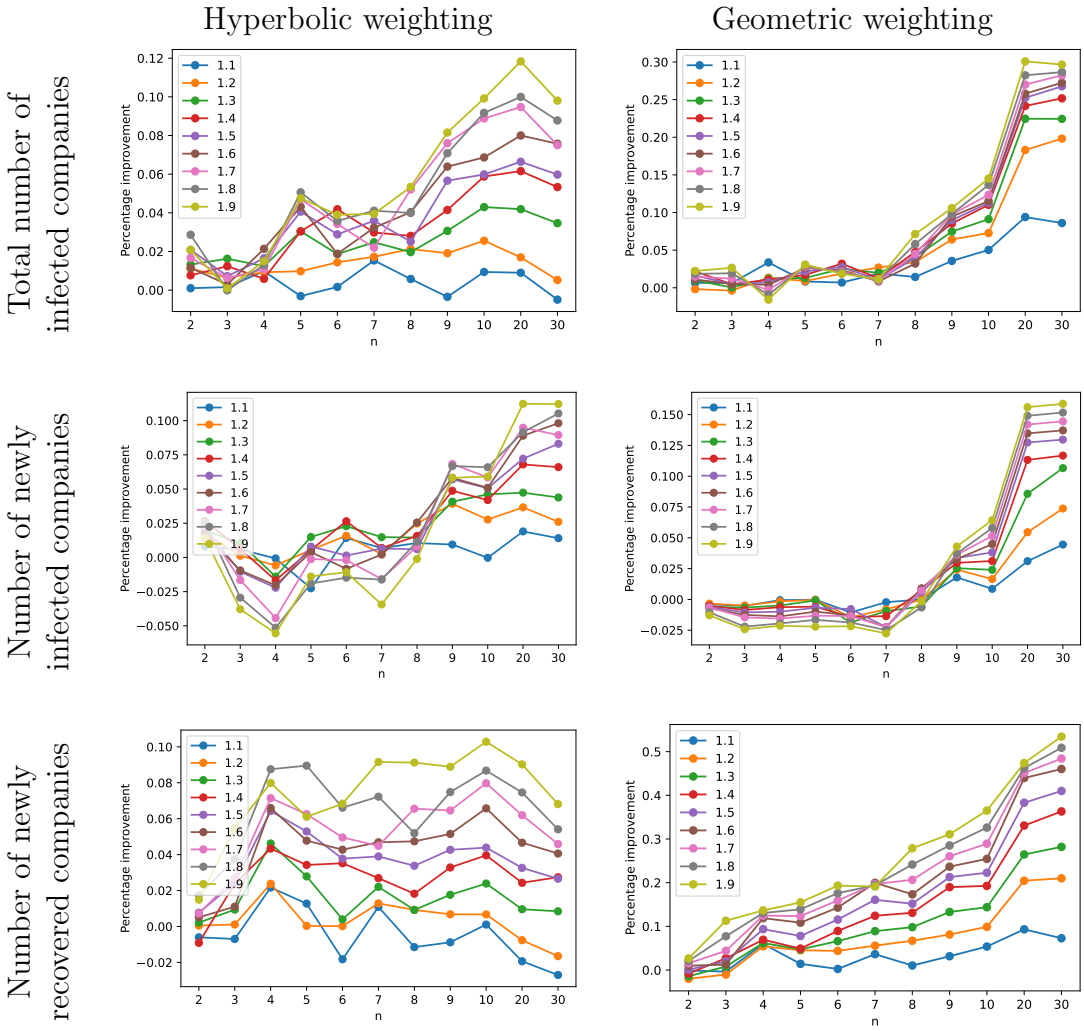

Figure B.10: The plots illustrate the percentage improvement when using the hyperbolic (left column) and geometric (right column) weighting schemes in the CPD+TiWML model for different values of $\rho > 1$, compared to the results when using the change point model, introduced in Chapter 4. These improvements are measured in predicting the total number of infected companies (top row), newly infected companies (middle row), and newly recovered companies (bottom row) during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 10$ days.

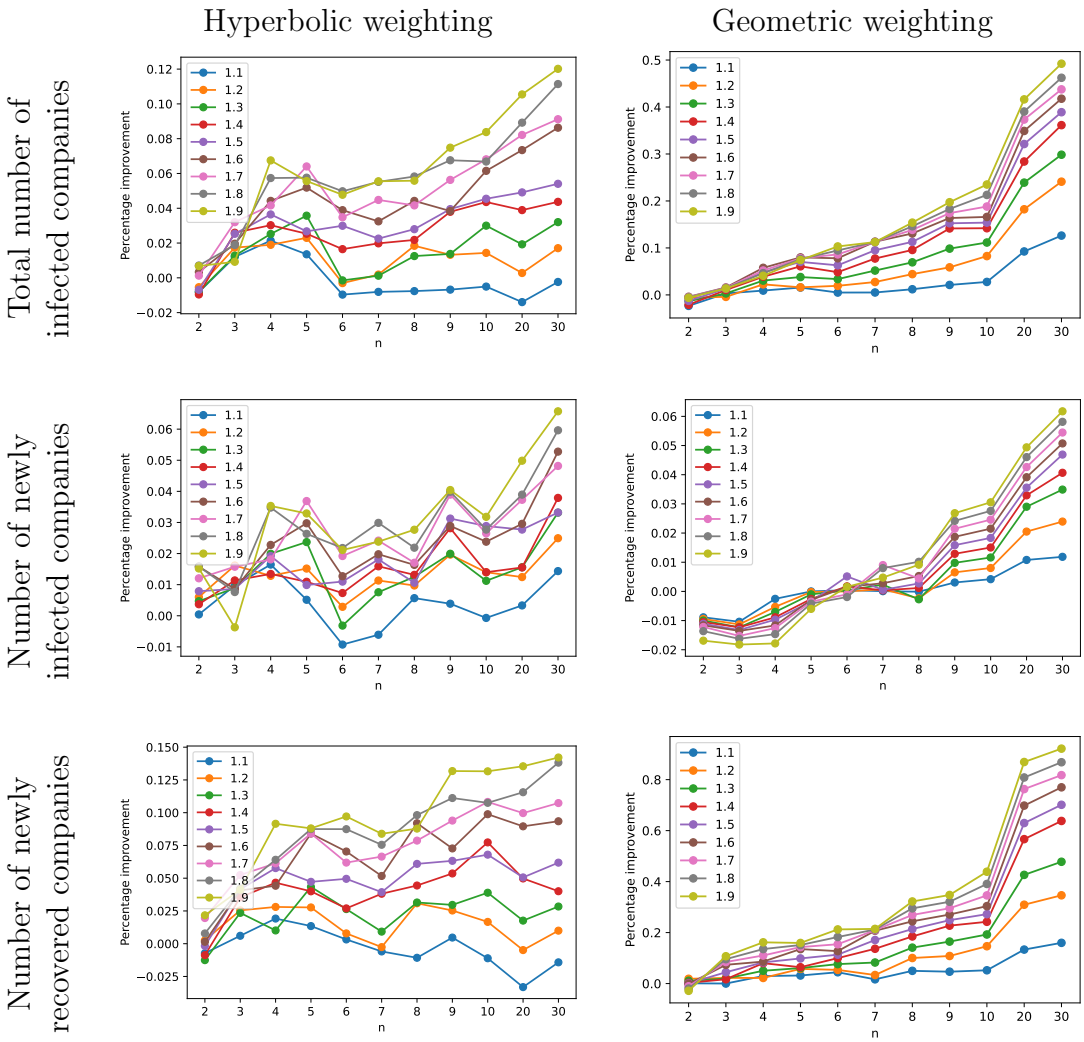

Figure B.11: The plots illustrate the percentage improvement when using the hyperbolic (left column) and geometric (right column) weighting schemes in the CPD+TiWML model for different values of $\rho > 1$, compared to the results when using the change point model, introduced in Chapter 4. These improvements are measured in predicting the total number of infected companies (top row), newly infected companies (middle row), and newly recovered companies (bottom row) during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 20$ days.

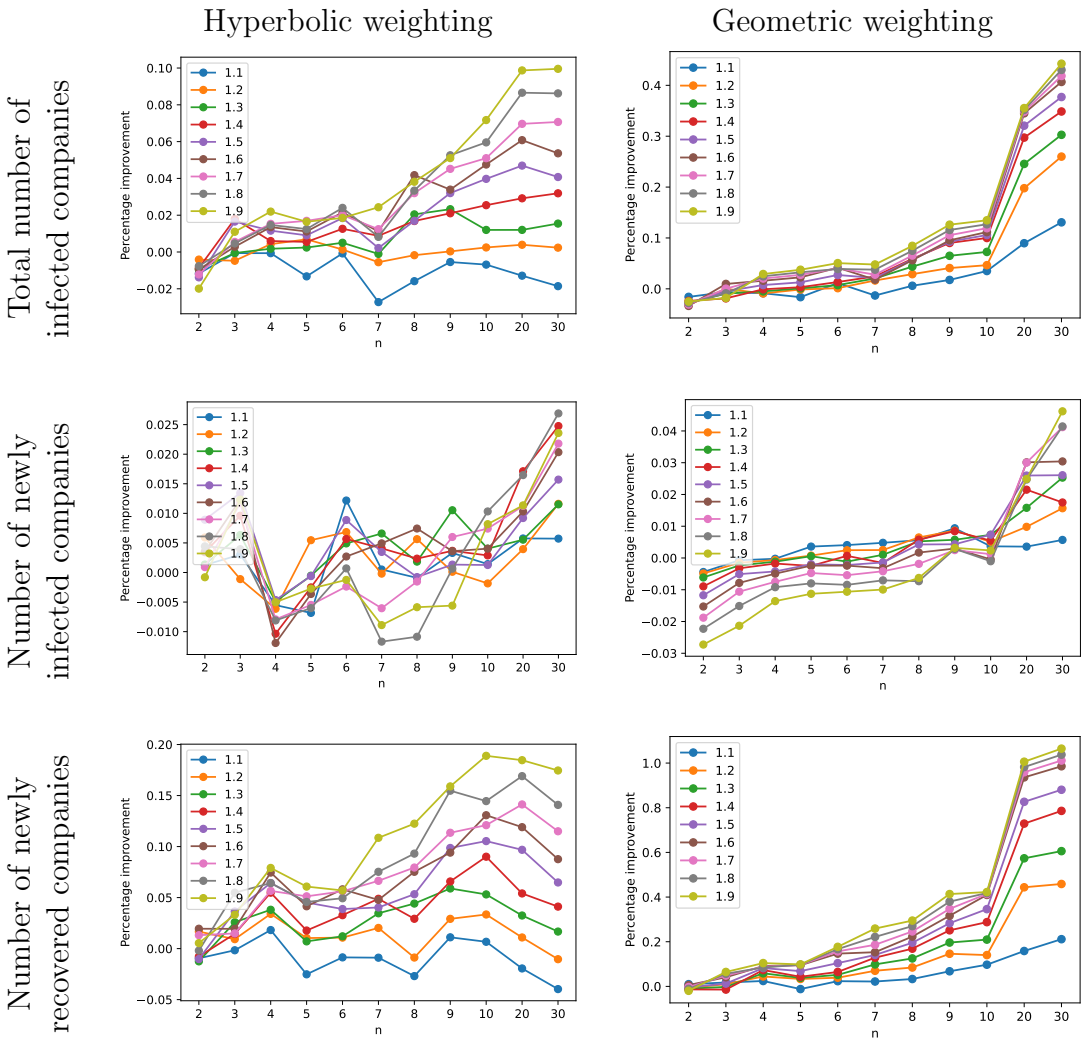

Figure B.12: The plots illustrate the percentage improvement when using the hyperbolic (left column) and geometric (right column) weighting schemes in the CPD+TiWML model for different values of $\rho > 1$, compared to the results when using the change point model, introduced in Chapter 4. These improvements are measured in predicting the total number of infected companies (top row), newly infected companies (middle row), and newly recovered companies (bottom row) during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 30$ days.

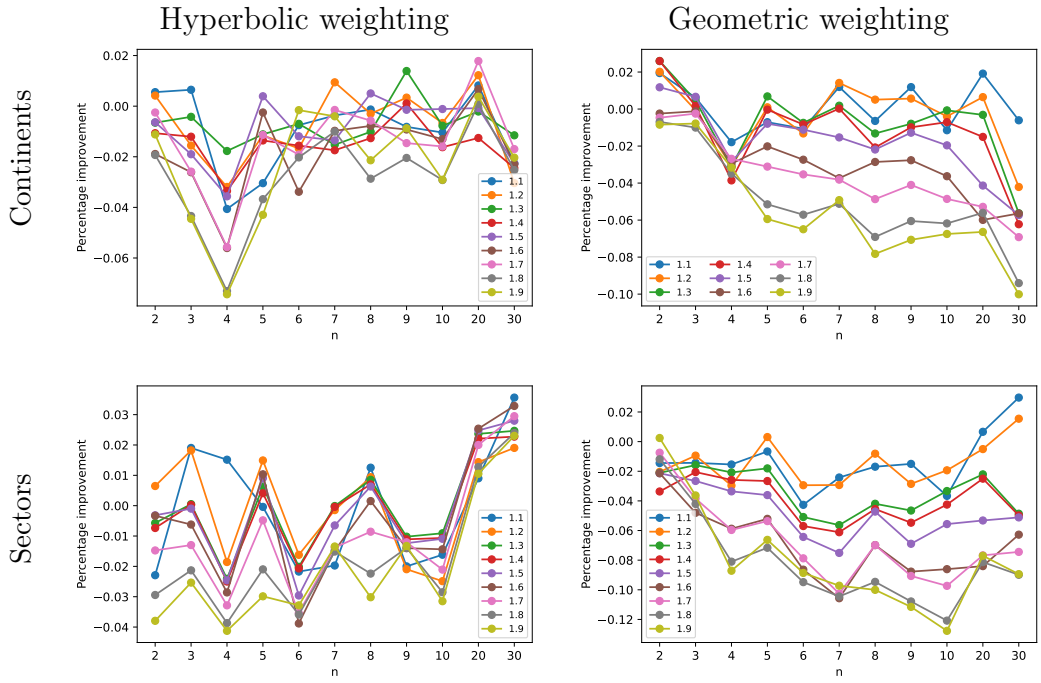

Figure B.13: The plots show the percentage improvement in the Sørensen–Dice coefficient when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured in predicting the continents (top row) and the sectors (bottom row) of the newly infected companies during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 10$ days.

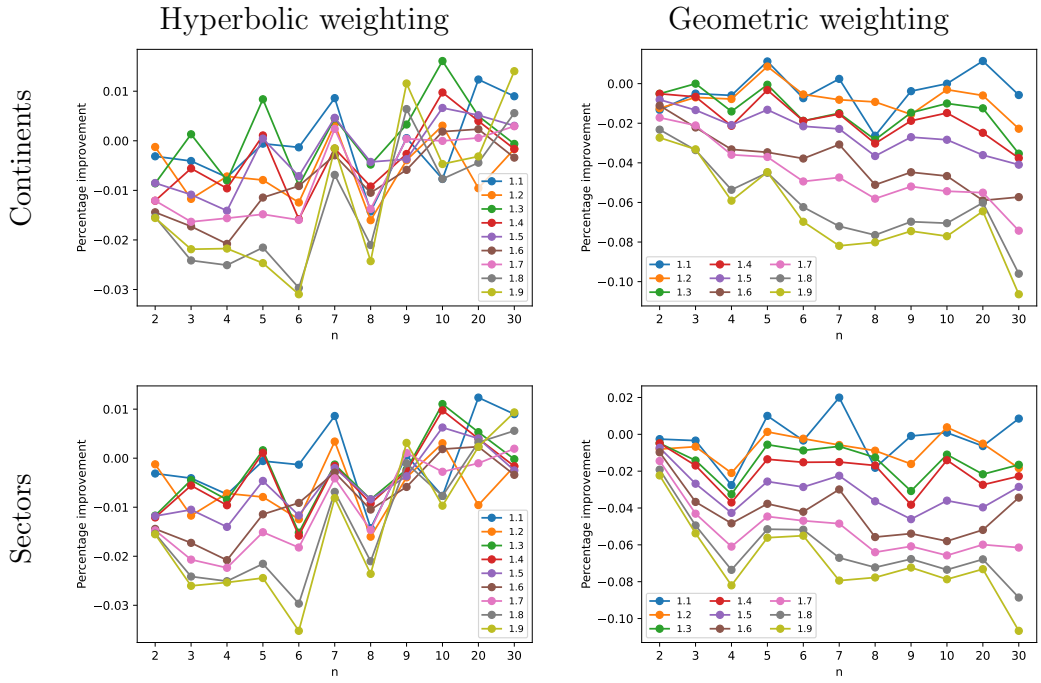

Figure B.14: The plots show the percentage improvement in the Sørensen–Dice coefficient when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured in predicting the continents (top row) and the sectors (bottom row) of the newly infected companies during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 20$ days.

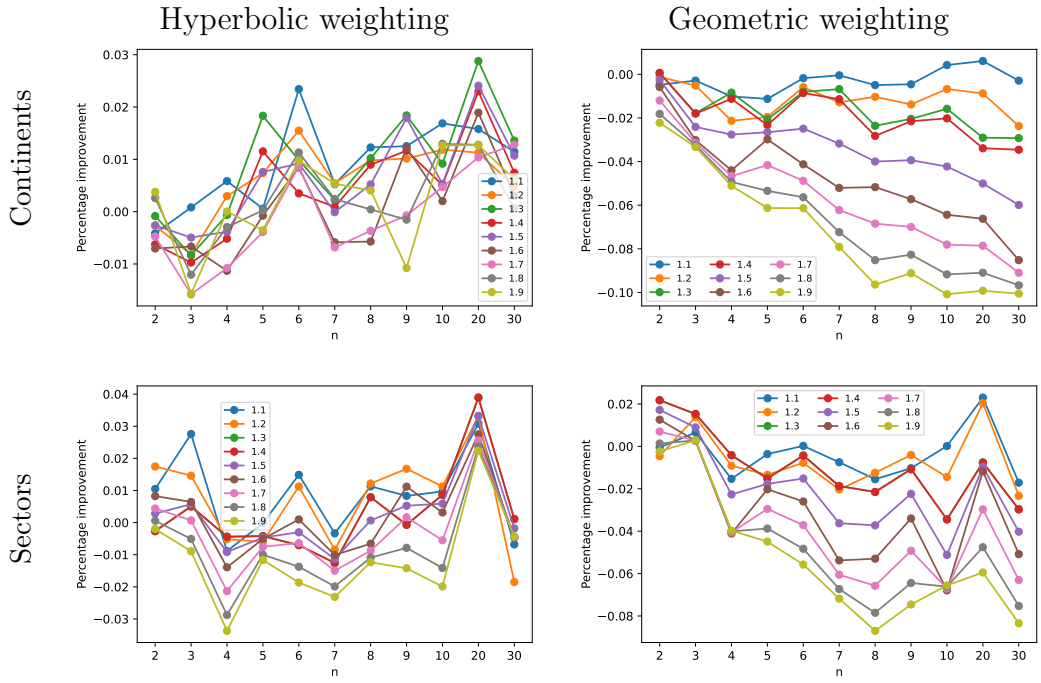

Figure B.15: The plots show the percentage improvement in the Sørensen–Dice coefficient when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured in predicting the continents (top row) and the sectors (bottom row) of the newly infected companies during the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 30$ days.

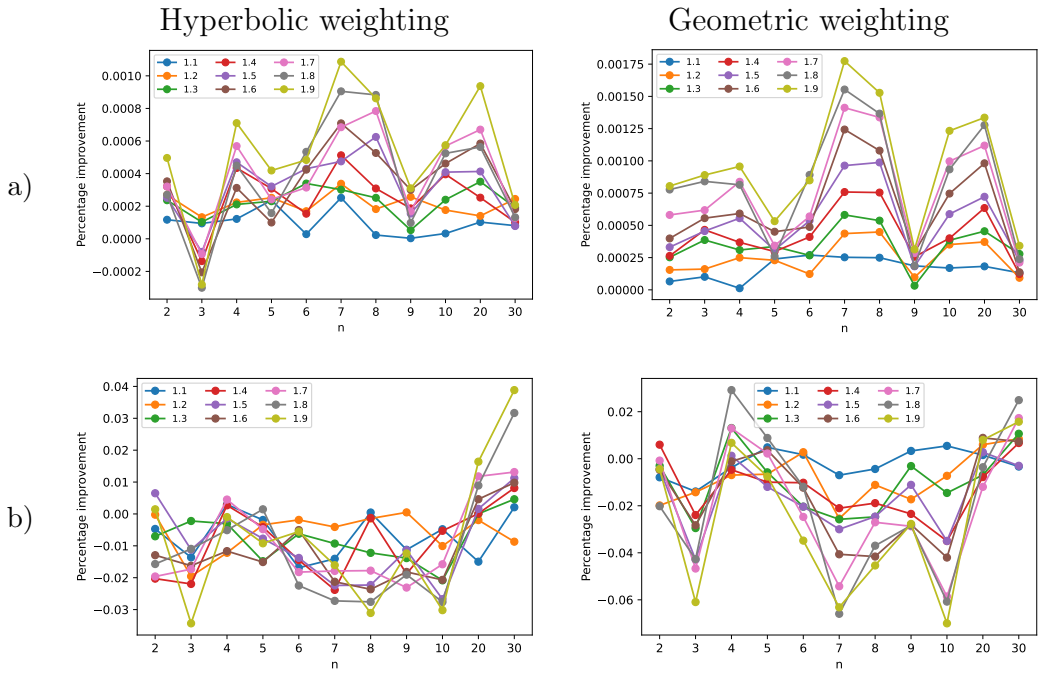

Figure B.16: The plots show the percentage improvement in the a) Accuracy and b) F_1-score when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured for the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 10$ days.

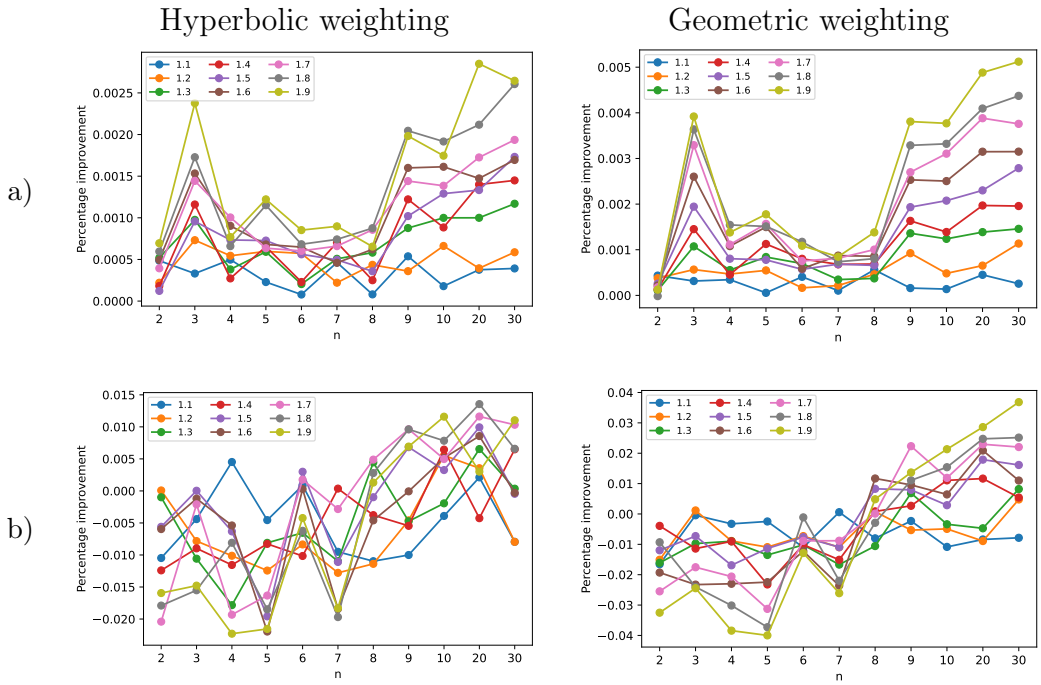

Figure B.17: The plots show the percentage improvement in the a) Accuracy and b) F_1-score when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured for the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 20$ days.

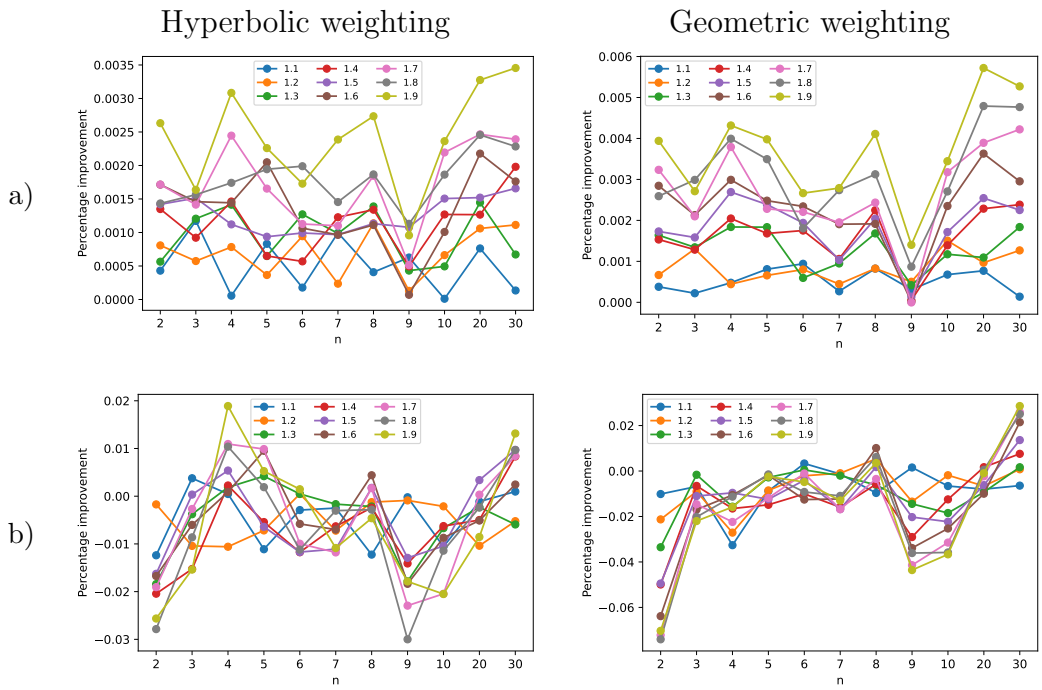

Figure B.18: The plots show the percentage improvement in the a) Accuracy and b) F_1-score when using the CPD+TiWML model with hyperbolic (left column) and geometric (right column) weighting schemes for different values of $\rho > 1$, compared to the results for $\rho = 1$ (equal weighting for all observations). The improvements are measured for the 2020 financial crisis, using infection data from the previous n days with a prediction horizon of $k = 30$ days.