



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Bozhidarova, M., Ball, F., van Gennip, Y., D O'Dea, R., & Stupfler, G. (2026). Enhancing financial crisis prediction: integrating change point detection for exogenous event identification. *Applied Network Science*, 11(1), Article 15. <https://doi.org/10.1007/s41109-025-00748-1>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

RESEARCH

Open Access



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

Malvina Bozhidarova^{1,2*}, Frank Ball², Yves van Gennip³, Reuben D O'Dea² and Gilles Stupfler⁴

*Correspondence:

Malvina Bozhidarova
mbozhidarova@russell.co.uk

¹Russell Group Limited,
NG1 1HS Nottingham, UK

²School of Mathematical Sciences,
University of Nottingham,
NG7 2RD Nottingham, UK

³Delft Institute of Applied
Mathematics (DIAM), Technische
Universiteit Delft, Delft, The
Netherlands

⁴Univ Angers, CNRS, LAREMA, SFR
MATHSTIC, F-49000 Angers, France

Abstract

This paper explores the integration of change point detection (CPD) techniques to improve the adaptability of financial contagion models 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 a real-time CPD algorithm, we enable our model to respond effectively to these disruptions, resulting in more robust and accurate predictions. We analyze stock price, geographical location, and economic sector data for a dataset of 398 companies to construct multiplex networks with four layers. On these networks, we implement a Susceptible–Infected–Recovered (SIR) transmission model to simulate the spread of financial shocks among companies, accounting for their interconnectedness. Using stock price data from the 2008 to 2020 financial crises, we evaluate the model's ability to predict the propagation of financial shocks through the network, where shocks are identified based on stock price volatility. We continuously monitor the data for anomalies and when a change point is identified, the model discards the older data before the change point and focuses on the more recent data. We demonstrate the effectiveness in incorporating change points for improving the model's predictive accuracy.

Keywords Epidemic modelling, Financial crisis, Multiplex network, SIR model, Change point detection

Introduction

Detecting significant changes in financial time series data is a critical aspect of quantitative analysis, especially in understanding the dynamics of market behaviour 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 (Chevapatrakul and Tee 2014). 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. For instance, in the 2008 financial crisis, governments worldwide implemented various decisions

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

and interventions to stabilize financial markets and mitigate the 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 (Mishkin 2009). During the 2020 financial crisis, government policy regarding lockdowns, travel restrictions and quarantine measures significantly impacted various industries, leading to market disruptions, supply chain interruptions and shifts in consumer behaviour. Many sectors, such as travel, hospitality, entertainment and retail experienced substantial declines in revenue during the crisis, impacting stock prices and overall market performance (Naseer et al. 2023). 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 (Makin and Layton 2021).

Other events, such as wars and geopolitical tensions, natural disasters and protests 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 (Nikkinen et al. 2008). 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. Following the earthquake and tsunami in Japan in 2011, the Fukushima nuclear disaster led to widespread disruptions in Japan's economy (Asongu 2012). 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 (Kawashima and Takeda 2012). Brexit, the withdrawal of the United Kingdom (UK) from the European Union (EU), also had a substantial impact on the financial markets (Hohlmeier and Fahrholz 2018). The uncertainty surrounding Brexit negotiations led to significant fluctuations in the value of the British pound (GBP) and UK-focused stocks, particularly those reliant on EU trade or with substantial exposure to the domestic economy, were especially sensitive to Brexit developments (Belke et al. 2018).

In this paper we enhance the model introduced in Bozhidarova et al. (2024) by incorporating change point detection (CPD) techniques in order to enable the model to react to such exogenous shocks. In Bozhidarova et al. (2024) we introduced an innovative framework for modelling financial contagion, based on an SIR (Susceptible – Infected – Recovered) epidemic model applied to a multiplex network of interconnected companies constructed from financial data. Our model employs a stochastic epidemic transmission mechanism, allowing financial crises to propagate both locally (to neighboring nodes within the network) and globally (to any company in the system), thereby simulating how a financial shock spreads from the original infected companies to the others. The approach introduced in the current paper enhances this model to enable the dynamic integration of significant events in real-time, leading to a more robust and adaptive predictions. The primary novelty of the current work lies in the integration of CPD, SIR dynamics and a data-driven multiplex network within a unified framework. This combination enables the model to simulate realistic contagion paths, while also detecting and adapting to abrupt structural shifts in financial markets. To our knowledge, while earlier research such as Acemoglu et al. (2015) and Battiston et al. (2016) focused on discussing theoretical properties of financial networks and perspectives for their practical use, this

is the first application of such an integrated approach in the context of financial contagion modelling.

The remainder of the paper is organized as follows. In Sect. 2 we describe the dataset and present the two components of the modelling framework: the procedure to construct the multiplex network and the transmission mechanism. The section summarizes the methodology, introduced in detail in Bozhidarova et al. (2024), while ensuring clarity for readers unfamiliar with the prior study. In Sect. 3 we build on previous work by incorporating change points into the model. We apply the model to the 2008 and the 2020 financial crises. The section begins by defining the concept of infection in a financial context and explores how the model can leverage historical data to predict future infections during each crisis. In particular, in Sect. 3.2 we introduce the ChangeFinder algorithm (Takeuchi and Yamanishi 2006) and demonstrate how it can be incorporated into the model to detect significant changes and adapt to them as soon as they happen. We demonstrate the effectiveness of incorporating CPD techniques into the model by notably improving forecasts of the number of affected companies. Section 4 provides a conclusion, discusses the implications of our findings, outlines the limitations of our approach, and proposes potential areas for future research. Details of our estimation approach and further finite-sample results are provided in the Appendix.

Modelling framework

Our primary objective is to model the dynamics of financial contagion within a population of companies. In particular, we aim to predict the future impact of financial crises on individual companies, financial sectors and continents, leveraging historical data from past crisis days. Our modelling framework consists of two main parts. First, we construct a multiplex financial network with nodes representing individual companies and edges within each layer representing different types of connections between companies. The multiplex network approach allows us to account for the various ways in which financial contagion may propagate. Next, we apply an SIR epidemic model to each network layer. In the standard SIR framework, the transmission and recovery probabilities, or their dependencies on other unknown parameters, are often assumed or set a priori, because the infection and/or recovery dates are not known (Cauchemez and Ferguson 2008). However, there has been a significant body of work on parameter estimation in epidemic models (Becker 1989; Held et al. 2019), in line with which we seek to reconstruct sequences of infection and recovery events, and therefore we treat transmission and recovery probabilities as unknown parameters to be inferred from crisis data. Specifically, the transmission probabilities represent the daily likelihood of an infected company spreading the contagion to a susceptible company via a given network layer, while the recovery probabilities represent the daily likelihood of an infected node recovering. These parameters are estimated using a maximum likelihood approach, which fits the model to the (reconstructed) observed sequences of infection and recovery events from past crisis episodes. This allows us to tailor the model to real-world dynamics rather than rely on assumed values. We then identify the most vulnerable companies, financial sectors and continents by simulating the spread of financial contagion using the SIR model with these estimated parameters over a given time horizon of interest.

Data

The analysis in this paper relies on the closing daily stock price of 398 companies, spanning from 17/01/2002 to 18/07/2022, representing $N = 5229$ trading days. We emphasise that, in this and all future instances, datasets defined over specified date ranges are considered to include both the start and end dates. The data are sourced 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. The companies are categorized into groups, based on the Orbis¹ company database. Firstly, they are grouped by the geographical location of their headquarters, resulting in six categories: Africa (2), Asia (77), Europe (115), North America (194), Oceania (9) and South America (1). Secondly, the companies are separated into 13 groups based on their primary economic sector: 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). The numbers in parentheses indicate the total number of companies in each respective group. More details showing the characteristics of the dataset can be found in Bozhidarova et al. (2024).

Network construction

We construct a multiplex network (see Fig. 1), in which nodes represent companies and each layer represents a distinct type of connection between them. The network consists of four layers: a tail dependence network layer, a continents layer, a sectors layer and a global layer. The following subsections explain the rationale and methodology used to construct these layers.

Tail dependence network layer

To study the likelihood of two companies experiencing extreme losses simultaneously we build complex financial networks based on the concept of tail dependence from Extreme Value Theory (EVT). This process consists of two main steps. We first calculate the strength of the tail dependence between each pair of companies' stock returns. We then use the PMFG (Planar Maximally Filtered Graph) approach to filter and retain the edge information necessary for constructing the network.

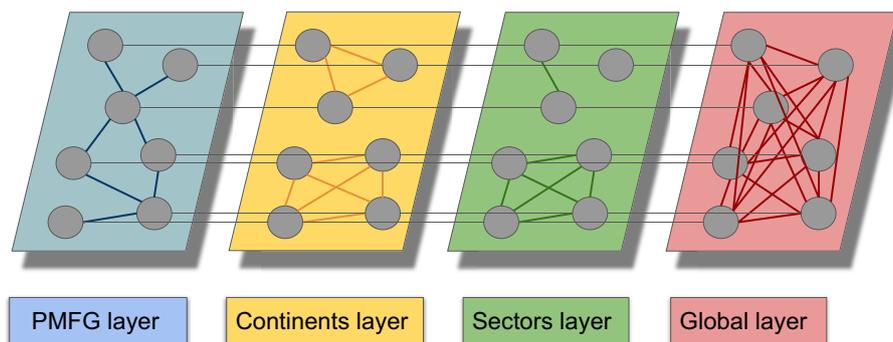


Fig. 1 Schematic diagram of the 4-layer multiplex network described in Sect. 2.2, comprising a weighted PMFG layer and unweighted continents, sectors and global layers. An inter-layer link is present between each two copies of the same node in different layers and between no other nodes

¹ Orbis is the world's largest database with a focus on private company financials, ownership and linkages.

Tail dependence estimation

Let $\{(-x_{i,t}, -x_{j,t}) : t = 2, \dots, N\}$ be the realisations of the bivariate negative stock return (X_i, X_j) , where $x_{i,t}$ is defined as follows:

$$x_{i,t} = \log(p_{i,t}) - \log(p_{i,t-1}). \quad (1)$$

Here $p_{i,t}$ denotes the closing stock price of company i at day t for $1 \leq t \leq 5229$. Defining $S_i = -1/\log F_i(X_i)$ to be the unit Fréchet variable, corresponding to each marginal distribution function F_i of X_i , the upper tail dependence coefficient (upper TDC) $\chi_{i,j}^U$ for each pair (S_i, S_j) and threshold $q \in (0, 1)$ can then be estimated as

$$\hat{\chi}_{i,j}^U = \hat{\chi}_{i,j}^U(q) = \frac{1 - 2q + \hat{C}_{i,j}(q, q)}{1 - q}. \quad (2)$$

In (2), $\hat{C}_{i,j}$ is defined as

$$\begin{aligned} & \hat{C}_{i,j}(q, q) \\ &= \frac{1}{N-1} \sum_{n=1}^{N-1} \mathbb{1}(r_i^n \leq N-1 - \lfloor (N-1)(1-q) \rfloor, r_j^n \leq N-1 - \lfloor (N-1)(1-q) \rfloor), \end{aligned}$$

where r_i^n and r_j^n are the ranks (the index of the element in an ascending list) of the n^{th} observations of S_i and S_j , respectively. In the case of equal ranks, we assign the average rank to all tied values. The coefficient $\chi_{i,j}^U$ ranges from 0 to 1, quantifying the strength of the tail dependence between S_i and S_j . Specifically, $\chi_{i,j}^U = 0$ means that the two variables S_i and S_j are upper tail independent, while $\chi_{i,j}^U > 0$ indicates upper tail dependence. More details can be found in Coles et al. (1999).

The analysis throughout the remainder of the paper is based on the estimated upper tail dependence coefficients $\hat{\chi}_{i,j}^U(0.95)$. This threshold choice aligns with the established literature that employs tail dependence for constructing financial networks (Le et al. 2021; Maghyreh and Abdoh 2020; Singh et al. 2017). Moreover, we construct two networks: the first constructed using all available data before the 2008 financial crisis, which includes the data from 17/01/2002 to 30/06/2007 and the second employing all available data before the 2020 financial crisis, i.e., from 17/01/2002 to 29/02/2020. The selection of the $q = 0.95$ value ensures that there are enough data in the upper tail of the distributions before both the 2008 and the 2020 financial crises. For completeness, in Appendix B we compare the estimated upper TDC values for $q \in \{0.90, 0.95, 0.99\}$. The results indicate that there is limited discrepancy between the upper TDC for $q = 0.90$ and $q = 0.95$. However, when examining the upper TDC at $q = 0.99$ in the pre-2008 dataset, the results indicate a limited number of extreme co-exceedances available in the data, which highlights the limitations of using very high quantiles.

The upper TDC values between each pair of companies i and j are then used to quantify the strength of dependence between the companies in our dataset. These values serve as a basis for constructing multiplex financial networks and implementing the SIR model for financial contagion: a higher TDC between two companies indicates a greater probability of crisis propagation.

Planar Maximally Filtered Graph

The Planar Maximally Filtered Graph (PMFG) method, introduced in Tumminello et al. (2005), is designed to filter complex networks by retaining only the most important

links while maintaining planarity (the property that a graph can be embedded in a plane without any edges crossing) and ensuring network connectedness. This approach assists with eliminating spurious (weak) connections, highlighting key topological properties such as community structures and reducing computational complexity. In the field of finance, PMFGs have been employed to detect market shifts and community structures (Musmeci et al. 2015), to study financial contagion (Pozzi et al. 2013) and to analyse financial networks describing correlations (or other dependencies) between financial assets (Fiedor 2014). Furthermore, as shown in Song et al. (2012), PMFGs are useful for reducing the complexity and dimensionality of financial networks, while preserving the original clustering structure. 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.

- 1: Start with the graph $G = (V, E)$, where V is the set of nodes (representing companies) and the edge set is $E = \emptyset$.
- 2: For each pair (S_i, S_j) ($i < j$) calculate $\hat{\chi}_{i,j}^U(0.95)$ via equation (2) and store the values in a list l , sorted from greatest to smallest.
- 3: Add edge $e_{i,j} = \hat{\chi}_{i,j}^U$ to E , corresponding to the index (i, j) of the first element of l , only if the network G remains planar after addition.
- 4: Remove the evaluated element $\hat{\chi}_{i,j}^U$ from the list l (regardless of whether the edge was added or not) and go back to Step 3.
- 5: Continue the procedure until no more edges can be added.

Algorithm 1 Construct PMFG layer

The resulting connections in the PMFG layer are, contrary to those in the additional layers discussed below, data-driven since they depend on the values $\hat{\chi}_{i,j}^U$.

Additional layers

Additionally to the tail dependence network (hereafter referred to as the PMFG for brevity) layer, we include layers to represent other known relationships between the companies and to capture alternative channels through which crises may propagate.

Sector and continents layers

The 2008 financial crisis highlighted the critical role of interconnectedness as it quickly spread from the US subprime mortgage market to the broader financial sector. This cascade led to significant losses for financial institutions, while the healthcare and technology sectors demonstrated relative resilience. The crisis affected continents unevenly, highlighting differences in vulnerability and recovery. Similarly, the 2020 crisis, triggered by the COVID-19 pandemic, had sector- and continent-specific effects, benefiting the healthcare sector but hurting travel, tourism, retail and airlines industries. To capture these dynamics, we introduce ‘sectors’ and ‘continents’ layers, where companies are connected if they belong to the same sector or continent, respectively.

Global layer

In the literature there are two types of financial contagion: ‘fundamentals-based’ contagion, known as spillover, where the transmission happens due to real and financial connections and ‘pure contagion’, where crises may propagate due to global effects not explicitly accounted for within the connectivity so far defined. To capture this, we add a complete network, referred to as the ‘global layer’. A schematic representation of our network is shown in Fig. 1.

Contagion model

We employ a discrete-time SIR epidemic model defined on the network of $C = 398$ companies to simulate the spread of financial crises. In this model, at any given time step, a company is in one of three states: susceptible (S), infected (I) or recovered (R). Let the functions $S(t), I(t)$ and $R(t) \in \mathbb{Z}_{\geq 0}$ represent the number of companies that are in the states S, I and R , respectively, at time t .

The process begins on day $t = 0$, with $M \geq 1$ initially infected companies ($I(0) = M$) and the remaining companies are susceptible ($S(0) = C - M, R(0) = 0$). At each subsequent day $t = 1, 2, 3, \dots$, an infected company i can infect each susceptible neighbour j on layer α independently with probability $w_{i,j}^{[\alpha]}$. Independently, each infected company recovers with probability γ and cannot be reinfected. Consistent with the multiplex nature of our network, the infection or recovery of a company occurs simultaneously across all layers. The process continues until there are no more infected companies.

The per-day transmission probability for each edge (i, j) in layer α ($\alpha = 1, 2, 3, 4$, representing the PMFG, continents, sectors and the global layer, respectively) is

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

where $\chi_{i,j}^U$ is estimated by $\hat{\chi}_{i,j}^U$ defined in (2), and $\beta_\alpha, 1 \leq \alpha \leq 4$, are parameters to be estimated (see Sect. 3.4). Specifically, $w_{i,j}^{[1]}$ is proportional to $\chi_{i,j}^U$, scaled by a parameter β_1 , such that edges with stronger joint tail dependence (reflecting the likelihood of joint extreme losses beyond a high threshold) correspond to higher probabilities of contagion transmission. For the other layers (continents, sectors, global), transmission probabilities are set using parameters β_α ($\alpha \in \{2, 3, 4\}$), constant across each layer, to capture broader structural influences. The definition of ‘infection’ in a financial context is given in Sect. 3.1.

Application to financial crises

In this section, we fit the model to the 2008 and the 2020 financial crises. Specifically, in Sect. 3.1, we define what is meant for a company in the dataset to be ‘financially infected’. Then, in Sect. 3.2, we describe the ChangeFinder algorithm we use for online CPD. In Sect. 3.3, we explain the process of incorporating change points into the model. Finally, in Sect. 3.4 we examine how the model can be used to predict future infections in both financial crises, using recent infection data.

Infection

We define a company as infected whenever the volatility of its stock returns during a specified period surpasses a predefined threshold (indicating stock price instability) and its average stock return for that period is negative (indicating financial loss). In more detail, the volatility of company i for a time horizon $T > 1$ at day t is

$$V_{i,t} = \sqrt{\frac{1}{T} \sum_{j=t-T}^{t-1} (x_{i,j} - \mu_{i,t})^2}, \tag{4}$$

where $x_{i,j}$ is defined in (1) and $\mu_{i,t}$ is the mean stock return over the same period. Thus, company i is considered infected on day t if $V_{i,t} \geq \sigma_i$ and $\mu_{i,t} < 0$, where σ_i is the 90%

quantile threshold for company i . In the following analysis we set $T = 21$ trading days (one trading month).

Using a rolling window of historical returns over the past 21 days is common in risk analysis (Alexander 2008), particularly for estimating volatility. This method effectively balances the need to capture recent changes in volatility while incorporating enough historical data to provide a stable estimate. Such a balance is crucial for identifying ‘infection’, as longer periods may include stock price fluctuations that no longer influence the market, whereas shorter periods could be more sensitive to noise. The choice of σ_i is determined by examining each company’s volatility curve over the entire available timeframe. This value serves as a reference point for assessing the relative volatility of the company’s current and future volatility values. However, it is important to note that determining this threshold in real-time, without knowledge of future volatility, is impractical. As a result, the quantile threshold is mainly 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 Fig. 2a. The red horizontal line corresponds to a 90% threshold line. We see that with a higher threshold σ_i , the final, post-2020 ‘infection period’ (high volatility period) would not be captured.

Figure 2b plots the number of infected companies per day, showing 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. The other infection peaks also correspond to different financial crises, including the early 2000s recession (in blue), the 2012 Cypriot financial crisis (in green) and the 2015–2016 stock market sell-off (in red).

Change point detection

Change point detection is a technique for identifying the points at which the probability distribution of a time series changes. It has been employed in statistical analysis for various applications, such as trend analysis (Aminikhanghahi and Cook 2017) and anomaly detection (Schmidl et al. 2022). Algorithms for change point detection are usually classified as ‘online’ or ‘offline’. The ‘offline’ algorithms take into account the whole available data series, 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,

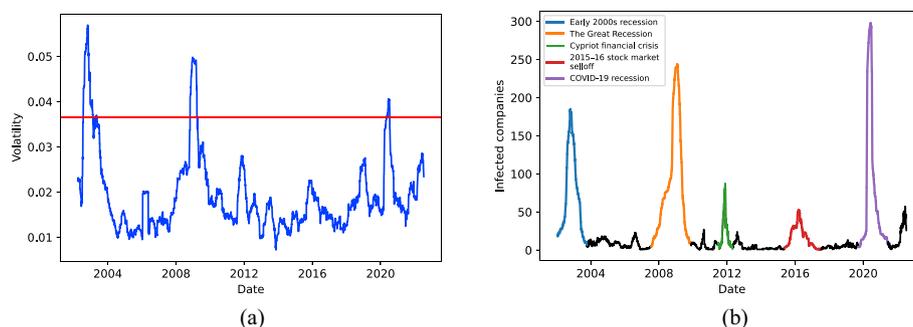


Fig. 2 **a** Volatility curve (in blue) and 90% quantile threshold line (in red) for an example company from the dataset. The volatility at day t for the period from 07/02/2002 to 29/02/2020 is calculated using the formula in (4) with a rolling window of $T = 21$ trading days; **b** Total number of infected companies per day for the period from 07/02/2002 to 29/02/2020. It can be seen that the highest volatility periods in (a) correspond to the highest peaks in (b)

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. However, such algorithms are still useful in practice, as they can be applied incrementally, allowing for timely updates when new data is processed and facilitating near-real-time adjustments.

In our study we utilize ChangeFinder (Takeuchi and Yamanishi 2006), an online change point detection algorithm. The ChangeFinder algorithm is particularly well suited for detecting abrupt shifts or changes in the underlying dynamics of non-stationary time series. This capability is crucial in applications like epidemic modelling, where the transmission rates and recovery rates are likely to evolve over time due to external factors such as interventions, new variants or changes in behaviour. By adapting quickly to these changes, ChangeFinder provides a robust method for real-time anomaly detection, enabling timely responses to unexpected shifts in the spread of an epidemic. Moreover, its ability to continuously update and refine its model of the time series ensures that it can track evolving trends without being overly influenced by outdated data.

ChangeFinder

Since 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 an AR model.

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 (Brockwell and Davis 2002). AR models represent the relationship between an observation and previous observations in a time series by expressing the current value as a linear combination of its past values and some random noise.

Let (z_1, \dots, z_n) be a time series with mean 0. Then, the AR model of order p , denoted by $AR(p)$ can be defined as

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

where α_i are the regression parameters, which indicate the degree with which z_t is correlated with z_{t-i} , and the ϵ_t are independent and identically distributed (iid) random variables generated from a Normal distribution with mean 0 and variance σ^2 . Let $x_{1:n} = (x_1, \dots, x_n)$ be the time series that we observe, where $x_t = z_t + \mu$. That is, the observed data is the $AR(p)$ process shifted by μ , so μ is the mean of the observed data. Let $x_{i:j} = (x_i, \dots, x_j)$ represent the segment of the time series spanning from time i to j , inclusive. Then, assuming the $AR(p)$ model, the conditional probability density function of x_t , given its past, is given by

$$P(x_t | x_{t-p:t-1}, \theta) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{(x_t - \omega_t)^2}{2\sigma^2} \right]. \quad (5)$$

Here, ω_t is given as follows:

$$\omega_t = \sum_{i=1}^p \alpha_i (x_{t-i} - \mu) + \mu, \quad (6)$$

and $\theta = (\alpha_1, \alpha_2, \dots, \alpha_p, \mu, \sigma)$ are the model parameters. Then, if $\hat{\omega}_t$ is the estimated value of ω_t obtained by substituting the estimated model parameters $\hat{\theta} = (\hat{\alpha}_1, \dots, \hat{\alpha}_p, \hat{\mu}, \hat{\sigma})$, estimated via maximum likelihood, into (6), the model-fitting error for $x_{1:n}$ is given by

$$I(x_{1:n}) = \sum_{t=1}^n (x_t - \hat{\omega}_t)^2.$$

For a given time t with $p + 2 \leq t \leq n - p - 1$, the model parameters θ may also be estimated separately via maximum likelihood estimation using the observed sequences $x_{1:t}$ and $x_{t+1:n}$. The model fitting errors $I(x_{1:t})$ and $I(x_{t+1:n})$ are then computed based on these different estimated parameters, reflecting in each case the model's effectiveness in explaining the data prior to, and beyond, time t .

Hence, if $I(x_{1:t}) + I(x_{t+1:n})$ is significantly smaller than $I(x_{1:n})$, 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 . In the above formulation we assume that we have only one change point, with the time series being stationary before and after it. However, this assumption can be overly simplistic and unrealistic. The SDAR model, which we discuss in the following section, deals with the situation in which we have multiple change points (Takeuchi and Yamanishi 2006).

SDAR Model

The SDAR (Sequentially Discounting Auto Regression model learning) algorithm can be used to address the computational problem mentioned above (Takeuchi and Yamanishi 2006). The SDAR algorithm is employed for adapting and discounting data in an online fashion, utilizing the AR model of order p as follows. For $t > p$, for an estimate $\hat{\theta}_t$ of θ given $x_{1:t}$, the SDAR algorithm is employed to calculate in an online manner the value of θ that maximizes the following quantity:

$$\sum_{i=p+1}^t (1-r)^{t-i} \log P(x_i | x_{i-p:i-1}, \theta), \quad (7)$$

where $r \in (0, 1)$ is a discounting factor. A smaller r indicates a stronger impact from past data. As mentioned in Iwata et al. (2019), the SDAR algorithm is useful for real-time learning with non-stationary time series data.

ChangeFinder algorithm

The ChangeFinder algorithm utilizes the SDAR algorithm to detect change points in real time. SDAR-based change point detection has been proven to be efficient in multiple studies (Iwata et al. 2019; Saaid et al. 2012; Wang and Xue 2017). The ChangeFinder algorithm consists of a two-step learning process which combines the detection of outliers and the detection of change points in a time series. An outlier is a point which indicates a momentary abnormal condition (significant increase or decrease), while a change point indicates that the abnormal condition does not return to its original state.

The first learning step of the algorithm uses the SDAR method to calculate the conditional negative log-likelihood of the fitted model at each data point, which shows how much it deviates from the learned model, where a higher value suggests a higher chance of the data point being an outlier. To smooth out short-term fluctuations and reduce the impact of isolated outliers, these quantities are averaged over a rolling window of fixed length T .

The second stage applies the same procedure to this new time series, effectively refining the detection process. By analyzing these refined scores over another sliding window, the algorithm identifies points where the data exhibits a sustained change, signaling a potential shift in the underlying trend. The decision whether a time point is classified as a change point or not is taken by implementing a specific decision rule: in our implementation in Sect. 3.3, we shall decide that a time point is a change point when the corresponding smoothed score is a local maximum in time and exceeds a pre-specified threshold value \mathcal{T} . The detailed ChangeFinder method is described in Algorithm 2.

- 1: Input parameters: order p , discounting rate r and rolling window length T .
- 2: Receive a new data point x_t at time t .
- 3: For each $t \geq p+1$, the conditional negative log-likelihood $S(x_t)$ at x_t is calculated,

$$S(x_t) = -\log P(x_t | \mathbf{x}_{t-p:t-1}, \hat{\theta}_{t-1}),$$

where the values of $\hat{\theta}_{t-1}$ are estimated using the SDAR method with discounting rate r , maximizing equation (7), with P given in (5). Here, a higher $S(x_t)$ indicates that x_t is more likely to be an outlier.

- 4: For each t , a rolling mean of these log-likelihood values 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 $t = T+1, T+2, \dots$ is computed:

$$y_t = \frac{1}{T} \sum_{i=t-T+1}^t S(x_i). \quad (8)$$

This process is needed in order to reduce the influence of isolated outliers contained in the time series (x_1, \dots, x_t) .

- 5: For each t , another conditional negative log-likelihood $S(y_t)$ at y_t is calculated,

$$S(y_t) = -\log P(y_t | \mathbf{y}_{t-p:t-1}, \tilde{\theta}_{t-1}),$$

where $\tilde{\theta}_{t-1}$ is estimated using the SDAR method with discounting rate r , maximizing equation (7), with P given in (5).

- 6: A smoothing step is applied to get a sequence of moving averages a_t for $t = T+1, T+2, \dots$:

$$a_t = \frac{1}{T} \sum_{i=t-T+1}^t S(y_i). \quad (9)$$

Here, a_t is the change point or anomaly score at time t . A higher value of a_t indicates that t is more likely to be a change point.

Algorithm 2 ChangeFinder algorithm (Takeuchi and Yamanishi 2006)

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. We first show the effectiveness of the ChangeFinder algorithm in detecting change points in the 2008 and 2020 financial crises in an online manner. We then propose an algorithm for incorporating change points into the model for adaptive parameter estimation. We finally demonstrate the effectiveness of change-point incorporation into the model for predicting future infections.

Real-time change detection in financial crises with ChangeFinder

Let $I(t)$ denote the total number of infected companies on day t in each of the two crises, estimated using the procedure in 3.1. To detect significant changes in the probability of infection and incorporate them into the model, we examine the time series $\{I(1), I(2), \dots, I(L)\}$, where L is the length of the studied crisis in days. We emphasize that the ChangeFinder algorithm does not use future data when detecting change points. At each time step $1 \leq t \leq L$, the algorithm receives a new data point $I(t)$ and updates its calculations based only on the observations available up to that point, ensuring that change points are identified in an online, sequential manner. In other words, each anomaly score a_i is calculated based only on the time series $\{I(1), I(2), \dots, I(i)\}$ and does not use the full time series $\{I(1), I(2), \dots, I(L)\}$.

We fit the ChangeFinder algorithm with order $p = 1$, discounting rate $r = 0.5$ and rolling window length $T = 7$ for the smoothing steps to both the 2008 and the 2020 financial crises. These parameter choices were made to strike a balance between computational efficiency and sensitivity to abrupt changes. In particular, the parameter $p = 1$ is chosen. Firstly, because it keeps the computational burden low. Secondly, with $p = 1$, the model adapts quickly to sudden changes: if a sudden spike or drop in the time series happens, the ChangeFinder algorithm can detect it with minimal latency. Choosing $T = 7$ strikes a balance between time delay and outlier filtering. 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 a time delay for detecting change points, but outliers are filtered. Finally, a moderate r value balances between sensitivity to changes and noise in the results. For completeness, Figs. 10 and 11 in Appendix B present the anomaly scores for different combinations of the ChangeFinder 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. It can be seen that the smaller the window size T and the larger the discount rate r , the noisier the obtained anomaly scores. The threshold \mathcal{T} for detecting change points is set to the 90% quantile of the entire anomaly score time series, ensuring that only the most extreme deviations are highlighted. However, as discussed in Sect. 4, the choice of this threshold remains an open problem. Here, the 90% level is selected primarily to demonstrate the proposed approach rather than to claim it as the optimal choice.

The results for the 2008 and 2020 financial crises, using the chosen parameters, are shown in Figs. 3a and b. In these figures, anomaly scores and change points are plotted sequentially, reflecting the real-time nature of the algorithm. The 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

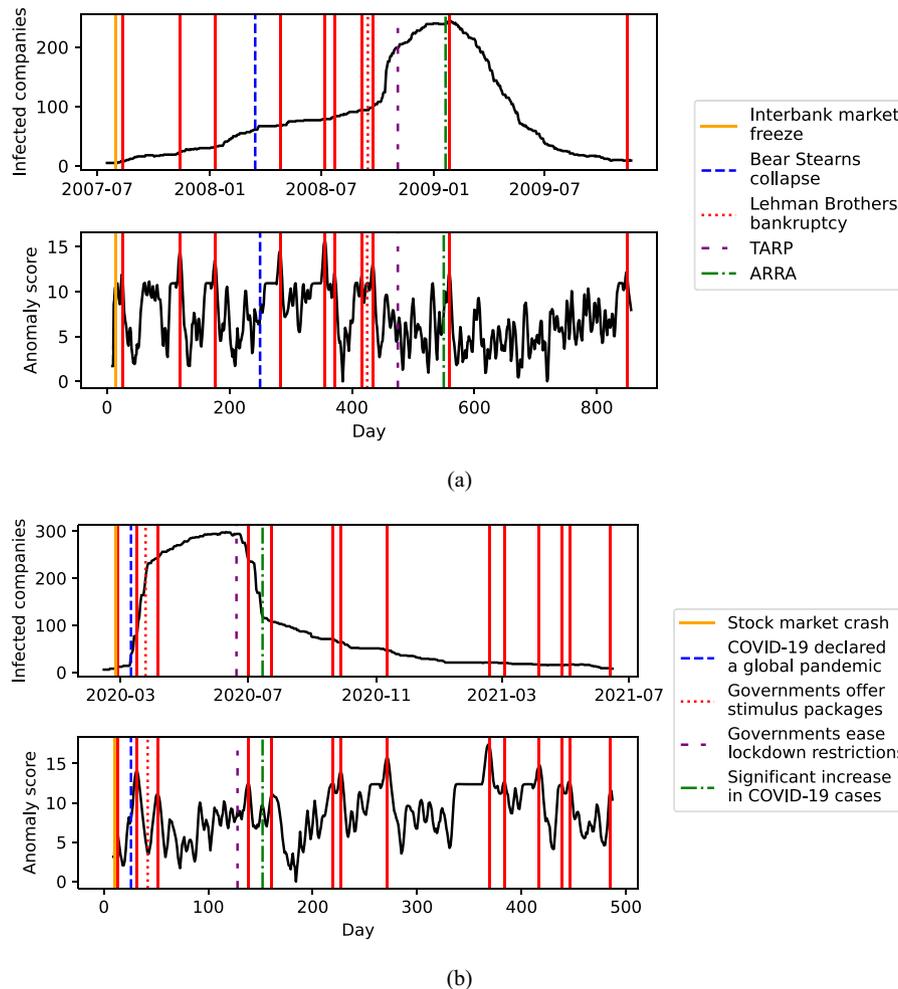


Fig. 3 Change point detection on the total number of infected companies in the **a** 2008 and **b** 2020 financial crises, respectively. The anomaly scores and change points are plotted sequentially as they are detected, reflecting the real-time nature of the algorithm. For each crisis the bottom plot illustrates the anomaly scores, as defined in (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

the bottom plot in each figure. The other coloured vertical lines correspond to significant events that happened during each crisis, marking key moments that had a substantial impact on the financial system.

During the 2008 financial crisis, some of the most notable events include the freeze in the interbank lending market, which caused liquidity shortages and heightened financial instability. This was followed by the collapse of Bear Stearns in March 2008, a key investment bank that had significant repercussions for global financial markets. The crisis deepened with the bankruptcy of Lehman Brothers in September 2008, which led to a sharp deterioration in market confidence and a global economic downturn. In response to the financial turmoil, the U.S. government implemented several rescue measures, such as the Troubled Asset Relief Program (TARP), which was launched to stabilize the banking sector by purchasing distressed assets. Additionally, the American Recovery and Reinvestment Act (ARRA) of 2009 aimed to stimulate the economy through fiscal spending and tax cuts, as well as investments in infrastructure and energy efficiency, seeking to mitigate the effects of the recession.

In the 2020 crisis, the World Health Organization (WHO) declared a global health emergency in March 2020. This was followed by national lockdowns in numerous countries, which significantly impacted businesses. As a result, governments worldwide began implementing stimulus packages to mitigate the economic damage. The introduction of financial support measures helped stabilize markets but did not immediately curb the health crisis. Subsequently, between June and July 2020, many countries began easing their lockdown restrictions, which allowed for some economic recovery. However, this recovery was short-lived, as a resurgence of COVID-19 cases occurred about a month later, signaling the start of a second wave of infections globally.

In both figures it can be seen that most of the time shortly after the mentioned 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.

Incorporating change points for adaptive parameter estimation

In this section we propose a method to adjust the model's parameters dynamically following the detection of a change point. We first construct two separate networks, as detailed in Sect. 2.2, representing the connections between the companies in the 2008 and 2020 financial crises. 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. For a detailed comparison of the two networks the reader is referred to Bozhidarova et al. (2024).

To make predictions of the number of infected companies in the future k crisis days, given infection data from the past m days we utilize a 'sliding window' technique. The procedure starts from the first day of the crisis, setting $t = 1$ (data window 1) and we set the initial change point to be $t_c = t$ (the first day of the data window). Then, for each $i \in \{t, \dots, t + m\}$ we calculate the anomaly score a_i , using the ChangeFinder algorithm, applied on the time series $\{I(1), I(2), \dots, I(i)\}$. In other words, each anomaly score a_i is calculated based on the time series available up to time i and does not use future data. If a_i is a local maximum, meaning $a_i > a_{i-1}$ and $a_i > a_{i+1}$, within the obtained anomaly scores time series $\{a_t, \dots, a_{t+m}\}$ and $a_i > \mathcal{T}$ (the anomaly score is over the threshold \mathcal{T}), then we classify i as a change point and set $t_c = i$. We disregard all data points before the change point and fit the model to the period from day t_c to day $t + m$ (data window t) obtaining maximum likelihood estimates $\hat{\beta}_j$ of the layer transition probabilities β_j for $1 \leq j \leq 4$ and $\hat{\gamma}$ for the recovery rate γ (we refer to Appendix A for further details). In other words, instead of fitting the model to all the previous m crisis days, if a change point is detected during the past m days of a crisis, we discard the data before the change point has happened and fit the model exclusively to the days following (and including) the change point. If more than one change point is detected during the past m days, we discard all the data prior to the most recent change point. The step-by-step procedure is described in Algorithm 3.

Next, we simulate $N_{\text{sim}} = 10,000$ realisations of the estimated SIR model for the upcoming k days (from day $m + 1$ to day $m + k$, denoted prediction window 1), with the initial data being that from day m . We then use the results from the simulations as predictions (see step 5 of Algorithm 3 for more details). We finally 'slide the window' forward by one day and refit the model to the period from day 2 to day $m + 1$ of the crisis

(data window 2), re-estimating $\hat{\beta}_j$ for $1 \leq j \leq 4$ and $\hat{\gamma}$ for the new window. We repeat the above steps for each subsequent data window, with the final prediction window covering the period from day $L - k$ to L , where L is the length of the crisis in days.

In the following section, we evaluate the model that incorporates change points (we henceforth denote this as the change point model) and compare its performance with that of the model from Bozhidarova et al. (2024), without incorporating change point detection.

- 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 $i \in \{t, \dots, t + m\}$, estimate $\hat{\beta}_j$ ($1 \leq j \leq 4$) (see Step 3(a) or Step 3(b) based on model choice):
 - (a) **Model without CPD:** Fit the model to the period from day t to day $t + m$ and estimate $\hat{\beta}_j$ ($1 \leq j \leq 4$) and $\hat{\gamma}$.
 - (b) **Model with CPD:** For $i \in \{t, \dots, t + m\}$ calculate the anomaly score a_i . If a_i is a local maximum within the collected time series of anomaly scores $\{a_t, \dots, a_{t+m}\}$ and $a_i > \mathcal{T}$ (the anomaly score is over the threshold \mathcal{T}), then i is classified as a change point. Set $t_c = i$. If more than one change point is detected, set t_c as the most recent change point detected. Fit the model to the period from day t_c (disregard all data points before the change point) to day $t + m$ and estimate $\hat{\beta}_j$ ($1 \leq j \leq 4$) and $\hat{\gamma}$.
- 4: Run $N_{\text{sim}} = 10,000$ SIR simulations on the network from day $t + m + 1$ to day $t + m + k$ (i.e., predict infections and recoveries in the next k days) using the values of $\hat{\beta}_j$ ($1 \leq j \leq 4$) and $\hat{\gamma}$ estimated in Step 4. To 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 over all N_{sim} simulations as a prediction for the number of newly infected (recovered, respectively) companies in the period from day $t + m + 1$ to day $t + m + k$ (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_{sim} simulations as a prediction for the total number of infected companies on the final day of the prediction window t .
- 5: Update $t = t + 1$.
- 6: If $t > L - k - m$ stop. Else go back to Step 2.

Algorithm 3 Model predictions for the future k days

Evaluating the change point model: predicting future infections

In this section we assess the accuracy of the change point model in predicting the number of infected companies in the future k crisis days, given data from the past m days of each crisis. The results show that incorporating change points improves predictive accuracy, particularly when estimating the number of infected companies. Unlike the baseline model which performs poorly during and after significant events, our enhanced approach effectively adapts to these critical moments, yielding better predictions.

We first compare in Fig. 4 how the change point model performs, in comparison with the model without incorporating the change point algorithm (Algorithm 3 with step 3a), in predicting the total number of infected companies (top row), number of newly infected companies (middle row) and the number of newly recovered companies (bottom row) in the future $k = 30$ days for different values of m . The outcomes are derived using Algorithm 3. The results reveal that the model incorporating the change point detection algorithm (shown in orange) consistently outperforms the model without the

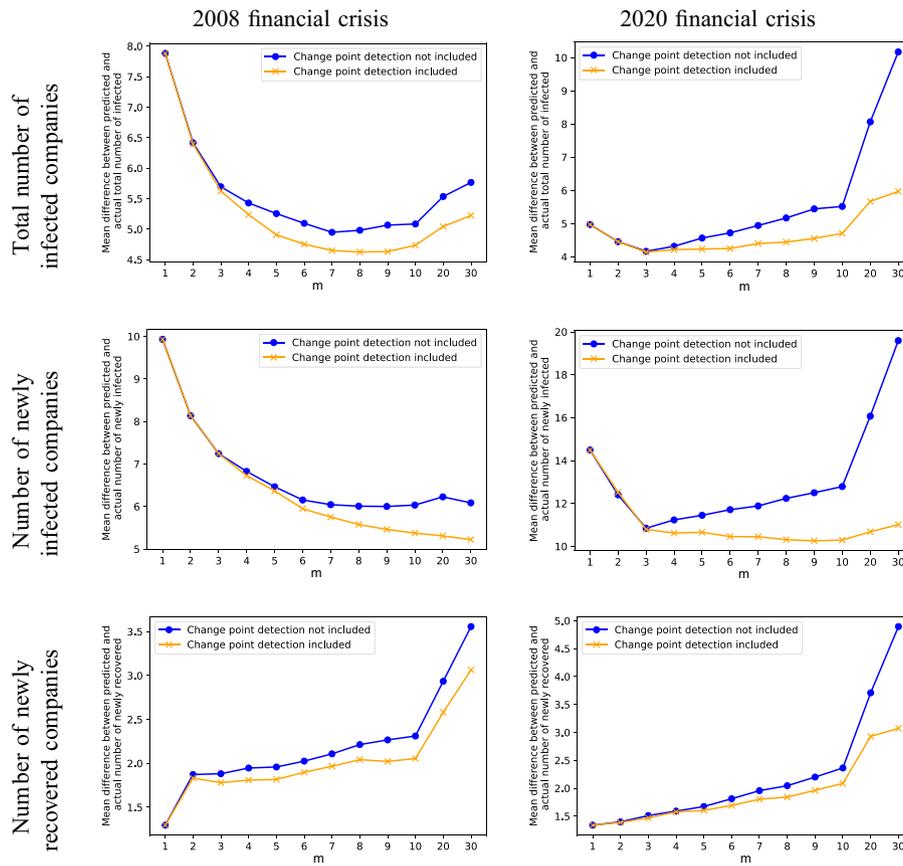


Fig. 4 Comparison between the model from Bozhidarova et al. (2024) and the change point detection model of the absolute difference after $k = 30$ days between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous m days

change points (shown in blue), with higher improvement for higher values of m . For completeness, plots for $k = 10$ and $k = 20$ can be found as Figs. 8 and 9 in Appendix B and show that the results are similar for all choices of k .

While the previous analysis focused on comparing the mean predictions of the models in each window, we now evaluate the models by considering the full distribution of simulated trajectories. To this end, we employ the Mahalanobis distance (Dodge 2008) to measure how closely the simulations capture the real observed dynamics. That is, for each sliding window we generate $N_{\text{sim}} = 10,000$ simulations for the future $k = 30$ days, and let $\mu_{\text{sim}} \in \mathbb{R}^k$ and $\Sigma_{\text{sim}} \in \mathbb{R}^{k \times k}$ represent, respectively, the k -dimensional mean column vector and $k \times k$ covariance matrix of these simulations. Let also $x_{\text{real}} \in \mathbb{R}^k$ denote the column vector of real observed numbers of infected companies in a given window. The Mahalanobis distance between the real trajectory and the simulated distribution is then computed as

$$D_M(x_{\text{real}}, \mu_{\text{sim}}) = \sqrt{(x_{\text{real}} - \mu_{\text{sim}})^\top \Sigma_{\text{sim}}^{-1} (x_{\text{real}} - \mu_{\text{sim}})}.$$

This distance quantifies how far the real trajectory lies from the centre of the simulated distribution while accounting for the variability and correlations of the model output. A

smaller D_M indicates that the observed dynamics are well captured by the simulations, whereas a larger distance signals a poor fit.

Evaluating D_M across all windows provides a rigorous, distribution-aware measure of model accuracy. Figure 5 shows that the model with change point detection consistently achieves lower Mahalanobis distances than the baseline model without change points. This confirms that incorporating change points enhances the model's ability to capture abrupt shifts in crisis dynamics and to align more closely with the real infection trajectories.

In addition, Fig. 6 illustrates the model predictions (shown as coloured lines) alongside the observed infection data (represented by the black curves) at each prediction window. The predictions are derived from the average of all $N_{\text{sim}} = 10,000$ simulations for the upcoming $k = 30$ days of each crisis, using infection data from the previous $m = 30$ days. In the left column, illustrating the predictions of the model from Bozhidarova et al. (2024), which does not incorporate CPD, it can be seen that significant errors occur shortly after major events, denoted by the vertical lines. For instance, in the 2008 crisis, substantial recovery prediction errors were observed during periods linked with the Lehman Brothers bankruptcy, TARP and ARRA. In the 2020 crisis, notable errors were noted shortly after times of stimulus packages and lockdown restrictions. Comparing the outcomes of the two models, we observe a clear improvement in prediction accuracy when CPD is incorporated, particularly noticeable in the periods following significant events, as indicated by the coloured rectangles in Fig. 6. In contrast, the model without CPD struggles to adapt to these sudden changes, leading to larger deviations between predicted and actual infections. These results illustrate the importance of incorporating CPD, as it allows the model to identify key turning points in the data and adjust its predictions accordingly.

In our analysis, we focus on a direct comparison between our proposed model which incorporates change point detection and a baseline version of the model without change point detection. This setup allows us to isolate the effect of incorporating structural breaks on predictive performance. For completeness, we also examined a classical integer-valued generalized autoregressive conditional heteroskedasticity (INGARCH) model in Appendix C. INGARCH models are well suited for discrete count data and capture short-term autocorrelation, but they do not account for abrupt structural changes and do not incorporate the network structure of company interactions. Because of these

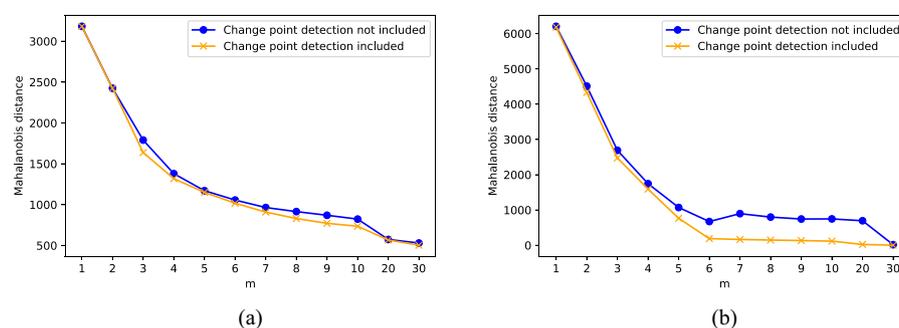


Fig. 5 Mahalanobis distances between the simulated infection trajectories and the observed data for **a** the 2008 crisis and **b** the 2020 crisis, evaluated for different window lengths m . The orange line corresponds to the model incorporating change point detection, while the blue line represents the baseline model without change points. Lower Mahalanobis distances indicate that the real trajectory lies closer to the distribution of simulated trajectories

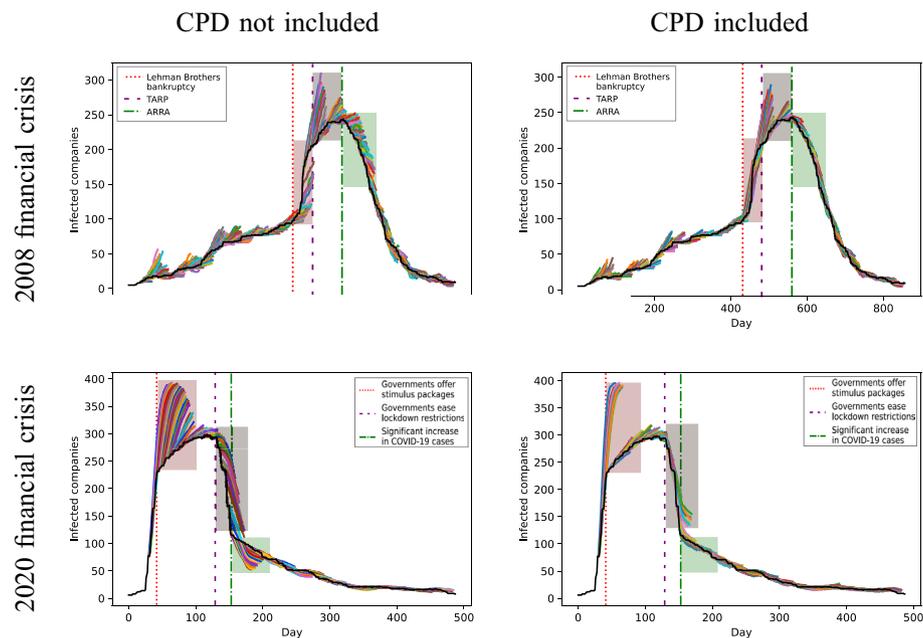


Fig. 6 Comparison between the model predictions from Bozhidarova et al. (2024) (left column) and the CPD model (right column). The colored curves represent the predicted average total number of infected companies for each sliding window over the next $k = 30$ days, based on $N_{\text{sim}} = 10,000$ simulation runs. The black lines indicate the actual number of infected companies, as defined in Sect. 3.1, serving as a benchmark for comparing the model's predictions. The coloured vertical lines indicate important events during each crisis and the coloured rectangles indicate the periods shortly after these events, where the improvements of the CPD model are prominent

limitations, INGARCH is included only as an auxiliary benchmark. As shown in Appendix C, INGARCH exhibits limited ability to capture the dynamics of sudden changes in financial crises, confirming that our network-based change point model is better suited to modelling infection propagation in financial networks.

Discussion and conclusion

In this paper, we incorporated the ChangeFinder algorithm (Takeuchi and Yamanishi 2006), into the model introduced in Bozhidarova et al. (2024) for change point detection in real time. 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. We investigated the effectiveness of our approach by applying it to two recent financial crises: the 2008 crisis, triggered by the subprime mortgage market collapse, and the 2020 crisis, driven by the COVID-19 pandemic. In each case, we analyzed how well our model estimates future infection risk over a horizon of k days, using data from the preceding m crisis days. Through a comparative analysis, we demonstrated the enhanced predictive performance of the change point model in forecasting the number of infected/recovered companies. The results indicate a consistent improvement over the model without change point detection, highlighting the benefits of integrating change points for a more robust and adaptive predictive model.

Although our model offers valuable predictive insights into how financial crises spread, it is crucial to acknowledge its limitations. As part of our modelling framework, we construct two distinct networks by measuring the tail dependence between stock returns of company pairs, identifying extreme events or shocks using a uniform threshold

($q = 0.95$) applied consistently across all companies. While this provides a consistent benchmark, it may not fully capture company-specific variability in risk profiles or volatility. We acknowledge that such an approach could overlook idiosyncratic extremes for companies with highly skewed or low-variance distributions. As a promising direction for future research, we propose incorporating dynamic or company-specific thresholds, which adapt to each company's historical distribution or evolving market conditions. It would, likewise, be important to validate the stability of tail dependence coefficients over time, since this is used in the construction of the PMFG layers in our model, although we note that since these coefficients are estimated using the data collected before the financial crises happen, non-stationarity due to financial stress should not be a concern. Violations of the stability property could be identified by, for example, considering the stationarity of parametric benchmarks such as variances and covariances.

Our current multiplex network framework assumes independence between layers. Extending the model to incorporate inter-layer dependencies, such as through tensor-based approaches (Avdjiev et al. 2019) or vector autoregressive (VAR) network models (Ahelegbey and Giudici 2022) would allow for richer representations of cross-sectoral and cross-regional interactions. We emphasize that our current model serves as a foundational benchmark that demonstrates the benefits of integrating online change point detection within multiplex contagion modelling. A clear direction for future research is to build upon this foundation by incorporating inter-layer dependencies and comparing the performance of such enhanced models with the existing framework.

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 3) could impact the algorithm's sensitivity to change points. Moreover, choosing the optimal parameter values depending on the dataset remains an important problem, providing a clear direction for future research. In addition, the ChangeFinder model framework could be extended by incorporating the ARIMA (AutoRegressive Integrated Moving Average) model or a broader model, like the state-space model (Monteiro and Costa 2023), which could be beneficial for more accurate change point detection.

The exponential discounting feature of the SDAR algorithm (Equation (7)) enhances responsiveness but, equally, makes it ill-suited for capturing seasonal patterns that may be present in financial time series, potentially limiting their effectiveness in detecting subtle or complex structural changes. Moreover, SDAR does not explicitly model volatility clustering, a common feature in financial markets. Alternative methods, such as GARCH-based (Berkes et al. 2004) or Bayesian (Adams and MacKay 2007) CPD approaches, could address these challenges more effectively by explicitly accounting for heteroskedasticity and incorporating prior information, at the price of a further modelling and computational effort. Future research could explore these approaches to enhance detection accuracy, particularly in datasets exhibiting highly volatile dynamics or nonlinear behavior.

Another drawback of the model is that it disregards 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 explore this, we applied the CPD model to make continent-, sector-, and company-specific predictions, which were not presented in

this paper. Although improvements in these cases were less pronounced, they highlight opportunities for further optimization. A detailed explanation of the accuracy metrics employed to evaluate these predictions can be found in Bozhidarova et al. (2024). Future work could focus on refining these methods to achieve more substantial improvements in these specific contexts.

To address these limitations, a natural direction for future research is introducing 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 predictive accuracy. Moreover, in this paper 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.

In addition, our current framework could be extended by incorporating Markov switching models (Guidolin 2011) or hidden Markov models (Cappé et al. 2005). These approaches allow for capturing regime changes and unobserved states within the system, providing a more dynamic and flexible representation of financial contagion.

Appendix A Parameter estimation

Once we have determined which companies are infected on each day using the procedure in Sect. 3.1, we estimate the model’s parameters by maximizing the likelihood function of the SIR model.

Transmission probability

Let us denote the set of infected neighbours in layer α of a susceptible company j at time t as $I_j^\alpha(t)$. We know that the probability that company $i \in I_j^\alpha(t)$ does not infect company j through layer α at time $t + 1$ is $1 - w_{i,j}^{[\alpha]}$ (see (3)). Therefore, the probability that company j does not become infected at time $t + 1$ is the probability that none of its infected neighbours in any of the layers $1 \leq \alpha \leq 4$ at time t infects j at time $t + 1$:

$$\begin{aligned}
 &P(j \text{ does not get infected at time } t + 1) \\
 &= \prod_{\alpha=1}^4 \prod_{i \in I_j^\alpha(t)} P(i \text{ does not infect } j \text{ on layer } \alpha \text{ at time } t + 1) = \prod_{\alpha=1}^4 \prod_{i \in I_j^\alpha(t)} (1 - w_{i,j}^{[\alpha]}). \tag{A1}
 \end{aligned}$$

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

$$P(j \text{ gets infected at time } t + 1) = 1 - \prod_{\alpha=1}^4 \prod_{i \in I_j^\alpha(t)} (1 - w_{i,j}^{[\alpha]}). \tag{A2}$$

Let S_t be the set of susceptible companies at time t . Let SI_{t+1} denote the subset of companies which get infected at time $t + 1$, so that $S_{t+1} = S_t \setminus SI_{t+1}$. Therefore the likelihood for the observed infections from time t to time $t + 1$ is:

$$\begin{aligned} \mathcal{L}_{t,t+1} &= \prod_{j \in \text{SI}_{t+1}} \text{P}(j \text{ gets infected at time } t + 1) \\ &\times \prod_{j \in \text{S}_{t+1}} \text{P}(j \text{ does not get infected at time } t + 1). \end{aligned} \tag{A3}$$

Substituting (A1) and (A2) into (A3) we get:

$$\mathcal{L}_{t,t+1} = \left\{ \prod_{j \in \text{SI}_{t+1}} \left(1 - \prod_{\alpha=1}^4 \prod_{i \in \text{I}_\alpha^\sigma(t)} (1 - w_{i,j}^{[\alpha]}) \right) \right\} \left\{ \prod_{j \in \text{S}_{t+1}} \prod_{\alpha=1}^4 \prod_{i \in \text{I}_\alpha^\sigma(t)} (1 - w_{i,j}^{[\alpha]}) \right\}. \tag{A4}$$

Suppose we want to estimate the parameters $\beta_i, 1 \leq i \leq 4$ for a period of N days from day d to day $d + N - 1$. By substituting $w_{i,j}^{[\alpha]}$ with its representation in terms of $\beta_i, 1 \leq i \leq 4$, shown in (3), and writing the likelihood of the observations as

$$\mathcal{L}_{d,d+N-1} = \prod_{t=d}^{d+N-1} \mathcal{L}_{t,t+1}, \tag{A5}$$

where $\mathcal{L}_{t,t+1} = \mathcal{L}_{t,t+1}(\beta_1, \beta_2, \beta_3, \beta_4)$ is defined in (A4), a maximization of (A5) numerically leads to the maximum likelihood estimate $(\hat{\beta}_1, \hat{\beta}_2, \hat{\beta}_3, \hat{\beta}_4)$ of $(\beta_1, \beta_2, \beta_3, \beta_4)$.

Recovery probability

Suppose that each company recovers on day t with probability γ . 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. Therefore the likelihood for the recoveries is:

$$\mathcal{L}_{t,t+1}^\gamma = (1 - \gamma)^{|\text{II}_{t+1}|} \times \gamma^{|\text{IR}_{t+1}|}. \tag{A6}$$

Hence, the recovery likelihood from day N to day $d + N - 1$ is:

$$\mathcal{L}_{N,d+N-1}^\gamma = \prod_{t=N}^{d+N-1} \mathcal{L}_{t,t+1}^\gamma = (1 - \gamma)^{\sum_{t=N}^{d+N-1} |\text{II}_{t+1}|} \times \gamma^{\sum_{t=N}^{d+N-1} |\text{IR}_{t+1}|}. \tag{A7}$$

The maximum likelihood estimate $\hat{\gamma}$, namely the value of γ which maximises (A7), is then:

$$\hat{\gamma} = \frac{\sum_{t=N}^{d+N-1} |\text{IR}_{t+1}|}{\sum_{t=N}^{d+N-1} (|\text{II}_{t+1}| + |\text{IR}_{t+1}|)} = \frac{\sum_{t=N}^{d+N-1} |\text{IR}_{t+1}|}{\sum_{t=N}^{d+N-1} |\text{I}_t|},$$

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

Appendix B Additional figures

Here we present additional figures related to the outcomes of our modelling approach. Figures 7a and b provide a comparison of the upper tail dependence values for different values of $q \in \{0.90, 0.95, 0.99\}$ for the two periods before the 2008 and 2020 financial crises, respectively. In both figures the results appear to be inconsistent between $q = 0.90$ and $q = 0.99$. A possible explanation is that using a lower threshold such as $q = 0.90$ may include values that are not truly ‘extreme’, mixing tail and non-tail behaviour. Besides, in the pre-2008 dataset, when examining the tail dependence at $q = 0.99$, the estimates often exhibit a discrete or banded structure, with repeated values appearing across many company pairs. This pattern reflects the limited number of extreme co-exceedances available in the data. In the absence of a priori theoretical justification for a specific value of q , the value $q = 0.95$ appears to strike an acceptable middle ground between $q = 0.90$ and $q = 0.99$, with the results between $q = 0.90$ and $q = 0.95$, on the one hand, and $q = 0.95$ and $q = 0.99$, on the other hand, being generally consistent. This sensitivity analysis highlights the limitations of empirical tail dependence estimation at very high quantiles and the difficulty of choosing a reasonable threshold level.

Figures 8 and 9 provide a comparison of the performance of the change point model and the model from Bozhidarova et al. (2024) in predicting the total number of infected companies (top row), number of newly infected companies (middle row) and the number of newly recovered companies (bottom row) in the future $k = 10$ and $k = 20$ days, respectively, for different values of m . It can be seen that the change point model yields better predictions, with higher improvement for higher values of m .

Figures 10 and 11 illustrate 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 illus-

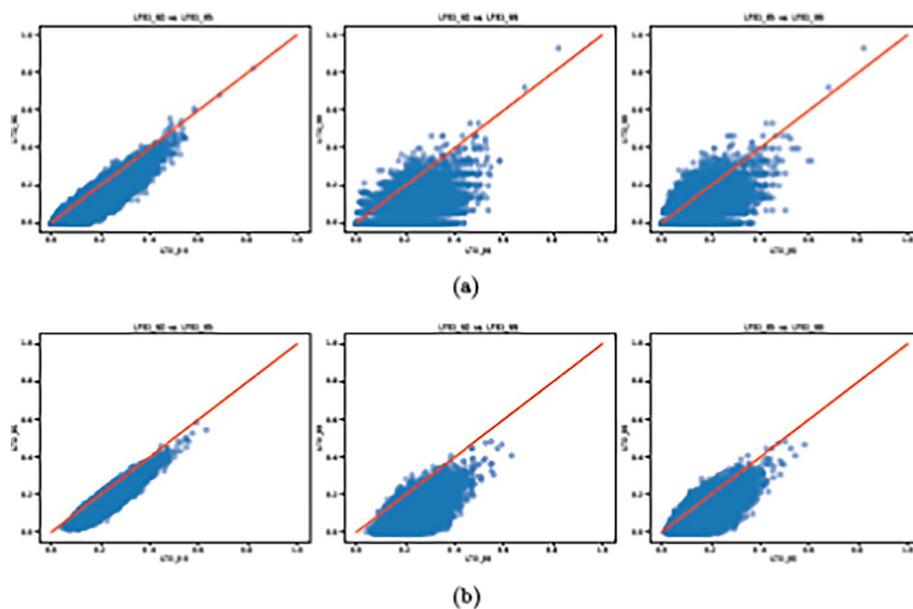


Fig. 7 Scatterplots comparing pairwise upper tail dependence estimates between company pairs at different quantile thresholds: $q = 0.90$ vs $q = 0.95$, $q = 0.90$ vs $q = 0.99$ and $q = 0.95$ vs $q = 0.99$ from left to right for a) 2008 and b) 2020 financial crisis. Each point represents a unique company pair. The red dashed line indicates the 45-degree reference line (perfect agreement). Points near this line suggest stable tail dependence across quantiles, while deviations indicate sensitivity of tail dependence to the depth of the tail (i.e., more extreme quantiles)

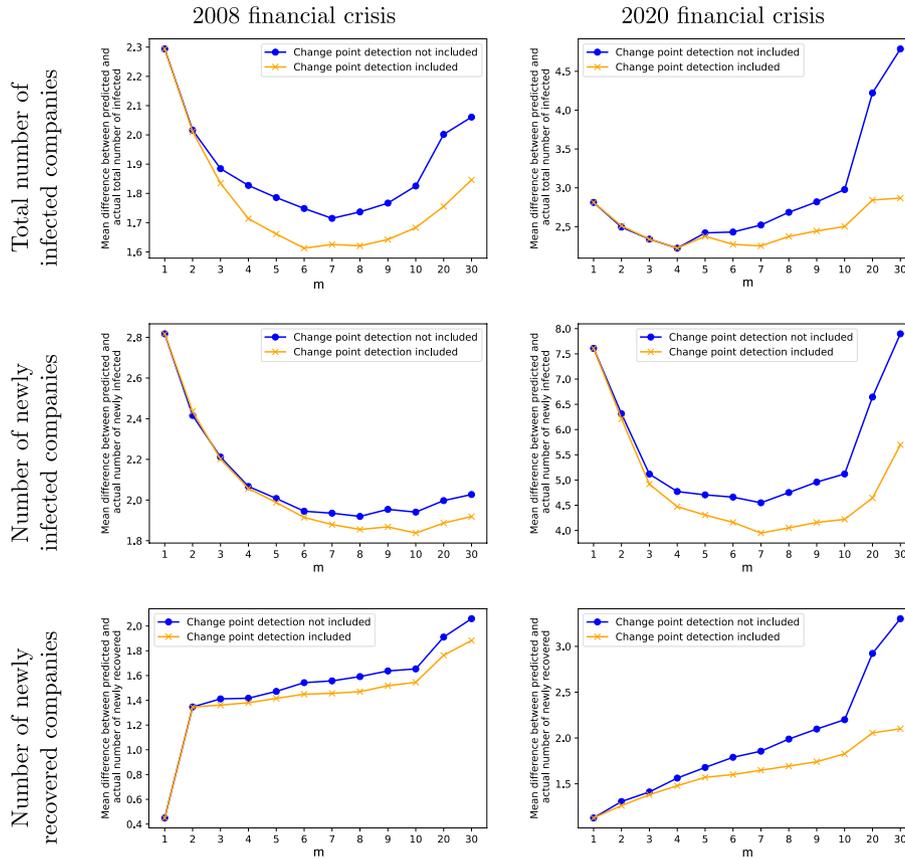


Fig. 8 Comparison between the model from Bozhidarova et al. (2024) and the change point detection model of the absolute difference after $k = 10$ days between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous m days

trate 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 it reduces the noise in the anomaly scores while still allowing for the detection of critical events.

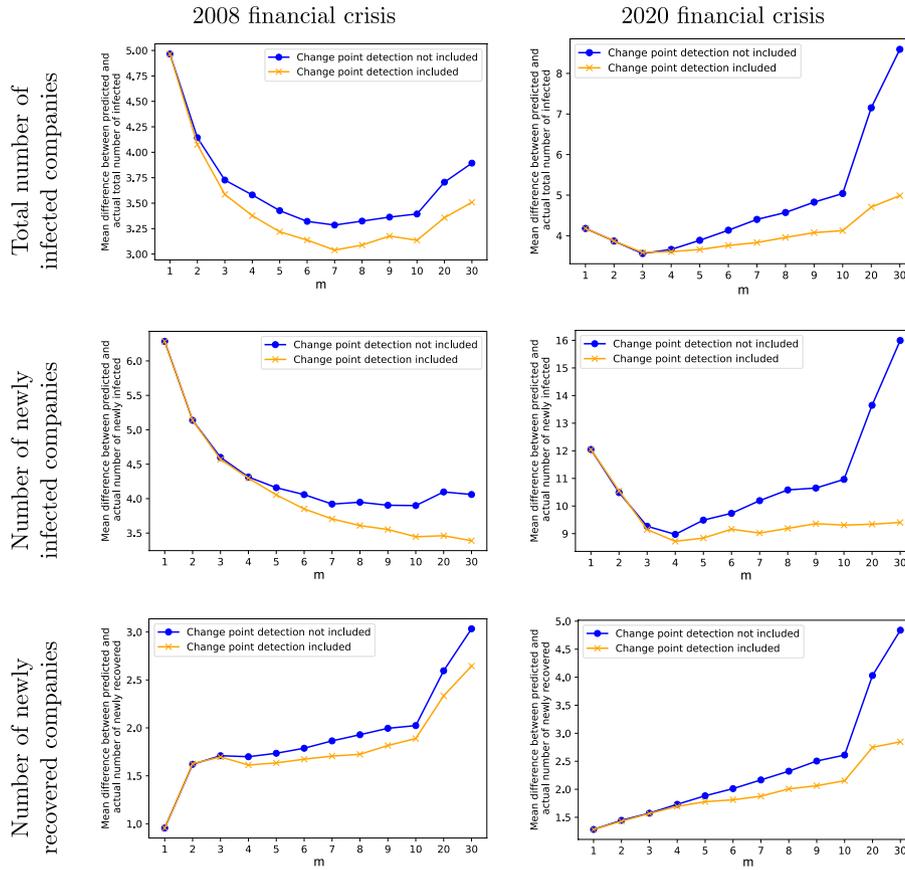


Fig. 9 Comparison between the model from Bozhidarova et al. (2024) and the change point detection model of the absolute difference after $k = 20$ days between predicted and actual total number of infected companies (top row), number of newly infected companies (middle row) and number of newly recovered companies (bottom row) for the 2008 (left column) and the 2020 (right column) financial crises, using the infections data from the previous m days

Appendix C Comparison with INGARCH model

For completeness, we compare our model to an integer-valued generalized autoregressive conditional heteroskedasticity (INGARCH) model. The Poisson INGARCH(p,q) process assumes that an integer-valued random variable X_t , conditioned on X_{t-1}, X_{t-2}, \dots , is Poisson distributed with mean M_t , where M_t satisfies

$$M_t = \omega + \sum_{i=1}^p \alpha_i X_{t-i} + \sum_{j=1}^q \beta_j M_{t-j}.$$

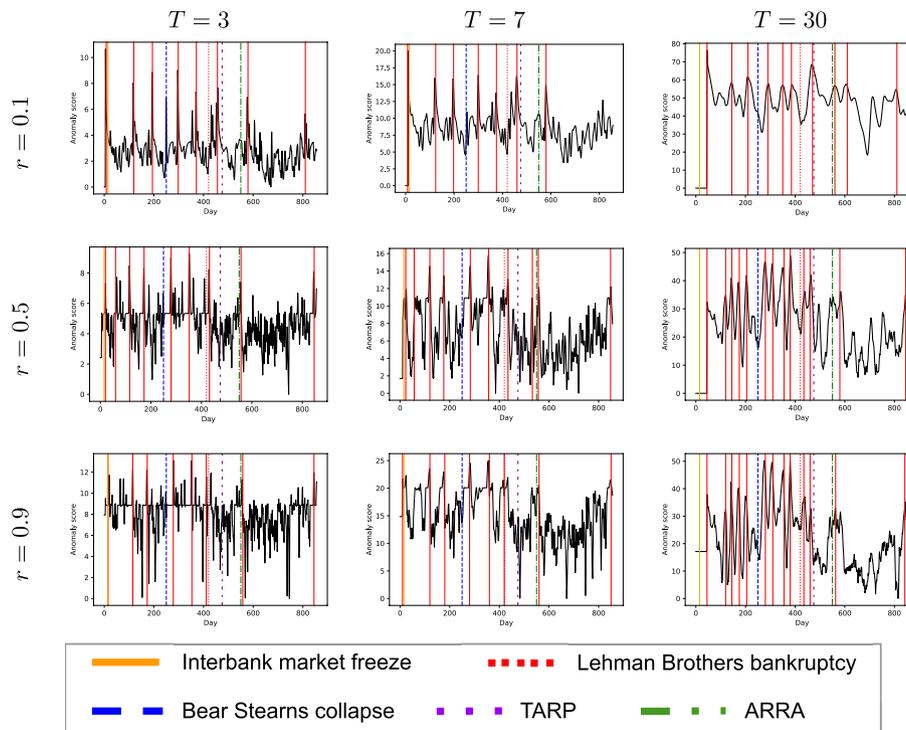


Fig. 10 The figures illustrate the anomaly scores output, using the ChangeFinder algorithm with different combinations of the discounting rate r and the rolling window size T . The ChangeFinder algorithm is run on the 2008 daily infections data. The solid red vertical lines illustrate points with high probability of being a change point, whereas the other vertical lines correspond to the significant events that happened during the crisis

See e.g. equation (4.6) in Davis and Liu (2016).

We evaluated the INGARCH($p = m, 1$) model, viewing X_t as the number of infected companies at time t , using sliding windows of various lengths m to forecast $k = 30$ days ahead. Figure 12a and Figure b illustrate the Mahalanobis distances between the simulated infection trajectories and the observed data for the 2008 and 2020 crises, respectively, across different window lengths m . The orange line corresponds to the model incorporating change point detection, whereas the green line represents the INGARCH model. Across both crises, the change point model consistently achieves smaller distances, indicating that it better captures the dynamics of the infection process and better adapts to abrupt shifts in the data.

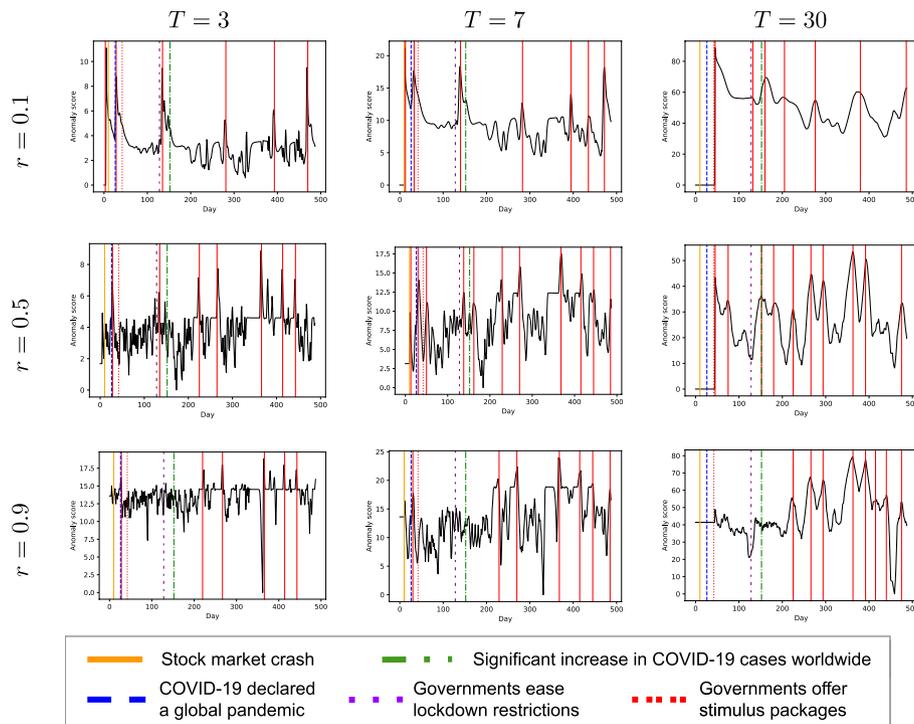


Fig. 11 The figures illustrate the anomaly scores output, using the ChangeFinder algorithm with different combinations of the discounting rate r and the rolling window size T . The ChangeFinder algorithm is run on the 2020 daily infections data. The solid red vertical lines illustrate points with high probability of being a change point, whereas the other vertical lines correspond to the significant events that happened during the crisis

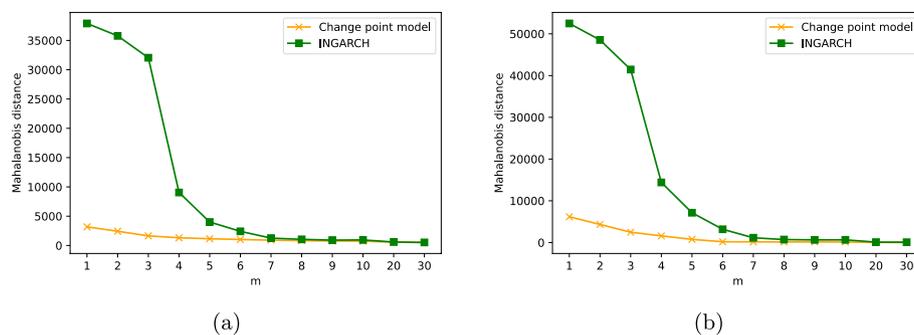


Fig. 12 Mahalanobis distances between the simulated infection trajectories and the observed data for (a) the 2008 crisis and (b) the 2020 crisis, evaluated for different window lengths m . The orange line corresponds to the model incorporating change point detection, while the green line represents the INGARCH model

Acknowledgements

We are grateful for the support of Russell Group Limited for providing the Moody’s dataset and for their valuable ideas and discussions throughout the project. We are grateful for access to the University of Nottingham’s Augusta HPC service for running simulations used in this study.

Author Contributions

MB collected the data, did the analysis and visualizations shown in the paper and wrote the manuscript. FB, YvG, GS and RDO did the supervision, reviewed and edited the manuscript. All authors read and approved the final manuscript.

Funding

MB has received funding in the form of an EPSRC CASE studentship EP/V520020/1 in collaboration with Russell Group Limited. GS has received financial support from the French National Research Agency under the grants ANR-23-CE40-0009 (EXSTA project) and ANR-11-LABX-0020-01 (Centre Henri Lebesgue), as well as from the TSE-HEC ACPR Chair and from the Chair Stress Test, RISK Management and Financial Steering of the Foundation Ecole Polytechnique.

YvG has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 777826.

Data Availability

The datasets used in this study are available in the GitHub repository, https://github.com/MBozhidarova98/Integrating_CP_PD_financial_crisis/tree/main/Datasets.

Code Availability

The codes used in the article are available in the GitHub repository, https://github.com/MBozhidarova98/Integrating_CP_PD_financial_crisis/tree/main.

Declarations

Conflict of interest

The authors declare no Conflict of interest.

Ethical approval

Not applicable

Consent for publication

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Received: 16 April 2025 / Accepted: 12 October 2025

Published online: 08 January 2026

References

- Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *Am Econ Rev* 105(2):564–608
- Adams RP, MacKay DJ (2007) Bayesian online changepoint detection. arXiv preprint [arXiv:0710.3742](https://arxiv.org/abs/0710.3742)
- Ahelegbey DF, Giudici P (2022) Netvix-a network volatility index of financial markets. *Physica A* 594:127017
- Alexander C (2008) *Market Risk Analysis, Pricing, Hedging and Trading Financial Instruments*. John Wiley & Sons, Chichester, England
- Aminikhanghahi S, Cook DJ (2017) A survey of methods for time series change point detection. *Knowl Inf Syst* 51(2):339–367
- Asongu SA (2012) The 2011 Japanese earthquake, tsunami and nuclear crisis: evidence of contagion from international financial markets. *J Financ Econ Policy* 4(4):340–353
- Avdjiev S, Giudici P, Spelta A (2019) Measuring contagion risk in international banking. *J Financ Stab* 42:36–51
- Battiston S, Farmer JD, Flache A, Garlaschelli D, Haldane AG, Heesterbeek H, Hommes C, Jaeger C, May R, Scheffer M (2016) Complexity theory and financial regulation. *Science* 351(6275):818–819
- Becker NG (1989) *Analysis of Infectious Disease Data*. Chapman and Hall/CRC, Boca Raton, Florida
- Belke A, Dubova I, Osowski T (2018) Policy uncertainty and international financial markets: the case of Brexit. *Appl Econ* 50(34–35):3752–3770
- Berkes I, Gombay E, Horváth L, Kokoszka P (2004) Sequential change-point detection in GARCH(p, q) models. *Economet Theor* 20(6):1140–1167
- Bozhidarova M, Ball F, Van Gennip Y, O'Dea RD, Stupfler G (2024) Describing financial crisis propagation through epidemic modelling on multiplex networks. *Proc Royal Soc A* 480(2287):20230787
- Brockwell PJ, Davis RA (2002) *Introduction to Time Series and Forecasting*. Springer, New York
- Cappé O, Moulines E, Rydén T (2005) *Inference in Hidden Markov Models*. Springer, New York
- Cauchemez S, Ferguson NM (2008) Likelihood-based estimation of continuous-time epidemic models from time-series data: application to measles transmission in London. *J R Soc Interface* 5(25):885–897
- Chevapatrakul T, Tee K-H (2014) The effects of news events on market contagion: evidence from the 2007–2009 financial crisis. *Res Int Bus Financ* 32:83–105
- Coles S, Heffernan J, Tawn J (1999) Dependence measures for extreme value analyses. *Extremes* 2:339–365
- Davis RA, Liu H (2016) Theory and inference for a class of nonlinear models with application to time series of counts. *Stat Sin* 26(4):1673–1707
- Dodge Y (2008) *The Concise Encyclopedia of Statistics*. Springer, New York
- Fiedor P (2014) Networks in financial markets based on the mutual information rate. *Phys Rev E* 89(5):052801
- Guidolin M (2011) *Markov Switching Models in Empirical Finance*. In: *Missing Data Methods: Time-Series Methods and Applications*. Emerald Group Publishing Limited, UK
- Held L, Hens N, O'Neill PD, Wallinga J (2019) *Handbook of Infectious Disease Data Analysis*. CRC Press, Boca Raton, Florida
- Hohlmeier M, Fahrholz C (2018) The impact of Brexit on financial markets-Taking stock. *Int J Financial Stud* 6(3):65
- Iwata T, Nakamura K, Tokusashi Y, Matsutani H (2019) 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, pp. 506–517. Springer
- Kawashima S, Takeda F (2012) The effect of the Fukushima nuclear accident on stock prices of electric power utilities in Japan. *Energy Economics* 34(6):2029–2038
- Le TH, Do HX, Nguyen DK, Sensoy A (2021) Covid-19 pandemic and tail-dependency networks of financial assets. *Financ Res Lett* 38:101800
- Maghyreh A, Abdoh H (2020) Tail dependence between Bitcoin and financial assets: evidence from a quantile cross-spectral approach. *Int Rev Financ Anal* 71:101545
- Makin AJ, Layton A (2021) The global fiscal response to COVID-19: risks and repercussions. *Econ Anal Policy* 69:340–349
- Mishkin FS (2009) Is monetary policy effective during financial crises? *American Econom Rev* 99(2):573–577

- Monteiro M, Costa M (2023) Change point detection by state space modeling of long-term air temperature series in Europe. *Stats* 6(1):113–130
- Musmeci N, Aste T, Di Matteo T (2015) Relation between financial market structure and the real economy: comparison between clustering methods. *PLoS ONE* 10(3):0116201
- Naseer S, Khalid S, Parveen S, Abbass K, Song H, Achim MV (2023) COVID-19 outbreak: impact on global economy. *Front Public Health* 10:1009393
- Nikkinen J, Omran MM, Sahlström P, Äijö J (2008) Stock returns and volatility following the September 11 attacks: evidence from 53 equity markets. *Int Rev Financ Anal* 17(1):27–46
- Pozzi F, Di Matteo T, Aste T (2013) Spread of risk across financial markets: better to invest in the peripheries. *Sci Rep* 3(1):1665
- Saaid F, Nur D, King R (2012) Change points detection of vector autoregressive model using SDVAR algorithm. *Applied Statistics Education and Research Collaboration (ASEARC) - Conference Papers*
- Schmidl S, Wenig P, Papenbrock T (2022) Anomaly detection in time series: a comprehensive evaluation. *Proc VLDB Endowment* 15(9):1779–1797
- Singh AK, Allen DE, Powell RJ (2017) Tail dependence analysis of stock markets using extreme value theory. *Appl Econ* 49(45):4588–4599
- Song W-M, Di Matteo T, Aste T (2012) Hierarchical information clustering by means of topologically embedded graphs. *PLoS ONE* 7(3):31929
- Takeuchi J, Yamanishi K (2006) A unifying framework for detecting outliers and change points from time series. *IEEE Trans Knowl Data Eng* 18(4):482–492
- Tumminello M, Aste T, Di Matteo T, Mantegna RN (2005) A tool for filtering information in complex systems. *Proc Natl Acad Sci* 102(30):10421–10426
- Wang L, Xue Z (2017) An automatic seismic event identification method by sequentially discounting autoregressive (SDAR). In: 2017 Workshop: Microseismic Technologies and Applications, Hefei, China, 4-6 June 2017, pp. 55–58. Society of Exploration Geophysicists

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.