International Journal of Financial Engineering
(2026) 2650028 (44 pages)

DOI: 10.1142/S2424786326500283

A structural framework for network-based default contagion in credit portfolios

B. Schiphorst [*,†,¶], M. R. H. Mandjes [*,‡], P. J. C. Spreij [*,§] and E. M. M. Winands [*,†]

[*]*Korteweg–de Vries Institute for Mathematics, University of Amsterdam*
Amsterdam, Science Park 107, 1098XH, The Netherlands

[†]*Rabobank, Utrecht, Croeselaan 18, 3521CB, The Netherlands*

[‡]*Mathematical Institute, Leiden University, Leiden*
P.O. Box 512, 2300RA, The Netherlands

[§]*IMAPP, Radboud University, Nijmegen*
Heyendaalseweg 135, 6525AJ, The Netherlands

Received: 20 November 2025; Accepted: 6 May 2026
Published:

Abstract

Portfolio credit risk models typically incorporate the default dependence between obligors that arises from common risk factors. However, they often ignore obligor-to-obligor default contagion effects, which can propagate through economic networks. We present a structural framework that can be layered onto existing factor models to capture higher-order network-based default contagion. Our approach yields efficient and tractable computations that are applicable to heterogeneous singly connected networks, which can represent empirically observed dependence structures such as supply chains, corporate groups and hierarchical governments. As a key contribution, we develop an expectation-maximization algorithm that estimates default contagion parameters from historical creditworthiness information. Our simulation-based numerical results show that the impact of default contagion on risk measures is well estimated by our framework.

Keywords: Portfolio credit risk; default contagion; structural model; expectation-maximization; singly connected network.

[¶]Corresponding author.
Email addresses: [¶]B.Schiphorst@uva.nl; budschiphorst@gmail.com

1. Introduction

Many financial institutions manage credit portfolios that are exposed to losses when obligors default. These portfolios face the risk of substantial aggregate losses when the defaults tend to cluster in time, even if the marginal default probability per obligor is relatively small. Effective risk management requires modeling the risk concentration due to default dependence between obligors. In this work, we consider a structural model that incorporates default dependence through exposure to common risk factors and through default contagion that propagates along a network. In particular, we develop a framework for estimating default contagion parameters from historical creditworthiness information, which can be used to estimate portfolio risk measures. We focus on the class of singly connected networks, in which there exists at most one single directed path between two obligors. This choice strikes a balance between practical relevance and model tractability, as it includes many empirically observed dependence structures while still allowing for tractable computations and data-driven parameter calibration.

Portfolio credit risk management typically focuses on default dependence that arises from common risk factors, such as macroeconomic or industry-specific influences. For example, a recession can lead to a simultaneous increase in default risk for a large group of obligors. A popular approach in practice is to use a factor model, which assumes that the defaults of obligors are conditionally independent given some (low-dimensional) underlying factor. In particular, industry practitioners often use Gaussian threshold models (see e.g., McNeil *et al.,* 2015, p. 465), which can be considered an extension of the structural framework by Merton (1974). For example, most banks use such a model to calculate (economic) capital requirements for credit risk in the Basel framework. In these models, default events occur when a latent creditworthiness variable, which can be interpreted as the distance-to-default, is below a specified default threshold. The default dependence between obligors originates from the shared exposure to common risk factors. These models are relatively easy to calibrate and can be used to estimate risk measures for heterogeneous credit portfolios using a simulation approach, which contributes to their widespread use.

However, standard factor models typically do not capture obligor-to-obligor default contagion, in which the default event of one obligor directly causes an increase in default risk of another obligor. This propagation of default contagion effects is observed in a variety of important dependence structures, e.g., from a parent company to subsidiaries (see e.g., Emery and Cantor, 2005), from a sovereign issuer to other domestic entities (see e.g., De Bruyckere *et al.*, 2013), or along supply chains (see e.g., Agca *et al.*, 2022). The latter is increasingly relevant for financial institutions, as there has been an increase in supply chain disruptions due to wars, volatile energy prices, and trade wars. This trend is likely to continue in the future due to increasing geopolitical tensions and emerging threats such as cyber risks and climate risks (both physical and transitional). Furthermore, default contagion has also received increased attention from regulatory authorities (see e.g., European Banking Authority, 2017).

Motivated by practical relevance and empirical evidence, our work contributes to the methodological literature on extending the standard factor models to incorporate default contagion effects. Our framework is applicable to a wide range of economic dependence structures, while allowing for data-driven calibration of default contagion parameters. Earlier work has already demonstrated that ignoring the presence of default contagion may lead to a significant underestimation of risk measures in credit portfolios (see e.g., Egloff *et al.*, 2007; Anagnostou *et al.*, 2018; Schiphorst *et al.*, 2024). However, the magnitude of the impact depends strongly on the network structure and the default contagion parameters that represent the strength of the dependence between obligors. Identifying the network structure corresponds to identifying the existence of a direct default dependence between pairs of obligors, which is often relatively straightforward in practice. For example, credit analysts can qualitatively assess the presence of legal or supply chain dependencies. In contrast, the estimation of default contagion parameters remains an open challenge, as the existing methodological literature reflects an important trade-off. At one end, frameworks that allow for complex network structures lack adequate methodology for parameter estimation (see e.g., Neu and Kühn, 2004; Egloff *et al.*, 2007; Anagnostou *et al.*, 2018). At the other end of the trade-off, existing estimation-focused frameworks require restrictive assumptions on the network structure (see e.g., Rösch and Winterfeldt, 2008; Schiphorst *et al.*, 2024). For example, Neu and Kühn (2004) suggest choosing parameters based on expert judgment, Egloff *et al.* (2007) propose an (ad hoc) proxy based on percentages of business volume, and Anagnostou *et al.* (2018) propose a proxy based on spike-like comovements in CDS spreads. Conversely, Rösch and Winterfeldt (2008) propose an estimation framework based on historical default rates, but require a strong homogeneity assumption of a single contagion parameter for the network connections. In addition, they impose a restriction on the dependence structure that is similar to the primary-secondary model in Jarrow and Yu (2001), in which contagion can only propagate from a primary group to a secondary group of obligors. The framework based on these strong assumptions has been adopted by several others (e.g., Lee and Poon, 2014; Batiz-Zuk *et al.*, 2015). Schiphorst *et al.* (2024) instead propose a moment-based estimation from historically observed creditworthiness information, which allows for relaxing the homogeneity assumption but still requires the primary-secondary network restriction. To bridge this gap in the literature, we develop a structural framework that allows for estimating default contagion parameters in a more general class of practically relevant networks.

Our proposed framework is designed to be layered onto existing structural factor models to additionally capture default contagion. As a key contribution, we derive an iterative algorithm to calibrate the model to exogenously given default probabilities via the latent distance-to-default of each obligor. Consequently, the framework has the desirable property that it can be layered onto an existing model without affecting the expected credit losses. Building on the calibration to given default probabilities, we propose a maximum-likelihood approach for estimating default contagion parameters from time-series of historical creditworthiness information. This approach does not rely on homogeneity assumptions, works for models with observed or latent factors, and is applicable to a wide range of dependence structures. More specifically, the framework applies to the class of singly connected networks, sometimes referred to as polytrees (see e.g., Pearl, 1988, p. 176). This class allows for efficient tractable computations and captures many important features of common economic dependence structures. For example, it can represent hierarchical dependencies with higher-order propagation of default contagion, such as corporate parent-subsidiary relationships, supply chains, and governmental hierarchies. In particular, the class generalizes the restrictive primary-secondary dependence structure described above, since it includes all such networks as special cases. Moreover, it could serve as an approximation for networks that only slightly deviate from a singly connected network structure.

The default contagion mechanism considered in our framework is inspired by the model of Neu and Kühn (2004) and the proposed estimation approach builds on an idea explored in Schiphorst *et al.* (2024), in which the latent distance-to-default is inferred from observed creditworthiness information. This work is complementary to Schiphorst *et al.* (2024), which considered a different contagion mechanism that yielded relatively simple closed-form results at the cost of being restricted to a primary-secondary structure. Overall, our contribution is primarily methodological and complements existing empirical work that has demonstrated the practical importance of default contagion in real economic networks (see e.g., Agca *et al.*, 2022). Our framework is designed to be readily applicable for institutions that possess the necessary data. To illustrate the applicability of our framework, we present a numerical example with empirically motivated choices for the network structure and other model parameters. We use simulated time-series of creditworthiness information to illustrate some convergence properties of our proposed estimation framework.

The remainder of this work is structured as follows. In Sec. 2, we present the structural model framework with default contagion effects and explain how the model can be used for portfolio credit risk management. In Sec. 3, we derive a framework for numerically inferring the latent distance-to-default from observed creditworthiness information. In Sec. 4, we present the framework for estimating the default contagion parameters from historical creditworthiness information. More specifically, we use an expectation-maximization approach to maximize the likelihood of time series of observed creditworthiness information. In Sec. 5, we demonstrate the application of the proposed framework and illustrate that portfolio risk measures are well approximated by using estimated default contagion parameters. Finally, in Sec. 6, we conclude on the obtained results and provide some final remarks.

2. Structural Model with Default Contagion

In this section, we present the structural model for the joint distribution of default events with dependence through exposure to common risk factors and through default contagion. First, Sec. 2.1 presents the details of the latent joint distance-to-default process and the default contagion network. Then, Sec. 2.2 outlines the details of the structural default event and the contagion mechanism. Finally, Sec. 2.3 describes how the model can be used in practice.

2.1. Distance-to-default

We consider a credit portfolio with obligors indexed by $\mathcal{I} = \{1, \ldots, N\}$ and introduce the latent joint distance-to-default process $V = [V_i]_{i \in \mathcal{I}}$. We focus on a structural Gaussian factor model, which is based on the structural model by Merton (1974) and is commonly used among industry practitioners (see e.g., Lütkebohmert, 2008, p. 23). More specifically, we let F and each ξ_i (for $i \in \mathcal{I}$) be independent standard Brownian motions and define the distance-to-default $V_i(t)$ of obligor $i \in \mathcal{I}$ at time t as

$$V_i(t) = V_i(0) + \sqrt{\rho_i} F(t) + \sqrt{1 - \rho_i} \xi_i(t). \tag{1}$$

The first term, $V_i(0)$, represents the initial distance-to-default of obligor $i \in \mathcal{I}$ at initial time 0. The larger (smaller) this initial value, the less (more) likely that the obligor will be in default at time T. The second term, $\sqrt{\rho_i} F(t)$, represents the dependence on the common factor F. The larger the factor loading $\sqrt{\rho_i}$, the stronger the indirect default dependence of obligor i with other obligors. The third term, $\sqrt{1 - \rho_i} \xi_i(t)$, represents idiosyncratic effects that are independent of other obligors and the common factor. In this model, the distance-to-default V_i, V_j of two different obligors $i, j \in \mathcal{I}$ are conditionally independent given the common factor F.

We note that the model in Eq. (1) is sometimes generalized by considering a multivariate process for the common factor F or alternative distributions such as Student's t processes for $(F, \{\xi_i\}_{i\in\mathcal{I}})$ (see e.g., McNeil *et al.*, 2015, pp. 442–443). However, this would yield a conceptually equivalent model, while making the model less recognizable to industry practitioners. Since our focus is on the default contagion extension, we decide not to pursue this possibility here.

To model default contagion effects, we consider a directed network with nodes in $\mathcal{I}$ and nonnegative edge weights $W = [W_{ij}]_{i,j\in\mathcal{I}}$. We will refer to this network as the default contagion network. This network has a directed edge from $j \in \mathcal{I}$ to $i \in \mathcal{I}$ whenever $W_{ij} > 0$, in which case we say that j is a parent of i and that i is a child of j. Similarly, if there exists a directed path from j to i, we say that j is an ancestor of i and that i is a descendant of j. For any obligor $i \in \mathcal{I}$, we denote the corresponding set of parents as

$$\mathcal{P}_i := \{j \in \mathcal{I} \mid j \text{ is a parent of } i\} = \left\{j \in \mathcal{I} \mid W_{ij} > 0\right\}.$$

We assume throughout this paper that we know the binary parent structure, i.e., that $\mathcal{P}_i$ is given for all obligors $i \in \mathcal{I}$. This assumption is further motivated in Sec. 4.2. Furthermore, we assume that the network is singly connected (i.e., a polytree), which means that there is at most one single directed path between each pair of obligors (see e.g., Pearl, 1988, p. 176). We illustrate in Sec. 2.3.1 that this assumption still yields a broad class of default contagion networks. Because a singly connected network is always acyclic, we can order the set of obligors $\mathcal{I} = \{1, \ldots, N\}$ such that the directed edges go only from lower-indexed to higher-indexed obligors. More precisely, we assume without loss of generality that

$$\mathcal{I} \text{ is ordered such that } j < i, \text{ for all parents } j \in \mathcal{P}_i \text{ of obligor } i \in \mathcal{I}. \quad (2)$$

2.2. Structural default events with contagion

We consider a two-period model, with initial time 0 and end time T (e.g., in 1 year). For each obligor $i \in \mathcal{I}$, we model a default event that may occur at time T, represented by the default indicator

$$\delta_i := 1\{\text{Default of obligor } i \text{ at time } T\}.$$

We take a structural approach and assume that a default event occurs when the latent distance-to-default $V_i(T)$ of obligor $i \in \mathcal{I}$ at time T is below a certain default threshold. We incorporate default contagion by letting the default threshold depend on the default events of any parents $j \in \mathcal{P}_i$. More specifically, we assume that

$$\delta_i = 1\left\{V_i(T) \le \sum_{j\in\mathcal{P}_i} W_{ij}\delta_j\right\}, \quad \forall\, i \in \mathcal{I}. \tag{3}$$

In the absence of default contagion, i.e.,when none of the parents $j \in \mathcal{P}_i$ default, the default event of obligor $i \in \mathcal{I}$ occurs if $V_i(T) \le 0$. Default contagion arises when a parent $j \in \mathcal{P}_i$ defaults and the default threshold of obligor i increases by W_{ij}, consequently increasing the default probability. We note that this contagion mechanism is conceptually similar to the model proposed by Neu and Kühn (2004), in which the distance-to-default is adjusted instead of the default threshold. In principle, other invertible transformations of the distance-to-default could be considered without significant conceptual changes to our framework.

In the considered approach, there exist multiple types of dependence between the default events of two obligors $i, j \in \mathcal{I}$. There may be indirect dependence due to their shared exposure to the common factor F, which is reflected by the distance-to-default correlation $\mathrm{Corr}(V_i(T), V_j(T)) = \sqrt{\rho_i \rho_j}$. Additionally, there can exist dependence between the default events of obligors due to default contagion, which can be direct or indirect. Indirect default contagion dependence arises when the obligors i, j share a common ancestor. There is direct default contagion dependence if there exists a directed path between obligor i and obligor j, i.e., when one is an ancestor of the other.

Unlike standard factor models without default contagion, the default events of two obligors $i, j \in \mathcal{I}$ are generally not conditionally independent given the factor F in this model. Instead, this conditional independence only holds if there exists no directed path between two obligors i, j and they do not share any common ancestors. In contrast, for a very large value W_{ij}, the default of the obligor $j \in \mathcal{P}_i$ almost guarantees the default of obligor $i \in \mathcal{I}$. This is because $\lim_{W_{ij} \to \infty} \mathbb{P}\left[V_i(T) \le W_{ij}\delta_j | \delta_j = 1\right] = 1$. The case where the default of one obligor guarantees the default of another obligor exists in real applications, for example in corporate parent-subsidiary structures.

2.3. Model use

In this subsection, we explain how the model can be used for portfolio credit risk management. First, Sec. 2.3.1 illustrates that the model can be applied to a wide range of credit portfolios. Second, Sec. 2.3.2 illustrates the significance of default contagion in computing risk measures.

2.3.1. *Singly connected networks*

In this work, we assume that the directed default contagion network is singly connected, i.e., that there exists at most a single directed path between any two obligors. Figure 1 displays a visual example of a (fragment of a) typical singly connected network, which is similar to the example presented by Pearl (1988, p. 176). This choice is made specifically with application in mind, as it allows for tractable model computations, but can also capture many important characteristics of empirically observed economic networks. These include sparseness of network connections and the existence of obligors with multiple parents/children. Furthermore, we highlight some special cases of singly connected networks: First, the class obviously includes the empty graph, which corresponds to the special case $W_{ij} = 0$ for all obligors $i, j \in \mathcal{I}$. This case represents the absence of default contagion effects, in which the model simplifies to a standard threshold model with conditionally independent default events.

As a second example, the class of singly connected networks includes all networks that satisfy the primary-secondary structure in which connections can go only from a group of primary obligors to a group of secondary obligors. This assumption has appeared in the literature as an attempt to keep models and their calibration relatively simple (see e.g., Jarrow and Yu, 2001; Rösch and Winterfeldt, 2008; Batiz-Zuk *et al.*, 2015; Lee and Poon, 2014; Schiphorst *et al.*, 2024). For example, the primary group could represent a group of important entities that

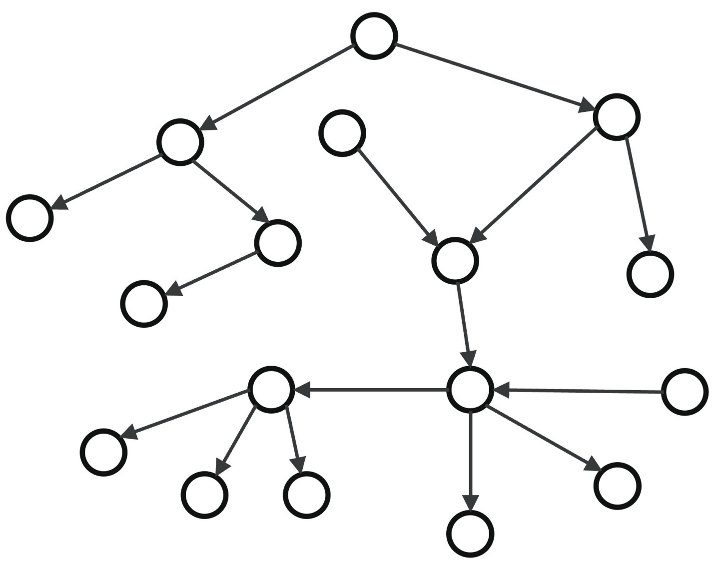

Fig. 1. A visual example of a directed singly connected network.

would affect many other smaller (secondary) entities in case of a default event. An important limitation of the primary-secondary assumption is that it inherently does not allow for higher-order propagation default contagion effects.

As another example, the class includes chain-like networks in which each obligor $i \in \mathcal{I}$ can have a connection to the next obligor $j = i + 1$. Such networks could represent supply chain structures in which higher-order default contagion has been empirically observed (see e.g., Agca *et al.*, 2022). Similarly, the class also includes more general networks that consist of a collection of tree-like structures. Such networks represent hierarchical forms of default dependence that could arise in many different situations, such as corporate parent-subsidiary relationships and layered supply chains. Another example is the hierarchical structure between national governments, regional governments, and municipalities. A specific example of a network consisting of a collection of connected tree-like structures is considered in more detail in Sec. 5.1.

Additionally, we note that the assumption of a singly connected network structure does not impose any restriction on the network weights. In this work, each connection in the default contagion network can have a different weight. This represents the heterogeneity of realistic credit portfolios and significantly expands the class of feasible default contagion networks compared to the assumption of a single homogeneous weight for a large group of connections in the default contagion network, as is considered in earlier estimation-focused methodological literature (e.g., Rösch and Winterfeldt, 2008; Batiz-Zuk *et al.*, 2015; Lee and Poon, 2014).

2.3.2. *Portfolio credit losses*

Within the considered model, portfolio credit losses L are represented by

$$L = \sum_{i \in \mathcal{I}} e_i \delta_i,$$

where e_i is the exposure to obligor $i \in \mathcal{I}$ that is lost in case of default and is assumed to be deterministic. We note that in credit risk modeling, each e_i is typically split into the product of an exposure at default EAD_i and a loss given default percentage LGD_i. The probability of a default event of obligor $i \in \mathcal{I}$ is denoted as $\mathrm{PD}_i = \mathbb{P}[\delta_i = 1]$. As is common practice in portfolio credit risk modeling, we treat the exposures lost in default $\{e_i\}_{i \in \mathcal{I}}$ and the default probabilities $\{\mathrm{PD}_i\}_{i \in \mathcal{I}}$ as exogenous inputs.

Expected losses. We note that expected portfolio losses are independent of the default contagion parameter W and always equal

$$\mathbb{E}[L] = \sum_{i \in \mathcal{I}} e_i \mathbb{E}[\delta_i] \equiv \sum_{i \in \mathcal{I}} e_i \mathrm{PD}_i.$$

As explained in Sec. 3, our framework preserves this property by calibrating the latent distance-to-default values such that the modeled default probabilities are equal to the exogenously given default probabilities. Although the calibrated distance-to-default values depend on W, the modeled default probabilities are, by construction, always the same. This is an important attractive property of our framework, which allows it to be layered onto existing portfolio credit risk models, without affecting the expected losses. Implicitly, it is assumed that the impact of default contagion is already incorporated in the given default probabilities.

Risk measures. The risk of significant losses is often quantified by risk measures such as the Value-at-Risk (VaR) or the Expected Shortfall (ES) (see e.g., Lütkebohmert, 2008, pp. 13–15). For a given confidence level q, these risk measures can be defined as

$$\mathrm{VaR}_q(L) := \inf\{x \in \mathbb{R} | \mathbb{P}[L > x] \leq 1 - q\}, \tag{4}$$

and

$$\mathrm{ES}_q(L) := \frac{1}{1-q} \int_q^1 \mathrm{VaR}_u(L) \mathrm{d}u. \tag{5}$$

Due to the heterogeneity of credit portfolios, analytic computation of these risk measures is generally not possible. Instead, risk measures can be estimated by simulating portfolio credit losses based on the joint distribution of default events.

Similarly to standard portfolio credit risk models that are commonly used in practice, the simulation of default events consists of two steps. First, a realization of the joint distance-to-default $V(T) = [V_i(T)]_{i \in \mathcal{I}}$ is drawn based on the model specified in Eq. (1). This initial step is identical to the one in standard factor models without default contagion. Second, the realization of $V(T)$ is used to determine the corresponding realization of default events, by iteratively evaluating Eq. (3) for all obligors along the default contagion network. In particular, the ordering assumed in Eq. (2) ensures that the default threshold of any obligor $i \in \mathcal{I}$ has already been computed before the default indicator δ_i is evaluated. We note that iterative evaluation of the default contagion thresholds is computationally inexpensive and can be added to existing implementations of standard factor models. In particular, we note that it is still possible to use variance reduction methods such as importance sampling.

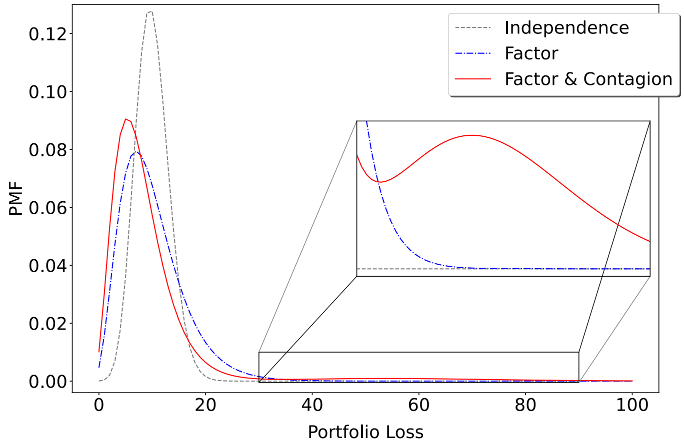

Fig. 2. Probability mass function (PMF) of total portfolio losses L in a stylized portfolio, comparing three dependence structures: independence, a standard factor model and a factor model with default contagion.

Loss distribution. We now demonstrate that default contagion can induce a fundamentally different shape of the loss distribution compared to standard correlation-based factor models. For this purpose, we consider a credit portfolio that consists of $N = 250$ obligors with homogeneous parameters $\mathrm{PD}_i \equiv 0.04$ and $e_i = 1$ for all $i \in \mathcal{I}$. Figure 2 shows the portfolio loss distribution for three cases that correspond to different types of dependence between obligors. The first case serves as a baseline and represents full independence between the obligors, i.e., without common factor dependence and without default contagion. The second case represents a standard single factor model with default dependence through a shared common factor, with factor loadings $\rho_i \equiv 0.05$ for all $i \in \mathcal{I}$. The third case represents the same factor model extended with additional default contagion dependence along a network with a single parent of all other obligors. More specifically, the network weights are chosen to equal $W_{i1} = 1.0$ for all $i \in \mathcal{I}$ and 0 otherwise.

In the case with default contagion, the loss distribution has a bimodal shape with an increased probability mass in the tail, driven by the binary default event of the central parent. Although the direct exposure e_1 to the central parent is relatively limited, their default event has a significant impact on total losses due to contagion effects, increasing the default risk of all other obligors. By construction, the expected total loss equals $\mathbb{E}[L] = 10$ in all three cases, since it is assumed that the effects of contagion are already incorporated in the exogenously given default probabilities. Consequently, the increased tail mass from default contagion is offset by a reduction in risk in the absence of the parent default, as represented by the leftward shift of the body of the distribution. Overall, this demonstrates that default contagion can have a significant impact on risk measures, which is revisited in a more complex network structure in Sec. 5.

3. Distance-to-Default and Creditworthiness

In this section, we explain how the initial joint distance-to-default $V(0) = v$ can be inferred from observed creditworthiness information in the form of forward-looking default probabilities. This initial value is required to simulate the joint distribution of default events, as explained in Sec. 2.3.2. Furthermore, the inferred distance-to-default is also used to estimate the default contagion parameters W, as outlined in Sec. 4.3.

First, Sec. 3.1 discusses creditworthiness information and how it can be obtained in practice. Then, Sec. 3.2 outlines how the creditworthiness information of an obligor relates to the initial distance-to-default in the considered structural model. More specifically, we outline an algorithm that numerically infers the latent initial distance-to-default of each obligor from observed creditworthiness information.

3.1. Creditworthiness

The creditworthiness of an obligor represents the (perceived) likelihood that they will not default. It is mathematically convenient to represent creditworthiness by a forward-looking default probability. That is, the creditworthiness of obligor $i \in \mathcal{I}$ can be represented by the forward-looking probability PD_i of a default event at the end time T, given the information available at the initial time $t = 0$. In the context of the structural model presented in Sec. 2, this is represented by

$$\mathrm{PD}_i(v) := \mathbb{P}[\delta_i = 1 | V(0) = v], \tag{6}$$

where the initial joint distance-to-default v represents the information available at the initial time 0.

In practice, estimated forward-looking default probabilities are often available. For example, financial institutions and credit rating agencies often have dedicated teams of credit analysts that produce these estimates. These analysts can use quantitative and qualitative information from various sources, such as financial statements, macroeconomic data, and industry trends. The default probability estimates are sometimes grouped into credit ratings, which are essentially a discretized form of creditworthiness information. Alternatively, forward-looking default probabilities can be inferred from market prices of credit default swaps (CDS) or defaultable bonds (see e.g., Hull, 2022, pp. 564–569).

3.2. Inferring distance-to-default

In this subsection, we outline how the latent distance-to-default can be inferred from observed creditworthiness information. First, Sec. 3.2.1 explains the general setup. Then, Sec. 3.2.2 presents some results for computing conditional default probabilities for all obligors $i \in \mathcal{I}$. Finally, Sec. 3.2.3 discusses how the conditional default probabilities can be used to (numerically) infer the latent distance-to-default.

3.2.1. *Distance-to-default and creditworthiness*

As expected, the forward-looking default probability PD_i defined in Eq. (6) is strictly decreasing in the initial distance-to-default v_i. In fact, there exists a one-to-one mapping between the vector of default probabilities $\mathrm{PD} = [\mathrm{PD}_i]_{i\in\mathcal{I}}$ and the initial joint distance-to-default $v = [v_i]_{i\in\mathcal{I}}$ for any given model parameters (ρ, W). Therefore, it is possible to infer the latent distance-to-default v from observed creditworthiness information $\mathrm{pd} = [\mathrm{pd}_i]_{i\in\mathcal{I}}$ by jointly solving the following equations:

$$\mathrm{PD}_i(v) \equiv \mathrm{pd}_i, \quad \forall\, i \in \mathcal{I}.$$

As we will illustrate below, it is useful to compute the default probability by integrating over the conditional default probability $\mathrm{PD}_{i|F}$ given the common factor F, i.e.,

$$\mathrm{PD}_i(v) = \int_{-\infty}^{\infty} \mathrm{PD}_{i|F}(v)\mathrm{d}\mathbb{P}[F], \tag{7}$$

where

$$\mathrm{PD}_{i|F}(v) := \mathbb{P}[\delta_i = 1 | F(T) = F, V(0) = v].$$

We first show how the distance-to-default can be inferred in the absence of default contagion, and then consider the more general case with default contagion.

Without default contagion. In the absence of default contagion, i.e., for an obligor $i \in \mathcal{I}$ without parents, the conditional default probability equals

$$\mathrm{PD}_{i|F}(v) = \Phi\left(\frac{-v_i - \sqrt{\rho_i}F}{\sqrt{1-\rho_i}\sqrt{T}}\right), \tag{8}$$

where Φ is the standard Gaussian distribution function. In this special case, the unconditional default probability conveniently simplifies as

$$\mathrm{PD}_i(v) = \int_{-\infty}^{\infty} \Phi\left(\frac{-v_i - \sqrt{\rho_i}F}{\sqrt{1-\rho_i}\sqrt{T}}\right)\mathrm{d}\mathbb{P}[F] = \Phi\left(-\frac{v_i}{\sqrt{T}}\right),$$

and the inferred distance-to-default v_i is given by

$$v_i = -\sqrt{T} \cdot \Phi^{-1}(\mathrm{pd}_i). \tag{9}$$

With default contagion. Unfortunately, the conditional default probability $\mathrm{PD}_{i|F}$ does not attain such a simple form in the more general case with default contagion, as it can then also depend on the default events of the parents $j \in \mathcal{P}_i$. Consequently, the unconditional default probability and the inferred distance-to-default cannot be computed in a closed form.

However, as we show in Sec. 3.2.2, the conditional default probabilities can still be computed efficiently for all obligors $i \in \mathcal{I}$ if the default contagion network is singly connected. Moreover, as we discuss in Sec. 3.2.3, we can approximate the unconditional default probability by numerical integration and numerically solve for the inferred distance-to-default.

3.2.2. *Conditional default probabilities*

We present several results for computing the conditional default probability $\mathrm{PD}_{i|F}$ of an obligor $i \in \mathcal{I}$. The presented results correspond to varying levels of restricting assumptions, where more restrictive assumptions yield simpler results. First, the more general case of an acyclic network is considered in Proposition 1. Then, the slightly stronger assumption of a (heterogeneous) singly connected network is considered in Proposition 2. The computation simplifies significantly in singly connected networks, due to the conditional independence of the default events of parents. Corollary 1 corresponds to the even stronger assumption of a primary-secondary structure with homogeneous parameters.

Proposition 1. *Suppose that the default contagion network represented by W is acyclic. Then the marginal conditional default probability $\mathrm{PD}_{i|F}$ of obligor $i \in \mathcal{I}$ is given by*

$$\mathrm{PD}_{i|F}(v) = \sum_{d \in \{0,1\}^{|\mathcal{P}_i|}} \Phi\left(\frac{\sum_{j \in \mathcal{P}_i} W_{ij} d_j - v_i - \sqrt{\rho_i} F}{\sqrt{1-\rho_i}\sqrt{T}}\right) \mathbb{P}\big[\delta_{\mathcal{P}_i} = d | F\big], \tag{10}$$

where $|\mathcal{P}_i|$ is the number of parents of i, Φ is the standard Gaussian CDF and $\delta_{\mathcal{P}_i} := [\delta_j]_{j \in \mathcal{P}_i}$ denotes the default indicators of parents $j \in \mathcal{P}_i$.

Proof. Consider a fixed initial distance-to-default $V_i(0) = v_i$. The conditional default probability $\mathrm{PD}_{i|F}$ can then be decomposed as

$$\begin{aligned}\mathrm{PD}_{i|F} &= \mathbb{P}[\delta_i = 1|F] \\ &= \sum_{d \in \{0,1\}^{|\mathcal{P}_i|}} \mathbb{P}\big[\delta_i = 1|F, \delta_{\mathcal{P}_i} = d\big]\mathbb{P}\big[\delta_{\mathcal{P}_i} = d|F\big].\end{aligned}$$

Furthermore, in the structural default model, we have that

$$\begin{aligned}\delta_i = 1 \quad &\overset{(3)}{\Longleftrightarrow} \quad V_i(T) \le \sum_{j \in \mathcal{P}_i} W_{ij}\delta_j \\ &\overset{(1)}{\Longleftrightarrow} \quad \xi_i(T) \le \frac{\sum_{j \in \mathcal{P}_i} W_{ij}\delta_j - v_i - \sqrt{\rho_i}F(T)}{\sqrt{1-\rho_i}}.\end{aligned}$$

So, by using $\xi_i(T) \sim \mathcal{N}(0, T)$ and $\xi_i(T) \perp (\delta_{\mathcal{P}_i}, F)$ we obtain

$$\mathbb{P}\big[\delta_i = 1|F, \delta_{\mathcal{P}_i} = d\big] = \Phi\left(\frac{\sum_{j \in \mathcal{P}_i} W_{ij}d_j - v_i - \sqrt{\rho_i}F}{\sqrt{1-\rho_i}\sqrt{T}}\right).$$

□

Essentially, the result in Proposition 1 splits the computation of the conditional default probability $\mathrm{PD}_{i|F}$, by conditioning on the (joint) default events of the parents $j \in \mathcal{P}_i$. In the case of a general acyclic network, the computational complexity mainly lies in the joint conditional default probabilities $\mathbb{P}\big[\delta_{\mathcal{P}_i} = d|F\big]$. The computation of them may depend on all ancestors of i and thus requires "global" network information. These joint probabilities could for example be computed by using a variable elimination algorithm (see e.g., Koller and Friedman, 2009, p. 287).

Proposition 2. *The following two statements are equivalent*:

(1) *The default contagion network represented by W is singly connected.*
(2) *For any obligor $i \in \mathcal{I}$, the default events of any two parents $j_1, j_2 \in \mathcal{P}_i$ are conditionally independent given F, i.e.*, $(\delta_{j_1} \perp \delta_{j_2}|F)$.

If either holds, then the joint conditional default probability of the parents $j \in \mathcal{P}_i$ decomposes as

$$\mathbb{P}\big[\delta_{\mathcal{P}_i} = d|F\big] = \prod_{j \in \mathcal{P}_i} \mathrm{PD}_{j|F}^{d_j}\big(1 - \mathrm{PD}_{j|F}\big)^{1-d_j}, \quad \forall\, d \in \{0,1\}^{|\mathcal{P}_i|}.$$

Proof. We first note that the default events of two obligors are conditionally independent given F if they do not share a common ancestor and neither is an ancestor of the other. These conditions are equivalent to the obligors being "d-separated given $\emptyset$" in the network, which is a criterion for determining independence in probabilistic graphical models (see e.g., Koller and Friedman, 2009, pp. 69–71).

(1) $\Rightarrow$ (2): Assume that the network represented by W is singly connected. Suppose, for contradiction, that there exists an obligor $i \in \mathcal{I}$ with parents $j_1, j_2 \in \mathcal{P}_i$ that are not conditionally independent given F. It follows that there exists a path between j_1 and j_2 or that there exists an obligor k that is a common ancestor of j_1 and j_2. Either would lead to a contradiction by violating the assumption of singly connectedness. A path from j_1 to j_2 would imply that there exist two paths from j_1 to i: one direct path and one indirect path via j_2. Similarly, a common ancestor k would imply that there exist two paths from k to i: one via j_1 and one via j_2.

$\neg$(1) $\Rightarrow$ $\neg$(2): Assume that the network represented by W is not singly connected. Then there exist obligors $k, l \in \mathcal{I}$ with (at least) two distinct directed paths from k to l. Since the distinct paths both end at obligor l, the paths must first diverge and then converge again at some obligor $i \in \mathcal{I}$. This obligor i thus has two distinct parents $j_1, j_2 \in \mathcal{P}_i$ that are part of the first and second paths, respectively. However, the parents j_1, j_2 are not conditionally independent given F, as they are both part of a path that starts with obligor k. □

Proposition 2 shows that the computation of the joint conditional default probabilities $\mathbb{P}\left[\delta_{\mathcal{P}_i} = d|F\right]$ of the parents $j \in \mathcal{P}_i$ simplifies significantly in the case of a singly connected network, due to the conditional independence of the parent default events. Essentially, it eliminates the need for "global" network information about the ancestors and instead only requires "local" information about the marginal parent probabilities. The result allows for efficient computation of the conditional default probability $\mathrm{PD}_{i|F}$ when $\mathrm{PD}_{j|F}$ has already been computed for all parents $j \in \mathcal{P}_i$. Assuming that the obligors in the default contagion network are ordered as in Eq. (2), the result can thus be used to iteratively compute the conditional default probabilities $\mathrm{PD}_{i|F}$ for all the obligors $i \in \mathcal{I}$ along the network. Crucially, the equivalence in Proposition 2 identifies the class of singly connected networks as the most general class that preserves this property. This is the primary motivation for our methodological decision to exclude more general network structures with prominent loops, such as interbank lending networks.

Corollary 1 illustrates how the result simplifies further in the case of a more restrictive assumption of a homogeneous two-group network. More specifically, it assumes that all obligors are split into a primary group $\mathcal{I}_P$ and a secondary group $\mathcal{I}_S$, such that all connections only go from primary to secondary obligors. Moreover, it assumes that all default contagion weights equal w and that the (conditional) default probabilities of all obligors $j \in \mathcal{I}_P$ are equal. A similar result was derived in Rösch and Winterfeldt (2008).

Corollary 1. *Suppose that there exists a partition* $\{\mathcal{I}_P, \mathcal{I}_S\}$ *of* $\mathcal{I}$ *and a weight* $w > 0$, *such that* $W_{ij} = w \cdot 1\{i \in \mathcal{I}_S \text{ and } j \in \mathcal{I}_P\}$, for all $i, j \in \mathcal{I}$. *Also, suppose that* $\mathrm{PD}_{j|F} \equiv \mathrm{PD}_{P|F}$, for all $j \in \mathcal{I}_P$. *Then, each obligor* $i \in \mathcal{I}_S$ *has conditional default probability*

$$\mathrm{PD}_{i|F}(v) = \sum_{k=0}^{N_P} \Phi\left(\frac{wk - v_i - \sqrt{\rho_i}F}{\sqrt{1-\rho_i}\sqrt{T}}\right)\binom{N_P}{k}\mathrm{PD}_{P|F}^k(1-\mathrm{PD}_{P|F})^{N_P-k},$$

where $N_P := |\mathcal{I}_P|$ *equals the number of obligors in the primary group.*

Proof. Due to the homogeneous default contagion parameter w, it is sufficient to consider only the total number of parent defaults

$$D_i := \sum_{j \in \mathcal{P}_i} \delta_j = \sum_{j \in \mathcal{I}_P} \delta_j.$$

Analogous to the proof of Proposition 1, we can decompose the conditional default probability $\mathrm{PD}_{i|F}$ as

$$\begin{aligned}\mathrm{PD}_{i|F} &= \mathbb{P}[\delta_i = 1|F] \\ &= \sum_{k=0}^{N_P} \mathbb{P}[\delta_i = 1|F, D_i = k]\mathbb{P}[D_i = k|F],\end{aligned}$$

where

$$\mathbb{P}[\delta_i = 1|F, D_i = k] = \Phi\left(\frac{wk - v_i - \sqrt{\rho_i}F}{\sqrt{1-\rho_i}\sqrt{T}}\right).$$

Since a primary-secondary network is singly connected, the parent default events are conditionally independent given F. Thus, conditional on F, D_i follows a binomial distribution with probability $\mathrm{PD}_{P|F}$. That is

$$\mathbb{P}[D_i = k|F] = \binom{N_P}{k}\mathrm{PD}_{P|F}^k(1-\mathrm{PD}_{P|F})^{N_P-k}.$$

□

We note that a similar result can be obtained when the assumption of equal conditional default probabilities in Corollary 1 is relaxed. In that case, the total number of parent defaults D_i instead follows a "Poisson binomial" distribution conditional on F. We refer the reader to the literature for details on computing Poisson binomial probabilities (see e.g., Hong, 2013).

Also, we note that for an obligor without parents, the conditional default probability is given by Eq. (8). In an acyclic network, there always exists at least one obligor that has no parents.

3.2.3. *Unconditional default probability*

As noted in Sec. 3.2.1, the unconditional default probability PD_i does not have a simple closed form in the presence of default contagion effects. Therefore, we use numerical integration to approximate the unconditional default probability PD_i from the conditional default probabilities $\mathrm{PD}_{i|F}$. The numerical integration approximation $\widehat{\mathrm{PD}}_i(v)$ is defined by

$$\mathrm{PD}_i(v) = \int_{-\infty}^{\infty} \mathrm{PD}_{i|F}(v)\mathrm{d}\mathbb{P}(F) \approx \sum_{r=1}^{R} \omega_r \mathrm{PD}_{i|F_r}(v) =: \hat{\mathrm{PD}}_i(v), \tag{11}$$

where the pairs $\{(\omega_r, F_r)\}_{r=1}^{R}$ of numerical integration weights ω_r and nodes F_r are defined as in Appendix A.1.

The approximation $\widehat{\mathrm{PD}}_i$ is used to numerically infer the distance-to-default. More specifically, the numerically inferred distance-to-default v is given by the unique solution to the equations

$$\widehat{\mathrm{PD}}_i(v) \equiv \mathrm{pd}_i, \text{ for all } i \in \mathcal{I}. \tag{12}$$

Appendix A.2 describes the details for consistently finding the numerical solution to Eq. (12) by using a standard univariate root-finding algorithm. More specifically, it is proven that there always exists a unique solution and also that the solution converges to the true distance-to-default as the numerical integration grid becomes finer. In addition, Appendix A.2 presents some derivatives that can be used for gradient-based root-finding algorithms.

Algorithm 1 outlines all the required steps for iteratively inferring the distance-to-default v for all obligors $i \in \mathcal{I}$ in a singly connected default contagion network. In essence, the algorithm leverages the ordering assumed in Eq. (2) and Proposition 2 to iteratively compute the distance-to-default along the network. For obligors without parents, the inferred distance-to-default is computed as in Eq. (9). For obligors with parents, the inferred distance-to-default is instead numerically inferred as explained above.

4. Estimating Default Contagion Parameters

We build on the inferred distance-to-default approach and propose a framework to estimate the default contagion parameters W from time series of observed creditworthiness information. The proposed framework is designed such that it can be layered onto existing factor models, regardless of whether the common factor is observed or unobserved (i.e., latent). First, Sec. 4.1 provides motivation for using time series of creditworthiness information to estimate default contagion parameters. Then, Sec. 4.2 outlines the assumptions of our proposed estimation framework. Finally, Sec. 4.3 outlines a maximum likelihood approach in which the default contagion parameters are iteratively estimated for all obligors along the default contagion network.

4.1. Motivation for time series of creditworthiness information

In principle, default contagion effects only truly arise in the event of a default. Naturally, attempts have been made to estimate default contagion parameters from historically observed default events (see e.g., Rösch and Winterfeldt, 2008). However, such an approach faces several issues. First, because default events are inherently rare occurrences, the sample of historically observed defaults used for statistical inference is usually small. Crucially, the impact of default contagion is most significant for large and centralized entities, which tend to have a relatively low probability of default. Second, historically observed default events are binary indicators, which only signal whether a default event occurred or not. Consequently, they provide only limited information about default contagion. Finally, another important issue arises due to the heterogeneity of default contagion dependencies across pairs of obligors: contagion effects observed after historical default events are generally not representative of the default dependencies between obligors in the current credit portfolio. Due to these issues, statistical inference based on historically observed default events requires unrealistic homogeneity assumptions.

These issues can be avoided by instead estimating default contagion parameters from creditworthiness information, represented by forward-looking default probabilities. In contrast to rare default events, changes in creditworthiness can occur frequently. For example, credit analysts frequently monitor and update the credit ratings of entities of interest. As another example, liquid CDS markets continuously incorporate new information that reflects the creditworthiness of the underlying entities. Furthermore, creditworthiness represents much more granular information than binary default event indicators. For example, credit analysts generally utilize multiple credit rating buckets and quoted CDS spreads can be considered to be nearly continuous. Finally, the need for homogeneity assumptions can be alleviated by the availability of granular obligor-level creditworthiness information for obligors in the current credit portfolio. The default contagion dependence between two specific obligors can be directly estimated from their historical creditworthiness time series.

To further motivate the use of time series of creditworthiness information, we illustrate how they reflect the underlying default contagion parameters. Consider an obligor $i \in \mathcal{I}$ with a strong dependence on a parent $j \in \mathcal{P}_i$, represented by $W_{ij} \gg 0$. Due to the default contagion mechanism modeled in Eq. (3), a default event of the parent j causes an increase in the default probability PD_i. Consequently, an increase (decrease) in the default probability PD_j of the parent causes an increase (decrease) in PD_i. This illustrates that in the presence of default contagion dependencies, default risk is propagated through creditworthiness information, even without the actual occurrence of a default event. All else being the same, a larger parameter W_{ij} leads to a stronger propagation of the default risk and therefore also to a stronger dependence between the creditworthiness of obligors i and j. It is important to note that this dependence is also partially caused by the shared exposure to common factors. The maximum likelihood estimation of the default contagion parameters can therefore intuitively be interpreted as follows: The default contagion weights W are estimated as the weights that best explain the observed time series of creditworthiness information of all obligors, taking into account the existence of common underlying factors.

4.2. Assumptions of the framework

In this work, we focus on the quantitative estimation of the parameters for known default contagion dependencies. Therefore, it is assumed that for each obligor $i \in \mathcal{I}$, the set of parents $\mathcal{P}_i$ is known *a priori*. Effectively, this assumes that it has already been identified between which obligors there exist default contagion dependencies. In practice, analysts can use both public sources of information and nonpublic information obtained through client onboarding processes. For example, dependencies within corporate groups could be identified via the existence of legal structures or intercompany loans. As another example, analysts could identify the dependencies of an obligor on key customers or suppliers. We note that many banks are expected to identify connected obligors to comply with supervisory guidelines (European Banking Authority, 2017).

Additionally, it is assumed that the factor loading parameter ρ is given *a priori*. Essentially, this assumes that a calibrated factor model is already available for the credit portfolio of interest. This is a reasonable assumption for many financial institutions, as factor models are common practice in industry. For example, most banks use a factor model to calculate (economic) capital requirements for credit risk in the Basel framework, and thus already use factor loading parameter estimates. For example, estimates could be obtained from internal models or from prescribed formulas (Basel Committee on Banking Supervision, 2005). This assumption makes our proposed framework attractive for practical use, as it can be used as an extension to existing factor models that have already been calibrated. Intuitively, the extension captures additional default dependencies that are not yet well captured by the already calibrated factor models.

Moreover, we assume that we have observed the creditworthiness information $\mathrm{pd}_i(t_m)$ for all obligors $i \in \mathcal{I}$ and for all time periods t_m indexed by $m = 1, \ldots, M$. More specifically, it is assumed that the creditworthiness information is represented as forward-looking default probabilities, i.e.,

$$\mathrm{pd}_i(t_m) = \mathbb{P}[\text{Default of obligor } i \text{ at time } T + t_m | V(t_m) = v(t_m)]$$

We note that the frequency at which creditworthiness information is available varies depending on the type of source. For example, CDS prices are often available at a daily or even intra-daily frequency. In contrast, credit ratings obtained from rating agencies or internal bank models are typically available on a monthly basis. However, it is not unlikely that the frequency of the latter will increase significantly in the future as the processes of determining creditworthiness become more automated and data-driven. The impact of the number of historical observations on the performance of the estimation is illustrated in Sec. 5.3.1.

Finally, to estimate the network weights from observed time series of creditworthiness information, we assume that our model is well-specified and stationary over time.

4.3. Maximum likelihood

The default contagion parameter W is estimated by maximizing the likelihood of observed time series of creditworthiness information. We focus on the general case in which the common factor process is latent and therefore propose the use of an expectation-maximization (EM) algorithm (Dempster *et al.*, 1977). However, we also briefly discuss how the algorithm simplifies when the common factor is observed.

Section 4.3.1 describes the EM algorithm, which maximizes the log-likelihood by iterating between an expectation step and a maximization step. Section 4.3.2 outlines details of the expectation step and Sec. 4.3.3 outlines details of the maximization step.

Notation. Before we outline the EM algorithm, we first introduce some notation that is used below. For conciseness, the values corresponding to a time period t_m are denoted with a superscript (m), e.g., $v_i^{(m)} := v_i(t_m)$. Values preceded by Δ indicate (normalized) increments over time, e.g.,

$$\Delta v_i^{(m)} := \frac{v_i(t_m) - v_i(t_{m-1})}{\sqrt{t_m - t_{m-1}}}, \quad \Delta F^{(m)} := \frac{F(t_m) - F(t_{m-1})}{\sqrt{t_m - t_{m-1}}}. \tag{13}$$

Observations are stacked together in matrix notation, where the rows correspond to different time periods, and (if applicable) the columns correspond to different obligors. Moreover, stacked observations corresponding to multiple time periods are denoted in bold to differentiate from the notation used in Secs. 2 and 3, where we did not yet consider multiple time periods. For example, $\boldsymbol{v}$ is a matrix with M rows $v^{(m)} := [v_i^{(m)}]_{i\in\mathcal{I}}$ and N columns $\boldsymbol{v}_i := [v_i^{(m)}]_{m=1}^M$. Similarly, pd is also a matrix of size $M \times N$, $\Delta\boldsymbol{v}$ is a matrix of size $(M-1)\times N$, and $\Delta\boldsymbol{F}$ is a matrix of size $(M-1)\times 1$. Furthermore, a subscript $< i$ or $> i$ represents a grouping of all obligors with a lower or higher index than the obligor i. For example, $\mathbf{pd}_{<i} := [\mathbf{pd}_j]_{j=1}^{i-1}$ and $\mathbf{pd}_{>i} := [\mathbf{pd}_j]_{j=i+1}^{N}$. Similarly, we denote $W_i := [W_{ij}]_{j\in\mathcal{P}_i}$ for all parent weights that correspond to obligor $i \in \mathcal{I}$ and $W_{<i} := [W_j]_{j=1}^{i-1}, W_{>i} := [W_j]_{j=i+1}^{N}$ for the weights that correspond to obligors with a lower or higher index, respectively.

We recall from Sec. 3 that the latent distance-to-default $\boldsymbol{v}$ is not directly observed, but must be inferred from creditworthiness information pd and default contagion parameters W. To make this explicit, we denote the inferred distance-to-default as $\boldsymbol{v}(W, \mathbf{pd})$. We note that for the true default contagion parameter W, it holds that $\boldsymbol{V} \equiv \boldsymbol{v}(W, \mathbf{pd})$.

4.3.1. *Expectation-maximization*

Essentially, EM is an iterative algorithm for computing maximum likelihood estimates in models with latent variables. We refer the reader to the literature for a more comprehensive understanding of EM algorithms and convergence properties (see e.g., Dempster *et al.*, 1977; Wu, 1983; McLachlan and Krishnan, 2008). In particular, we note that the application of EM algorithms is well-established in models with latent factors (see e.g., Rubin and Thayer, 1982).

The considered EM algorithm consists of the following steps. It starts with an initial estimate $\widehat{W}$ for the default contagion parameter W. The algorithm then iteratively updates the estimate by switching between the expectation step and the maximization step until the estimate has converged sufficiently. The expectation step computes the expectation of log-likelihood $\ell(W|\mathbf{pd}, \Delta\boldsymbol{F})$ over the conditional distribution of the latent historical factor increments $\Delta\boldsymbol{F}$, given the historical creditworthiness information pd and the latest estimate $\widehat{W}$. The maximization step updates the estimate $\widehat{W}$ by maximizing the expectation of the log-likelihood with respect to W.

In the context of this work, the EM algorithm has the following intuitive interpretation. Since the latent factor increments $\Delta\boldsymbol{F}$ are not observed, it is not possible to directly evaluate the log-likelihood. Therefore, the expectation step essentially approximates the log-likelihood by substituting the latent factor increments with estimated factor increments. The maximization step then updates the default contagion weights by maximizing the approximated log-likelihood. The algorithm switches between improving the estimate for the latent factor increments and improving the estimate for the default contagion weights.

4.3.2. *Expectation step*

Proposition 3 expresses the conditional expectation of the log-likelihood in terms of the inferred distance-to-default increments $\Delta\boldsymbol{v}(W, \mathbf{pd})$, the estimated factor increments $\mathbb{E}_{\widehat{W}}[\Delta\boldsymbol{F}|\mathbf{pd}]$ and a Jacobian term to account for the transformation to creditworthiness pd. The computation of the estimated factor increments is described in Lemma 1 below and the computation of the Jacobian term is described in Appendix A.3.1.

The decomposition into the components ℓ_i is useful for the proposed iterative approach to the maximization step, as outlined in Sec. 4.3.3.

Proposition 3. *The expectation of the log-likelihood can be decomposed as*

$$\mathbb{E}_{\widehat{W}}[\ell(W|\mathbf{pd}, \Delta\boldsymbol{F})|\mathbf{pd}] = \alpha + \sum_{i\in\mathcal{I}} \ell_i(W), \tag{14}$$

where α is a constant that does not depend on W and

$$\ell_i(W) := -\frac{(\Delta\boldsymbol{v}_i(W) - 2\sqrt{\rho_i}\widehat{\Delta\boldsymbol{F}}(\widehat{W}))^\top \Delta\boldsymbol{v}_i(W)}{2(1-\rho_i)} + \log(\boldsymbol{J}_i(W)), \tag{15}$$

where $\boldsymbol{v}(W) = \boldsymbol{v}(W, \mathbf{pd})$, $\widehat{\Delta\boldsymbol{F}}(\widehat{W}) := \mathbb{E}_{\widehat{W}}[\Delta\boldsymbol{F}|\mathbf{pd}]$, $\boldsymbol{J}_i(W) := \left|\frac{\partial \boldsymbol{v}_i(W,\mathbf{pd})}{\partial \mathbf{pd}_i}\right|$.

Proof. We start by factorizing the joint density $p[\mathbf{pd}, \Delta\boldsymbol{F}]$ to obtain

$$\begin{aligned}\ell(W|\mathbf{pd}, \Delta\boldsymbol{F}) &= \log(p[\mathbf{pd}, \Delta\boldsymbol{F}]) \\ &= \log(p[\mathbf{pd}|\Delta\boldsymbol{F}]) + \log(p[\Delta\boldsymbol{F}]).\end{aligned}$$

Since $p[\Delta\boldsymbol{F}]$ does not depend on W, we can ignore it. Also, $p[\mathbf{pd}|\Delta\boldsymbol{F}]$ can be further factorized over all obligors $i \in \mathcal{I}$ to obtain

$$\log(p[\mathbf{pd}|\Delta\boldsymbol{F}]) = \sum_{i\in\mathcal{I}} \log(p[\mathbf{pd}_i|\Delta\boldsymbol{F}, \mathbf{pd}_{<i}]).$$

For each $i \in \mathcal{I}$, we perform a change of variables from the observed default probabilities $\mathbf{pd}_i$ to the inferred distance-to-default $\boldsymbol{v}_i(W, \mathbf{pd}_i)$ and obtain

$$\log\left(p[\mathbf{pd}_i | \Delta\boldsymbol{F}, \mathbf{pd}_{<i}]\right) = \log\left(p[\boldsymbol{v}_i | \Delta\boldsymbol{F}, \mathbf{pd}_{<i}]\right) + \log\left|\frac{\partial \boldsymbol{v}_i}{\partial \mathbf{pd}_i}\right|. \tag{16}$$

To simplify the first term in Eq. (16), we note the following two points. First, conditional on $\Delta\boldsymbol{F}$, the distance-to-default $\boldsymbol{V}_i$ of obligor i is independent of the creditworthiness information $\mathbf{pd}_{<i}$ of the lower-indexed obligors $j < i$. Second, the distance-to-default $\boldsymbol{V}_i$ has independent increments. Hence, we have that

$$\begin{aligned}\log\left(p[\boldsymbol{v}_i | \Delta\boldsymbol{F}, \mathbf{pd}_{<i}]\right) &= \log\left(p[\boldsymbol{v}_i | \Delta\boldsymbol{F}]\right)\\ &= \log\left(p[\Delta\boldsymbol{v}_i | \Delta\boldsymbol{F}]\right).\end{aligned}$$

Noting the conditional distribution $[\Delta\boldsymbol{V}_i | \Delta\boldsymbol{F}] \sim \mathcal{N}\left(\sqrt{\rho_i}\Delta\boldsymbol{F}, (1-\rho_i)I\right)$ of the $(M-1)$-dimensional column vector $\Delta\boldsymbol{V}_i$, we get that the log-density equals

$$\begin{aligned}\log\left(p[\Delta\boldsymbol{v}_i | \Delta\boldsymbol{F}]\right) &= -\frac{M-1}{2}\log\left(2\pi(1-\rho_i)\right)\\ &\quad - \frac{\left[\Delta\boldsymbol{v}_i - \sqrt{\rho_i}\Delta\boldsymbol{F}\right]^\top\left[\Delta\boldsymbol{v}_i - \sqrt{\rho_i}\Delta\boldsymbol{F}\right]}{2(1-\rho_i)}\\ &= \alpha_i - \frac{\left(\Delta\boldsymbol{v}_i - 2\sqrt{\rho_i}\Delta\boldsymbol{F}\right)^\top\Delta\boldsymbol{v}_i}{2(1-\rho_i)},\end{aligned}$$

where α_i is a constant that does not depend on $\Delta\boldsymbol{v}_i$ and thus not on W.

Finally, Eq. (15) follows from taking the expectation over the conditional distribution of $\Delta\boldsymbol{F}$, which does not affect the Jacobian term. □

Estimated factor increments. Lemma 1 computes the expectation of the latent factor increments $\Delta\boldsymbol{F}$, given creditworthiness information $\mathbf{pd}$ and default contagion parameter W. Essentially, the latent factor increments are estimated as a weighted average of the inferred distance-to-default increments $\Delta\boldsymbol{v}_i$ over all obligors $i \in \mathcal{I}$, where more weight is given to obligors with a higher factor loading $\sqrt{\rho_i}$. Intuitively, the estimate becomes more accurate as the number of obligors N increases, as the idiosyncratic components of the distance-to-default increments average out over a large group of obligors. This is illustrated by the numerical results discussed in Sec. 5.3.2.

We note that if the factor increments are directly observed, the expectation step becomes redundant and the EM algorithm simplifies to a standard maximum likelihood approach. More specifically, the estimated factor increments can then simply be replaced by the observed increments $\Delta\boldsymbol{F}$ in all the steps of the estimation procedure.

Lemma 1. *The conditional expectation* $\widehat{\Delta \boldsymbol{F}}(W) := \mathbb{E}_W[\Delta \boldsymbol{F}|\mathbf{pd}]$ *of the latent factor increments* $\Delta \boldsymbol{F}$ *given the observed creditworthiness information* $\mathbf{pd}$ *and default contagion parameter W equals*

$$\widehat{\Delta \boldsymbol{F}}(W) = \Delta \boldsymbol{v}(W, \mathbf{pd}) D_{1-\rho}^{-1} \sqrt{\rho}^{\top} \left(1 + \sqrt{\rho} D_{1-\rho}^{-1} \sqrt{\rho}^{\top}\right)^{-1}, \tag{17}$$

where $\sqrt{\rho} = [\sqrt{\rho_i}]_{i \in \mathcal{I}}$ *is a row vector and* $D_{1-\rho}$ *is a diagonal matrix with elements* $D_{ii} = 1 - \rho_i$.

Proof. We recall that given the default contagion parameter W, there exists a one-to-one mapping between the creditworthiness $\mathbf{pd}$ and the latent distance-to-default $\boldsymbol{V} \equiv \boldsymbol{v}(W, \mathbf{pd})$. Hence $\mathbb{E}_W[\Delta \boldsymbol{F}|\mathbf{pd}] = \mathbb{E}[\Delta \boldsymbol{F}|\boldsymbol{V} = \boldsymbol{v}(W, \mathbf{pd})]$.

To compute the latter expectation, we first focus on a given historical time point t_m. It follows from Eq. (1) that the distance-to-default increments satisfy $\Delta V^{(m)} = \sqrt{\rho} \Delta F^{(m)} + \Delta \xi^{(m)} D_{\sqrt{1-\rho}}$, where $D_{\sqrt{1-\rho}}$ is a diagonal matrix with elements $[\sqrt{1-\rho_i}]_{i \in \mathcal{I}}$. Since the (normalized) increments all have unit variance, we have

$$\left[\Delta F^{(m)} \Delta V^{(m)}\right] \sim \mathcal{N}\left(0, \begin{bmatrix} 1 & \sqrt{\rho} \\ \sqrt{\rho}^{\top} & D_{1-\rho} + \sqrt{\rho}^{\top} \sqrt{\rho} \end{bmatrix}\right),$$

It is a well-known result on conditional Gaussian distributions Bishop (see e.g., 2006, p. 93) that

$$\Delta F^{(m)} | \Delta V^{(m)} \sim \mathcal{N}\left(\Delta V^{(m)} D_{1-\rho}^{-1} \sqrt{\rho}^{\top} \Sigma, \Sigma\right),$$

where $\Sigma = (1 + \sqrt{\rho} D_{1-\rho}^{-1} \sqrt{\rho}^{\top})^{-1}$. Furthermore, because the increment pairs $\left\{\left(\Delta F^{(m)}, \Delta V^{(m)}\right)\right\}_{m=2}^{M}$ are independent over different time periods, we obtain

$$\begin{aligned} \mathbb{E}[\Delta F^{(m)} | \boldsymbol{V} = \boldsymbol{v}(W, \mathbf{pd})] &= \mathbb{E}[\Delta F^{(m)} | \Delta V^{(m)} = \Delta v^{(m)}(W, \mathbf{pd})] \\ &= \Delta v^{(m)}(W, \mathbf{pd}) D_{1-\rho}^{-1} \sqrt{\rho}^{\top} \Sigma. \end{aligned} \tag{18}$$

Finally, the result in Eq. (17) follows from row-wise stacking of the result in Eq. (18) for different time periods m. □

4.3.3. *Maximization step*

The maximization step of the EM algorithm consists of updating an old parameter estimate $\widehat{W}^{(\text{old})}$ to the new parameter estimate $\widehat{W}^{(\text{new})}$, by maximizing the conditional expectation of the log-likelihood as in Proposition 3.

Iterative maximization. In large credit portfolios, the default contagion parameter W is a sparse and high-dimensional matrix. Consequently, direct optimization of the conditional log-likelihood with respect to W can be numerically unstable. Therefore, we instead propose an optimization approach in which the default contagion parameters $W_i = [W_{ij}]_{j \in \mathcal{P}_i}$ are iteratively estimated for each obligor $i \in \mathcal{I}$ along the default contagion network.

More specifically, we use the decomposition in Proposition 3, to iteratively maximize each log-likelihood component $\ell_i(W)$ with respect to W_i. As explained in Appendix A.3.2, each component $\ell_i(W) = \ell_i(W_i, W_i)$ depends only on the default contagion parameters corresponding to the obligors $j \leq i$. Therefore, we can iteratively compute

$$\widehat{W}_i = \arg\max_{W_i} \ \ell_i(W_i, \widehat{W}_{<i}), \tag{19}$$

where $\widehat{W}_{<i} := [\widehat{W}_j]_{j=1}^{i-1}$ and $\ell_i(W_i, W_{<i})$ is as defined in Eq. (15).

Since each obligor $i \in \mathcal{I}$ usually has at most a few parents, the maximization problem in Eq. (19) is low-dimensional and can be solved using standard gradient-based numerical optimization techniques. Hence, the iterative approach effectively avoids any numerical issues that arise due to W being sparse and high-dimensional. See Appendix A.4 for a derivation of the gradient of $\ell_i(W_i, W_{<i})$ with respect to W_i.

We note that an alternative approach could be to assume a parametric form $W = W(\theta, X)$, where θ is a lower-dimensional parameter and X represents exogenous explanatory variables, e.g.,business volumes or relationship type indicators. In such an approach, the expected log-likelihood could be optimized directly with respect to θ. However, this work instead focuses on the most challenging case in which W is fully heterogeneous and no additional explanatory variables are given.

5. Application

In this section, we demonstrate the application of our proposed framework and illustrate important insights that are relevant to risk management. First, Sec. 5.1 describes the considered credit portfolio and default contagion network. Second, Sec. 5.2 discusses the impact of default contagion on portfolio risk measures and illustrates that they are well estimated by our framework. Third, Sec. 5.3 illustrates the convergence of estimating the individual default contagion parameters and the latent common factor.

5.1. The considered portfolio

To showcase our methodological contribution, we consider a network of heterogeneous obligors that highlights the generality of our estimation framework, as the structure is beyond the scope of existing estimation-focused frameworks restricted to primary-secondary networks (e.g., Rösch and Winterfeldt, 2008; Schiphorst *et al.*, 2024). Overall, our focus is on insights that are robust to changes to the configuration.

5.1.1. *Connected tree structures*

As explained in Sec. 2.3.1, singly connected networks generalize many network structures considered in earlier literature and can capture many important network characteristics that may be encountered in practical applications. These include multilayered hierarchical dependence structures, sparse and heterogeneous network connections, and obligors with multiple parents and/or children. To reflect this, we consider a network of k tree structures that are connected by three central entities. Figure 3 displays a visual example of the network structure for $k = 4$. Each tree consists of a "root" obligor with three children that all have two children themselves. The three central entities are common parents for all the root obligors and thus common ancestors of all other obligors. In total, the network consists of $3 + 10k$ obligors.

Moreover, we represent the heterogeneity of the connected obligors through the model parameters. We consider independent random network weights $W_{ij} \sim \text{Unif}(1.0, 1.5)$ from a range of values inspired by earlier literature (e.g., Neu and Kühn, 2004; Egloff *et al.*, 2007; Agca *et al.*, 2022). Also, we consider credit risk parameters that are representative of a typical credit portfolio. The central entities $i = 1, 2, 3$ have a relatively low default probability $\text{pd}_i = 0.005$, a factor

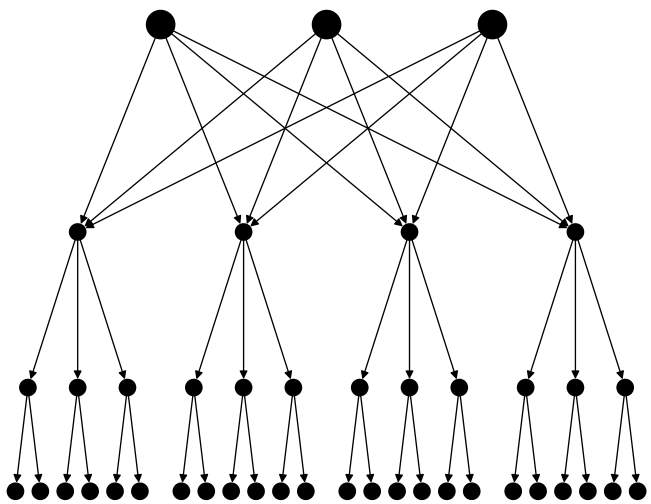

Fig. 3. Example of the network structure with $k = 4$ connected trees.

loading parameter $\rho_i = 0.15$, and an exposure lost at default $e_i = \text{LGD}_i \cdot \text{EAD}_i = 1000$. Similarly, all other obligors $i > 3$ have independently drawn uniform random parameters $\text{pd}_i \sim \text{Unif}(0.005, 0.01)$, $\rho_i \sim \text{Unif}(0.1, 0.2)$, and $e_i \sim \text{Unif}(100, 1000)$. The default probabilities represent credit ratings ranging from lower investment grade to upper noninvestment grade, the range of factor loadings roughly corresponds to the parameters for corporate and sovereign exposures prescribed in the Basel framework (Basel Committee on Banking Supervision, 2005), and the exposures lost at default represent heterogeneous loan sizes that can differ in order of magnitude.

We emphasize that the insights discussed below are robust to changes in this specific configuration. In particular, we note that the exposures lost at default have no impact on the estimation of default contagion parameters, whereas the factor loadings have only a limited impact: Lower factor loadings lead to a more noisy estimation of the latent factor increments, but simultaneously correspond to a lower overall impact of the common factor. Moreover, similar results are obtained for other practically relevant values for the default probabilities and network weights. Finally, we note that the performance of the EM algorithm does depend on the number of time periods of observed creditworthiness information and the number of obligors in the portfolio, as illustrated in Sec. 5.3.

5.1.2. *Central entities*

The central entities represent the well-established empirical observation that most economic networks tend to have heavy-tailed degree distributions (see e.g., Bacilieri *et al.*, 2025), which essentially means that a small number of entities have a disproportionally large number of connections. Networks with similar central entities have also been considered in earlier literature on estimating the impact of default contagion on portfolio risk measures. For example, (Egloff *et al.*, 2007) consider a so-called "Heavy Gravity Portfolio" with two obligors that act as central firms in a supply chain structure and (Anagnostou *et al.*, 2018) consider the default event of a sovereign issuer that affects many other domestic entities.

In the considered network, the three central entities introduce a significant risk concentration, as they pose a common risk source for all other obligors in the network. A default event of one of the central entities causes default contagion effects that propagate throughout the network, significantly increasing the default probability of all other obligors and thus causing large portfolio credit losses. In contrast, if none of the central entities default, all k tree structures are conditionally independent given the common factor and the impact of default contagion on portfolio losses is relatively limited. Although the impact of a default event of a central entity can be quite significant, their default probability often tends to be relatively small. Consequently, their default contagion impact is expected to be concentrated mainly in the tail of the probability distribution of portfolio credit losses. This characteristic makes networks with central entities especially relevant for portfolio credit risk.

5.1.3. *Simulated time series*

As discussed in Sec. 4, our proposed framework estimates default contagion parameters from creditworthiness information in the form of time series of forward-looking default probabilities. Since our contribution is primarily methodological, we conduct our analysis based on simulated data. This avoids the need for proprietary datasets and allows us to explore the impact of significantly increasing the number of observations. We briefly outline the steps to generate the time series of creditworthiness information corresponding to the considered portfolio of connected obligors.

First, the forward-looking default probabilities $\mathrm{pd}_i^{(1)}$ of all obligors $i \in \mathcal{I}$ in the first time period are set equal to those specified in Sec. 5.1.1. Then Algorithm 1 is used to calibrate the initial distance-to-default values $V_i^{(1)}$ for all obligors $i \in \mathcal{I}$. Using these as a starting point, the model specification in Eq. (1) is used to simulate the joint distance-to-default $V^{(m)}$ for all subsequent time periods $m = 2, \ldots, M$. Finally, for each simulated distance-to-default $V_i^{(m)}$, the forward-looking 1-year default probability $\mathrm{pd}_i^{(m)}$ is computed using the techniques described in Sec. 3.

5.2. Portfolio risk measures

We illustrate the impact of default contagion on portfolio risk measures and show the results of our proposed estimation approach.

5.2.1. *The approach*

We compute the VaR and ES risk measures, as defined in Eqs. (4) and (5), respectively, using the approach described in Sec. 2.3.2. We do this for the considered credit portfolio with $k = 10$ trees (i.e., $N = 103$ obligors), for varying quantiles q and for three different sets of default contagion weights W. The first set of weights W_{true} is as specified in Sec. 5.1 and represents the presence of default contagion. The second set of weights $W_{\text{naive}} = 0$ represents the (naive) assumption of no default contagion. Finally, the third set of weights $\widehat{W}_{\text{est}} = \widehat{W}$ is estimated from the generated data described in Sec. 5.1.3. Figures 4 and 5 display the results for the VaR and ES risk measures, respectively.

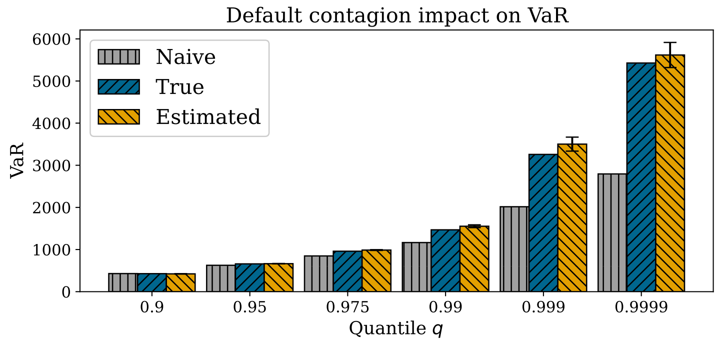

Fig. 4. A comparison of the Value-at-Risk (VaR) for varying quantiles q and for default contagion weights $W_{\text{naive}}, \widehat{W}_{\text{true}}, \widehat{W}_{\text{est}}$. The estimation of $\widehat{W}_{\text{est}}$ is based on $M = 25$ time periods of generated creditworthiness information.

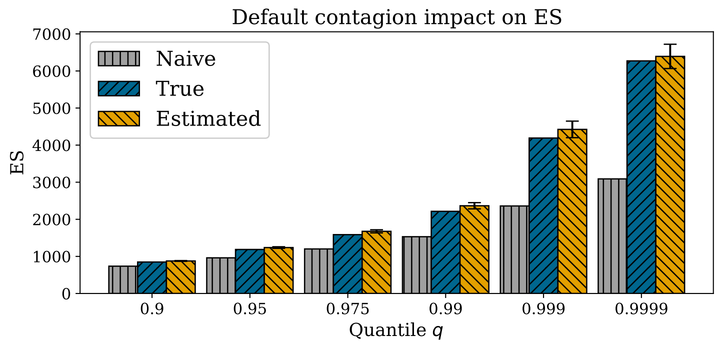

Fig. 5. A comparison of the Expected Shortfall (ES) for varying quantiles q and for default contagion weights $W_{\text{naive}}, \widehat{W}_{\text{true}}, \widehat{W}_{\text{est}}$. The estimation of $\widehat{W}_{\text{est}}$ is based on $M = 25$ time periods of generated creditworthiness information.

5.2.2. *Default contagion impact*

The impact of default contagion on portfolio risk measures is represented by the difference between the risk measures computed with default contagion weights W_{true} and W_{naive}. Overall, the results illustrate that default contagion can have a significant impact on both VaR and ES risk measures. This has also been observed in earlier literature (see e.g., Egloff *et al.*, 2007; Anagnostou *et al.*, 2018; Schiphorst *et al.*, 2024) and highlights the importance of default contagion in portfolio credit risk management.

Interestingly, the results also indicate that the impact of default contagion is greater for higher quantiles. For example, the (relative) impact on VaR is minimal for the quantile $q = 0.9$, but highly material for the quantile $q = 0.9999$. Since default contagion is triggered by default events, which are rare occurrences, the impact tends to be concentrated in the tail of the probability distribution. As explained in Sec. 5.1, this effect is amplified by the presence of central entities.

Furthermore, we note that the (relative) impact of default contagion is greater for ES_q than for VaR_q. This is because ES_q, as defined in Eq. (5), is essentially the average VaR over all higher quantiles in the interval $(q, 1)$, for which the impact is greater.

5.2.3. *Estimated risk measures*

The accuracy in estimating risk measures in the presence of default contagion is represented by the difference between the risk measures corresponding to W_{true} and $\widehat{W}_{\mathrm{est}}$. The displayed error bars represent the standard deviation of the risk measure estimate in the case of estimated W.

Overall, the results illustrate that the use of estimated default contagion weights leads to a good approximation of risk measures and provides a substantial improvement over the naive assumption of no default contagion. Importantly, the risk measures are accurately estimated even for extreme losses corresponding to higher quantiles q, which are not typically observed in historical data. This reinforces the arguments in Sec. 4.1 for estimating default contagion from creditworthiness information, which provides a signal for contagion risk even in the absence of actual default events.

In particular, the results illustrate that using only $M = 25$ time periods of observed creditworthiness information, e.g., roughly two years at a monthly frequency, is already sufficient for a reasonable estimate of risk measures. These relatively low data requirements make the proposed framework attractive even for practical applications in which no daily data is available, as is often the case for risk management of standard credit portfolios. As discussed in Sec. 4.3.3, the estimation accuracy could be improved further by assuming a parametric form for W and incorporating additional information from exogenous explanatory variables.

We note that the data requirements for accurately estimating portfolio risk measures are substantially lower than for accurately estimating individual default contagion weights in W, as discussed in Sec. 5.3.1. This coincides with the observations in (Schiphorst *et al.*, 2024) and is well explained by the following quote Neu and Kühn (2004, p. 14): "*Only statistical properties of the model parameters — not individual parameters themselves — must be captured with reasonable accuracy in order to allow reliable predictions about its global statistical properties such as, indeed, the properties of loss distributions.*"

5.3. EM algorithm

The EM algorithm, described in Sec. 4.3, estimates the default contagion parameters and the latent factor increments. We illustrate that the estimation errors converge to zero when increasing the number of observed time periods and the number of obligors, respectively.

5.3.1. *Default contagion weights*

The results in Fig. 6 illustrate that the root mean square error (RMSE) in estimating the individual weights in W converges to zero as the number of observed time periods M increases. In particular, the logarithmic plot indicates that the convergence rate is close to the order of $M^{-0.5}$, which is expected for maximum likelihood estimation. This confirms that the proposed EM algorithm can accurately estimate individual default contagion weights for a heterogeneous credit portfolio, even if the common factor is not observed and no additional explanatory variables are used.

The results are shown for two distinct cases: the completely heterogeneous network represented by W_{true} and a completely homogeneous network with a single common weight $w \sim \text{UNIF}(1.0, 1.5)$ for all network connections. The difference between the two cases illustrates the trade-off between model flexibility and data requirements for accurate estimation of individual weights. In practical applications, one could of course choose a model specification that lies between these two extremes.

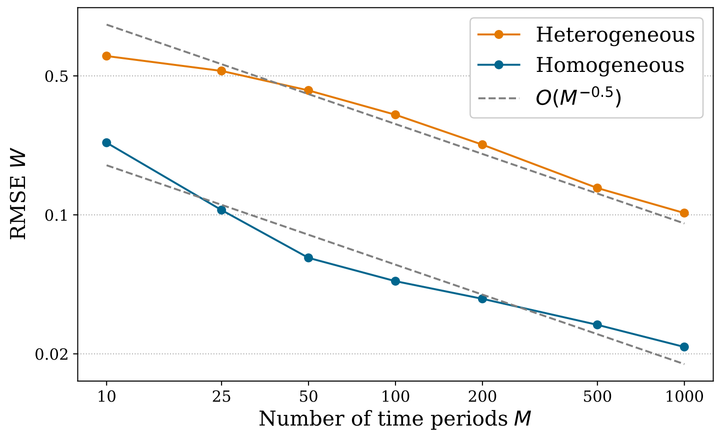

Fig. 6. A log–log plot with the RMSE of the estimated default contagion weights $\widehat{W}$ for a varying number of time periods M of observed creditworthiness information and for two distinct cases: the completely heterogeneous network represented by W_{true} and a completely homogeneous network with a single common weight $w \sim \text{UNIF}(1.0, 1.5)$ for all network connections.

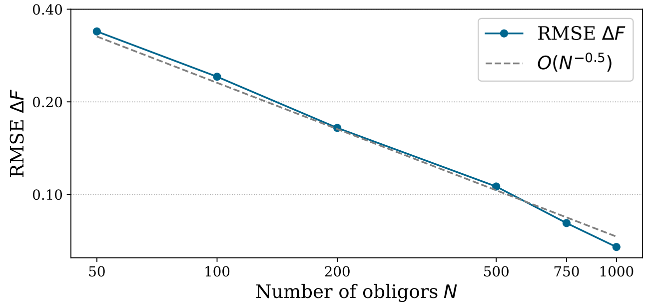

Fig. 7. A log–log plot with the RMSE in the estimated factor increments $\widehat{\Delta \boldsymbol{F}}$ for varying numbers of obligors, estimated from the distance-to-default increments inferred from the generated creditworthiness information and the true weights W_{true}.

5.3.2. *Latent factor*

The results in Fig. 7 illustrate that the error in estimating the latent factor increments $\Delta \boldsymbol{F}$ converges to zero as the number of obligors increases. In particular, the logarithmic plot indicates a convergence to zero at a rate that is close to the order of $N^{-0.5}$.

This confirms that the expectation step of the proposed EM algorithm can accurately estimate the latent common factor increments from the distance-to-default increments. As explained in Sec. 4.3.2, the estimate $\widehat{\Delta \boldsymbol{F}}$ is essentially a weighted average of the inferred distance-to-default of all obligors and therefore becomes more accurate as the number of obligors increases.

6. Conclusion

In this work, we consider a novel framework for modeling and estimating network-based default contagion in credit portfolios, which can be layered onto existing structural factor models. Motivated by the practical relevance and empirical evidence of default contagion, our framework addresses an important gap in methodological literature, where existing frameworks for estimating default contagion parameters are restricted to primary-secondary network structures (see e.g., Rösch and Winterfeldt, 2008; Batiz-Zuk *et al.*, 2015; Lee and Poon, 2014; Schiphorst *et al.*, 2024). By focusing on singly connected networks, our framework strikes a balance between being applicable to many important dependence structures, while still allowing for tractable computations and estimation of default contagion parameters.

From an empirical perspective, we explain that singly connected networks can represent many observed dependence structures, such as corporate groups, governments, and supply chains. Crucially, our framework can capture higher-order propagation of default contagion effects, which is typically explicitly excluded through primary-secondary network assumptions. This highlights the practical relevance of our methodological contributions. From a theoretical perspective, Proposition 2 shows that singly connectedness is equivalent to the property of every obligor having conditionally independent parent obligors. This property significantly simplifies computations and constitutes the most important motivation for our decision to exclude more general networks with loops, such as interbank lending networks. We use this result to calibrate the latent distance-to-default of heterogeneous obligors to exogenously given default probabilities through an iterative algorithm. Building on this, we propose an EM algorithm to estimate the contagion parameters from historical creditworthiness information via the inferred latent distance-to-default processes. This estimation approach allows for a completely heterogeneous default contagion network and works even when the common factor is not directly observed.

We showcase our methodological contributions through an empirically motivated numerical experiment, with a heterogeneous credit portfolio and a rich network structure that incorporates higher-order default contagion effects. The results illustrate that default contagion can have a significant impact on portfolio risk measures, highlighting its importance for portfolio credit risk management. Moreover, we observe and explain that the impact of default contagion tends to be concentrated primarily in the tail of the probability distribution of portfolio losses. This further motivates our estimation approach based on creditworthiness information, which offers a frequent and granular signal of default contagion, even in the absence of actual default events. Indeed, the simulation-based results indicate that our framework can accurately estimate portfolio risk measures in the presence of default contagion from a relatively limited amount of data.

In addition, the results show that our estimation approach can accurately estimate individual default contagion weights even in the most challenging case: a completely heterogeneous network without the use of additional explanatory variables. However, we demonstrate that homogeneity assumptions can substantially reduce the amount of data required for an accurate estimation of individual weights. In practice, one could choose a model specification between the two extremes of complete homogeneity and complete heterogeneity, as a trade-off between data requirements, model flexibility, and computational feasibility.

Acknowledgments

The authors thank the anonymous reviewers for their constructive comments and helpful suggestions that helped improve this manuscript.

Appendix A.

A.1. Numerical integration

We outline the proposed methodology for determining the pairs $\{(\omega_r, F_r)\}_{r=1}^R$ of weights ω_r and nodes F_r for the numerical integration approximation in Eq. (11). We are specifically interested in weights that satisfy $\omega_r > 0$ and $\sum_{r=1}^R \omega_r = 1$, as those properties are used below in Lemma 2.

We use a simple yet effective approach that could be considered a special case of univariate Quasi-Monte Carlo (see e.g., Asmussen and Glynn, 2007, p. 265). Since $F \sim \mathcal{N}(0, T)$, we can use the inversion method to obtain $F \sim \sqrt{T}\Phi^{-1}(U)$, where $U \sim \text{Unif}(0,1)$ and Φ is the standard Gaussian CDF. We then choose values $u_r = \frac{r}{R+1}$ that uniformly spread over the interval (0,1) to obtain

$$\omega_r = \frac{1}{R}, \quad F_r = \sqrt{T}\Phi^{-1}(u_r) = \sqrt{T}\Phi^{-1}\left(\frac{r}{R+1}\right). \tag{A.1}$$

If the default contagion network is singly connected, it follows from Propositions 1 and 2 that the numerical integration approximation $\widehat{\text{PD}}_i$ is given by

$$\begin{aligned}\widehat{\text{PD}}_i(v) &:= \sum_{r=1}^R \omega_r \text{PD}_{i|F_r}(v) \\ &= \sum_{r=1}^R \omega_r \sum_{d\in\{0,1\}^{|\mathcal{P}_i|}} \Phi\big(\theta_{i,F_r,d}(v_i)\big) \prod_{j\in\mathcal{P}_i} \text{PD}_{j|F_r}{}^{d_j}(1-\text{PD}_{j|F_r})^{1-d_j},\end{aligned} \tag{A.2}$$

where

$$\theta_{i,F,d}(v_i) := \frac{\sum_{j\in\mathcal{P}_i} W_{ij} d_j - v_i - \sqrt{\rho_i}F}{\sqrt{1-\rho_i}\sqrt{T}}.$$

While we propose a univariate Quasi-Monte Carlo approach for our single-factor model context, full Monte Carlo approaches could be considered for higher-dimensional factor model extensions, possibly in conjunction with variance reduction techniques such as importance sampling.

A.2. Inferring the distance-to-default

We present some details for numerically solving Eq. (12) to infer the latent distance-to-default v_i from an exogenously given default probability pd_i. Throughout this subsection, we consider a fixed obligor $i \in \mathcal{I}$ and fixed numerical integration pairs $\{(\omega_r, F_r)\}_{r=1}^R$. Also, we assume that $\text{PD}_{j|F_r}$ is given for all $j \in \mathcal{P}_i$ and $r = 1, \ldots, R$, as would be the case in Algorithm 1.

A.2.1. *Existence and uniqueness*

Lemma 2 shows that there exists a unique distance-to-default for any given default probability $\text{pd}_i \in (0,1)$.

Lemma 2. *Let* $\text{PD}_i(\cdot)$ *be as specified in Eq.* (7) *and let* $\widehat{\text{PD}}_i(\cdot)$ *be as specified in Eq.* (A.2). *Then, for any given* $\text{pd}_i \in (0,1)$, *there exists a unique solution* v_i *to the equation* $\text{PD}_i(v_i) \equiv \text{pd}_i$ *and a unique solution* $\hat{v}_i$ *to the* (*approximated*) *equation* $\widehat{\text{PD}}_i(\hat{v}_i) \equiv \text{pd}_i$.

Proof. We focus on proving the existence and uniqueness of the solution $\hat{v}_i$ that solves $\widehat{\text{PD}}_i(\hat{v}_i) \equiv \text{pd}_i$. The proof for the (nonapproximated) solution v_i that solves $\text{PD}_i(v_i) \equiv \text{pd}_i$ is analogous, with the summations over the numerical integration nodes replaced by the corresponding integrals.

We first note that each $\theta_{i,F_r,d}$ is strictly decreasing in v_i. It follows that each $\text{PD}_{i|F_r}$ is also strictly decreasing in v_i since the conditional default probabilities $\text{PD}_{j|F_r}$ of the parents $j \in \mathcal{P}_i$ are positive and do not depend on v_i. Finally, it follows that $\widehat{\text{PD}}_i$ is strictly decreasing in v_i, since all the weights ω_r are positive. The latter proves the uniqueness of any solution.

To prove the existence of the solution $\hat{v}_i$, we show that the function $\widehat{\text{PD}}_i$ has range $(0,1)$. We do this by noting that $\widehat{\text{PD}}_i$ is continuous and has lower and upper limit 0 and 1, respectively:

$$\lim_{v_i \to \infty} \widehat{\text{PD}}_i(v_i) = \sum_{r=1}^{R} \omega_r \lim_{v_i \to \infty} \text{PD}_{i|F_r}(v_i) = \sum_{r=1}^{R} \omega_r 0 = 0,$$

$$\lim_{v_i \to -\infty} \widehat{\text{PD}}_i(v_i) = \sum_{r=1}^{R} \omega_r \lim_{v_i \to -\infty} \text{PD}_{i|F_r}(v_i) = \sum_{r=1}^{R} \omega_r 1 = 1,$$

where we used the fact that the numerical integration weights are normalized such that $\sum_{r=1}^{R} \omega_r = 1$. □

Lemma 3 shows that the approximate distance-to-default $\hat{v}_i$ converges to the true distance-to-default v_i as the numerical integration grid becomes finer.

Lemma 3. *Let* v_i *be the solution that solves* $\text{PD}_i(v_i) \equiv \text{pd}_i$ *and let* $\hat{v}_i$ *be the solution that solves* $\widehat{\text{PD}}_i(\hat{v}_i) \equiv \text{pd}_i$, *which implicitly depends on the number of nodes R in the numerical integration approximation. Then, the approximate* $\hat{v}_i$ *converges to the true* v_i *as the numerical integration grid becomes finer*, *i.e.*,*as* $R \to \infty$.

Proof. We consider a fixed $\varepsilon > 0$. Since $\text{PD}_i(\cdot)$ is strictly decreasing in the distance-to-default, we have $\text{PD}_i(v_i + \varepsilon) < \text{pd}_i$ and $\text{PD}_i(v_i - \varepsilon) > \text{pd}_i$. Since $\widehat{\text{PD}}$

$(\cdot)$ converges point-wise to $\text{PD}_i(\cdot)$, it follows that $\widehat{\text{PD}}_i(v_i+\varepsilon) < \text{pd}_i$ and $\widehat{\text{PD}}_i(v_i-\varepsilon) > \text{pd}_i$ for sufficiently large R. Finally, since $\widehat{\text{PD}}_i(\cdot)$ is also strictly decreasing, we have $v_i - \varepsilon < \hat{v}_i < v_i + \varepsilon$. Hence, for sufficiently large R, it holds that $|\hat{v}_i - v_i| < \varepsilon$. □

A.2.2. *Root finding and derivatives*

We outline how to numerically solve the equation $\text{PD}_i(v_i) \equiv \text{pd}_i$. As explained in the proof of Lemma 2, $\text{PD}_i(v_i)$ is strictly decreasing in v_i. Therefore, a simple bracketing approach could be used. However, results are typically obtained faster with gradient-based root-finding algorithms. Therefore, we present some derivatives of PD_i with respect to v_i. In addition, we note that $\text{PD}_i(v_i) = \text{PD}(v_i, W_i)$ also depends on the default contagion weights $W_i = [W_{ij}]_{j\in\mathcal{P}_i}$ and present some derivatives with respect to these weights. The latter are useful for the maximization step of the EM algorithm outlined in Sec. 4.3.3.

We note that the variables v_i and W_{ij} do not affect the conditional default probabilities $\text{PD}_{j|F}$ of the parents $j \in \mathcal{P}_i$. Therefore, if the variables $x_1, \ldots, x_k$ are equal to either v_i or W_{ij}, we obtain that

$$\frac{\partial^k \text{PD}_i(v_i)}{\partial x_1 \ldots \partial x_k} = \int_{-\infty}^{\infty} \sum_{d\in\{0,1\}^{|\mathcal{P}_i|}} \frac{\partial^k \Phi\left[\theta_{i,F,d}(v_i)\right]}{\partial x_1 \ldots \partial x_k} \prod_{j\in\mathcal{P}_i} \text{PD}_{j|F}{}^{d_j}(1-\text{PD}_{j|F})^{1-d_j} \mathrm{d}\mathbb{P}[F].$$

To compute the different derivatives of PD_i, we can thus substitute in any of the following results:

$$\begin{aligned}
\frac{\partial \Phi\left[\theta_{i,F,d}(v_i)\right]}{\partial v_i} &= \frac{-1}{\sqrt{1-\rho_i}\sqrt{T}} \phi\left[\theta_{i,F,d}(v_i)\right],\\
\frac{\partial^2 \Phi\left[\theta_{i,F,d}(v_i)\right]}{\partial {v_i}^2} &= \frac{-\theta_{i,F,d}(v_i)}{(1-\rho_i)T} \phi\left[\theta_{i,F,d}(v_i)\right],\\
\frac{\partial \Phi\left[\theta_{i,F,d}(v_i)\right]}{\partial W_{ij}} &= \frac{d_j}{\sqrt{1-\rho_i}\sqrt{T}} \phi\left[\theta_{i,F,d}(v_i)\right],\\
\frac{\partial^2 \Phi\left[\theta_{i,F,d}(v_i)\right]}{\partial v_i \partial W_{ij}} &= \frac{\theta_{i,F,d}(v_i) d_j}{(1-\rho_i)T} \phi\left[\theta_{i,F,d}(v_i)\right],
\end{aligned}$$

where ϕ is the standard Gaussian PDF. Similarly, the derivatives of $\widehat{\text{PD}}_i$ can be computed by instead using

$$\frac{\partial^k \widehat{\text{PD}}_i(v_i)}{\partial x_1 \ldots \partial x_k} = \sum_{r=1}^{R} \omega_r \sum_{d\in\{0,1\}^{N_i}} \frac{\partial^k \Phi\left[\theta_{i,F_r,d}(v_i)\right]}{\partial x_1 \ldots \partial x_k} \prod_{j\in\mathcal{P}_i} \text{PD}_{j|F_r}{}^{d_j}(1-\text{PD}_{j|F_r})^{1-d_j}.$$

We note that these derivatives have a very similar structure as the function $\widehat{\mathrm{PD}}_i$ itself and can therefore be evaluated without adding a significant amount of computation time.

Finally, we recall from Eq. (9) that in the absence of default contagion, the inferred distance-to-default equals $v_i = \sqrt{T}\Phi^{-1}(\mathrm{pd}_i)$. Based on experimental results, this appears to be an effective initial guess for the root-finding algorithm when numerically inferring the distance-to-default in the presence of default contagion.

A.2.3. *Implicit function*

We note that Lemma 2 naturally induces an implicit function $v_i = v_i(W_i, \mathrm{pd}_i)$ that satisfies

$$\mathrm{PD}_i(v_i(W_i, \mathrm{pd}_i), W_i) = \mathrm{pd}_i,$$

for all possible default contagion parameters $W_i = [W_{ij}]_{j\in\mathcal{P}_i}$ and for all given default probabilities $\mathrm{pd}_i \in (0, 1)$. By using the chain-rule, we obtain

$$\frac{\partial v_i(W_i, \mathrm{pd}_i)}{\partial W_i} = -\left(\frac{\partial \mathrm{PD}_i(v_i(W_i, \mathrm{pd}_i), W_i)}{\partial v_i}\right)^{-1} \frac{\partial \mathrm{PD}_i(v_i(W_i, \mathrm{pd}_i), W_i)}{\partial W_i},$$

and

$$\frac{\partial v_i(W_i, \mathrm{pd}_i)}{\partial \mathrm{pd}_i} = \left(\frac{\partial \mathrm{PD}_i(v_i(W_i, \mathrm{pd}_i), W_i)}{\partial v_i}\right)^{-1}, \tag{A.3}$$

where the derivatives of PD_i are given in A.2.2.

A.3. Evaluating the log-likelihood

We provide some more details for the evaluation of the (conditional expectation of the) log-likelihood in Eq. (14).

A.3.1. *Jacobian term*

The evaluation of the log-likelihood component ℓ_i in Eq. (15) involves computing the Jacobian term $\log(\boldsymbol{J}_i(W))$ corresponding to the transformation from the latent distance-to-default $\boldsymbol{v}_i$ to the observed unconditional default probabilities $\mathbf{pd}_i$. Lemma 4 shows that this term can be expressed in terms of components $h_i^{(m)}(W)$, which are defined as

$$h_i^{(m)}(W) := \frac{\partial \mathrm{PD}_i^{(m)}(v_i^{(m)}(W, \mathbf{pd}), W)}{\partial v_i}, \tag{A.4}$$

and which can be computed using the results in A.2.2. We note that these components are also evaluated for the gradient-based root-finding algorithm used in Algorithm 1 and could thus be reused for evaluating the log-likelihood.

Lemma 4. *The Jacobian term in Eq.* (15) *can be computed as*

$$\log(\boldsymbol{J}_i(W)) = -\sum_{m=1}^{M} \log(|h_i^{(m)}(W)|). \tag{A.5}$$

Proof. We first note that for each time-period t_m the inferred distance-to-default $v_i^{(m)}(W, \mathbf{pd}) = v_i^{(m)}(W, \mathrm{pd}^{(m)})$ only depends on the default probabilities $\mathrm{pd}^{(m)}$ of that same period. Consequently, $\frac{\partial \boldsymbol{v}_i(W, \mathbf{pd})}{\partial \mathbf{pd}_i}$ is a diagonal matrix with diagonal elements that correspond to different time periods $m = 1, \ldots, M$. Therefore, the Jacobian determinant decomposes as

$$\boldsymbol{J}_i(W) := \left| \frac{\partial \boldsymbol{v}_i(W, \mathbf{pd})}{\partial \mathbf{pd}_i} \right| = \prod_{m=1}^{M} \left| \frac{\partial v_i^{(m)}(W, \mathrm{pd}^{(m)})}{\partial \mathrm{pd}_i^{(m)}} \right|.$$

The result in Eq. (A.5) follows from substituting the result in Eq. (A.3) and taking the logarithm. □

A.3.2. *Partial dependence on W*

We recall that the computation of the inferred distance-to-default in Algorithm 1, is performed iteratively for all obligors $i = 1, \ldots, N$, which are ordered as in Eq. (2). For each obligor i, the inferred distance-to-default $\boldsymbol{v}_i$ is computed using the parameters W_i and the inferred distance-to-default $\boldsymbol{v}_{<i}$ of the lower-indexed obligors. Consequently, the computation of the inferred distance-to-default $\boldsymbol{v}_i(W) = \boldsymbol{v}_i(W_i, W_{<i})$ only (directly) depends on the parameters W_i and (indirectly) on the parameters $W_{<i}$ of lower-indexed obligors. Analogously, the same holds for the Jacobian determinant $\boldsymbol{J}_i(W) = \boldsymbol{J}_i(W_i, W_{<i})$. Finally, we note that the log-likelihood component $\ell_i(W)$ defined in Eq. (15) depends on W only via the latent distance-to-default v_i and the Jacobian determinant $\boldsymbol{J}_i$. Hence, we also have $\ell_i(W) = \ell_i(W_i, W_{<i})$.

A.4. Gradient

Lemma 5 computes the derivative of $\ell_i(W_i, W_i)$ with respect to W_i, which can be used for a gradient-based optimization for each iteration of the maximization step, which is specified in Eq. (19). The derivative is decomposed into the derivatives of the inferred distance-to-default increments $\Delta \boldsymbol{v}_i(W, \mathbf{pd})$ and the Jacobian term $\log(\boldsymbol{J}_i(W))$, which are in turn expressed in terms of the derivatives computed in A.2.2.

Lemma 5. *Let* $\ell_i(W) = \ell_i(W_i, W_{<i})$ *be as defined in Eq.* (15), *then the derivative with respect to* W_i *is given by*

$$\frac{\partial \ell_i(W)}{\partial W_i} = -\frac{\left(\Delta \boldsymbol{v}_i - \sqrt{\rho_i}\Delta \widehat{\boldsymbol{F}}\right)^\top}{1-\rho_i}\frac{\partial \Delta \boldsymbol{v}_i}{\partial W_i} - \sum_{m=1}^{M}\left|\frac{1}{h_i^{(m)}}\frac{\partial h_i^{(m)}}{\partial W_i}\right|, \tag{A.6}$$

where

$$\frac{\partial \Delta v_i^{(m)}(W)}{\partial W_i} = \frac{1}{\sqrt{t_m - t_{m-1}}}\left(\frac{\partial v_i^{(m)}(W)}{\partial W_i} - \frac{\partial v_i^{(m-1)}(W)}{\partial W_i}\right), \tag{A.7}$$

and

$$\begin{aligned}\frac{\partial h_i^{(m)}(W)}{\partial W_i} = {} & \frac{\partial^2 \mathrm{PD}_i^{(m)}(v_i^{(m)}(W), W)}{\partial v_i^2}\frac{\partial v_i^{(m)}(W)}{\partial W_i} \\ & + \frac{\partial^2 \mathrm{PD}_i^{(m)}(v_i^{(m)}(W), W)}{\partial v_i \partial W_i}.\end{aligned} \tag{A.8}$$

Proof. Since ℓ_i depends on W_i only through $\Delta \boldsymbol{v}_i(W)$ and $\boldsymbol{J}_i(W)$, we have

$$\frac{\partial \ell_i(W_i)}{\partial W_i} = \frac{\partial \ell_i(W_i)}{\partial \Delta \boldsymbol{v}_i}\frac{\partial \Delta \boldsymbol{v}_i}{\partial W_i} + \frac{\partial \log(\boldsymbol{J}_i(W))}{\partial W_i}.$$

The result in Eq. (A.6) follows from applying the product rule to Eq. (15) to obtain

$$\frac{\partial \ell_i(W_i)}{\partial \Delta \boldsymbol{v}_i} = -\frac{\Delta \boldsymbol{v}_i^\top + \left(\Delta \boldsymbol{v}_i - 2\sqrt{\rho_i}\Delta \widehat{\boldsymbol{F}}\right)^\top}{2(1-\rho_i)} = -\frac{\left(\Delta \boldsymbol{v}_i - \sqrt{\rho_i}\Delta \widehat{\boldsymbol{F}}\right)^\top}{1-\rho_i},$$

and applying the chain rule to Eq. (A.5) to obtain

$$\frac{\partial \log(\boldsymbol{J}_i(W))}{\partial W_i} = -\sum_{m=1}^{M} \left| \frac{1}{h_i^{(m)}(W)} \frac{\partial h_i^{(m)}(W)}{\partial W_i} \right|.$$

Finally, the result in Eq. (A.7) follows directly from the structure in Eq. (13) and the result in Eq. (A.8) follows from applying the chain rule to Eq. (A.4). □

A.5. Algorithms

Algorithm 1. Numerically inferring the distance-to-default $v = [v_i]_{i \in \mathcal{I}}$ from given default probabilities $\mathrm{pd} = [\mathrm{pd}_i]_{i \in \mathcal{I}}$.

Input: Model parameters (ρ, W) and default probabilities $[\mathrm{pd}_i]_{i \in \mathcal{I}}$.
Output: Inferred $v = [v_i]_{i \in \mathcal{I}}$ that solves $\widehat{\mathrm{PD}}_i(v) \equiv \mathrm{pd}_i$ for all $i \in \mathcal{I}$.
Initialisation: Choose the numerical integration pairs $\{(\omega_r, F_r)\}_{r=1}^{R}$ as in Eq. (A.1) and order the obligors $\mathcal{I} = \{1, \ldots, N\}$ as in Eq. (2).
Computation:

```
for i = 1, ..., N do
    if i has no parents (i.e., P_i = ∅) then
        Compute v_i = −√T · Φ^{−1}(pd_i).
        for r = 1, ..., R do
            Compute PD_{i|F_r} using Eq. (8).
        end for
    else if i has parents (i.e., P_i ≠ ∅) then
        Numerically solve for v_i in Eq. (11) (see A.2.2).
        for r = 1, ..., R do
            Compute PD_{i|F_r} using Eq. (10).
        end for
    end if
end for
```

.

Algorithm 2. Estimating default contagion parameters W from historical creditworthiness information $\mathbf{pd}$.

Input: Factor loading parameters $\rho = [\rho_i]_{i\in\mathcal{I}}$, Historical creditworthiness information $\mathbf{pd} = [\mathrm{pd}^{(m)}]_{m=1}^{M}$, and the parent structure $\mathcal{P} = [\mathcal{P}_i]_{i\in\mathcal{I}}$.
Output: Estimated default contagion parameters $\widehat{W} = [\widehat{W}_i]_{i\in\mathcal{I}}$.
Initialisation: Choose the numerical integration pairs $\{(\omega_r, F_r)\}_{r=1}^{R}$ as in Eq. (A.1), order the obligors $\mathcal{I} = \{1,\ldots,N\}$ as in Eq. (2), choose an initial estimate $\widehat{W}^{(0)}$ and set a tolerance $\varepsilon > 0$.
Computation:
Set $\widehat{W}^{(\mathrm{new})} \leftarrow \widehat{W}^{(0)}$.
repeat
Set $\widehat{W}^{(\mathrm{old})} \leftarrow \widehat{W}^{(\mathrm{new})}$.

Expectation step:
Infer $\Delta\boldsymbol{v}(W^{(\mathrm{old})}, \mathbf{pd})$ using Algorithm 1 and Eq. (13).
Compute $\mathbb{E}_{W^{(\mathrm{old})}}[\Delta\boldsymbol{F}|\mathbf{pd}]$ using Lemma 1.

Maximization step:
for $i = 1,\ldots,N$ **do**
Update $\widehat{W}_i^{(\mathrm{new})}$ using Eq. (19).
end for

until $\|\widehat{W}^{(\mathrm{new})} - \widehat{W}^{(\mathrm{old})}\| < \varepsilon$

Set $\widehat{W} \leftarrow \widehat{W}^{(\mathrm{new})}$
return $\widehat{W}$

ORCID

B. Schiphorst https://orcid.org/0009-0000-9786-5649
M. R. H. Mandjes https://orcid.org/0000-0001-6783-4833
P. J. C. Spreij https://orcid.org/0000-0002-6416-6320
E. M. M. Winands https://orcid.org/0009-0008-4965-567X

References

Agca, S, V Babich, JR Birge and J Wu (2022). Credit shock propagation along supply chains: Evidence from the CDS market, *Management Science*, 68, 6506–6538.

Anagnostou, I, S Sourabh and D Kandhai (2018). Incorporating contagion in portfolio credit risk models using network theory, *Complexity*, 2018, 6076–6173.

Asmussen, S and PW Glynn (2007). *Stochastic Simulation: Algorithms and Analysis*, Springer.

Bacilieri, A, A Borsos, P Astudillo-Estévez, M Hoefer and F Lafond (2025). Firm-level production networks: What do we (really) know? INET Oxford Working Paper, No. 2025-14.

Basel Committee on Banking Supervision (2005). An explanatory note on the Basel II IRB risk weight functions, Technical Report, Bank for International Settlements, https://www.bis.org/bcbs/irbriskweight.pdf.

Batiz-Zuk, E, G Christodoulakis and SH Poon (2015). Credit contagion in the presence of non-normal shocks, *International Review of Financial Analysis*, 37, 129–139.

Bishop, CM (2006). *Pattern Recognition and Machine Learning*, Springer.

De Bruyckere, V, M Gerhardt, G Schepens and R Vander Vennet (2013). Bank/sovereign risk spillovers in the european debt crisis, *Journal of Banking & Finance*, 37, 4793–4809.

Dempster, AP, NM Laird and DB Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm, *Journal of the Royal Statistical Society. Series B, Methodological*, 39, 1–38.

Egloff, D, M Leippold and P Vanini (2007). A simple model of credit contagion, *Journal of Banking & Finance*, 31, 2475–2492.

Emery, K and R Cantor (2005). Default correlation among non-financial corporate affiliates, *Journal of Fixed Income*, 15, 87–96.

European Banking Authority (2017). Guidelines on connected clients, https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/large-exposures/guidelines-connected-clients.

Hong, Y (2013). On computing the distribution function for the poisson binomial distribution, *Computational Statistics & Data Analysis*, 59, 41–51.

Hull, JC (2022). *Options, Futures, and Other Derivatives*, Eleventhglobal edn., Pearson.

Jarrow, RA and F Yu (2001). Counterparty risk and the pricing of defaultable securities, *The Journal of Finance*, 56, 1765–1799.

Koller, D and N Friedman (2009). *Probabilistic Graphical Models: Principles and Techniques*, The MIT Press.

Lee, Y and SH Poon (2014). Modeling the credit contagion channel and its consequences via the standard portfolio credit risk model, *Journal of Credit Risk*, 10, 33–62.

Lütkebohmert, E (2008). *Concentration Risk in Credit Portfolios*, Springer.

McLachlan, GJ and T Krishnan (2008). *The EM Algorithm and Extensions*, John Wiley & Sons.

McNeil, AJ, R Frey and P Embrechts (2015). *Quantitative Risk Management: Concepts, Techniques and Tools*, Revised edn., Princeton University Press.

Merton, RC (1974). On the pricing of corporate debt: The risk structure of interest rates, *The Journal of Finance*, 29, 449–470.

Neu, P and R Kühn (2004). Credit risk enhancement in a network of interdependent firms, *Physica A: Statistical Mechanics and its Applications*, 342, 639–655.

Pearl, J (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann.

Rösch, D and B Winterfeldt (2008). Estimating credit contagion in a standard factor model, *Risk*, 21, 78–82.

Rubin, DB and DT Thayer (1982). EM algorithms for ML factor analysis, *Psychometrika*, 47, 69–76.

Schiphorst, B, M Mandjes, P Spreij and E Winands (2024). A structural credit risk model with default contagion, *Proceedings of Mathematical and Statistical Methods for Actuarial Sciences and Finance, (MAF 2024)*, 280–285.

Wu, CFJ (1983). On the convergence properties of the EM algorithm, *The Annals of Statistics*, 11, 95–103.