Contagion in Financial Systems: A Bayesian Network Approach[*]

Carsten Chong[†] and Claudia Klüppelberg[†]

Abstract. We develop a structural default model for interconnected financial institutions in a probabilistic framework. For all possible network structures we characterize the joint default distribution of the system using Bayesian network methodologies. Particular emphasis is given to the treatment and consequences of cyclic financial linkages. We further demonstrate how Bayesian network theory can be applied to detect contagion channels within the financial network, to measure the systemic importance of selected entities on others, and to compute conditional or unconditional probabilities of default for single or multiple institutions.

Key words. Bayesian network, financial contagion, measure of systemic risk, multivariate default risk, probability of default, structural default risk model, systemic risk

AMS subject classifications. 91B30, 62-09, 91G80

DOI. 10.1137/17M1116659

1. Introduction. The 2007–2008 financial crisis has unmistakably revealed on which fragile grounds the global financial system was built at that time. Contrary to the usual perception that diversification reduces risk, it was precisely the interconnectedness of the system that made the spread of shocks across large scales possible. Henceforth, a lot of effort has been put into understanding the origin, the mechanisms, and the consequences of financial contagion. It is still a matter of current research to analyze the complex effects of financial interconnection and whether it is a blessing or a curse for the stability of the overall financial system. For example, while the classical papers by Allen and Gale (2000) and Freixas, Parigi, and Rochet (2000) argue that interconnection strengthens the resilience of financial systems, recent work by Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) and Elliott, Golub, and Jackson (2014) demonstrate the robust yet fragile nature of financial systems: Densely connected systems can prove to be more robust but also to be more vulnerable to shocks depending on the actual network structure. Therefore, a profound understanding of interconnectedness and the risk of financial contagion is inevitable for bankers, regulators, and other decision makers in order to set up effective detection and prevention mechanisms for possible future crises.

In the literature on balance sheet contagion, one of the most prominent default cascade models is that of Eisenberg and Noe (2001). Under mild assumptions on the operative cash flows the financial institutions generate, the authors use a fixed-point argument to establish the existence of a unique clearing payment vector, which can be explicitly computed via a fictitious default algorithm. Extending the model by introducing bankruptcy costs,

[†]Center for Mathematical Sciences, Technical University of Munich, 85748 Garching, Germany (carsten.chong@tum.de, cklu@tum.de, www.statistics.ma.tum.de).

Rogers and Veraart (2013) show that the clearing vector is no longer uniquely determined in general. Instead, the fictitious default algorithm can be modified to produce the greatest and the least clearing vector. Although Eisenberg and Noe (2001) also discuss qualitative implications of stochastic operative cash flows, the core analysis of these two papers remains deterministic.

Imposing maximal bankruptcy costs (i.e., creditors lose the total value of a claim upon default of the counterparty), Gai and Kapadia (2010) show that the relationship between connectedness and stability in financial systems exhibits a phase transition: Interconnection can both absorb and amplify shocks. Analytical results, however, are only obtained on homogeneous random graphs where exposures are evenly distributed among all counterparties. We refer to Hurd (2016) for extensions and comparisons between the Eisenberg–Noe and the Gai–Kapadia model.

A nonmonotonic effect of interconnection on stability is also confirmed by Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015). Drawing on a variant of Eisenberg and Noe's model, they show that for small shocks, highly diversified systems are less susceptible to contagion than sparsely connected networks but that the opposite is true for big shocks. Methodologically, the authors impose a two-state evolution for the firms and identify the optimal network for the small as well as the big shock regime. Assuming that interconnection arises through asset cross-holdings rather than mutual liabilities, Elliott, Golub, and Jackson (2014) derive the same qualitative result on various random graph models.

Another interesting development of the Eisenberg–Noe model is carried out by Gouriéroux, Héam, and Monfort (2012). The authors first extend the existence and uniqueness of a clearing vector to the situation where interfirm share holdings are included in the model. Additional to these (again purely deterministic) considerations, they further discuss the consequences for the financial system when random shocks occur, giving two quantitative results. First, they compute the conditional distribution of the firm values when it is known which firms default and which not. Second, they provide an enumeration-type formula for the probability of default for a given firm in the network. However, they do not reveal how the terms in their expansion (which are the probabilities of joint default configurations) can be computed. Furthermore, it is unclear how to determine which institutions default or not at first place because this in turn depends on the firm values. Finally, let us also mention Glasserman and Young (2015), who give quantitative estimates on the probability of default in various contagion models.

In addition to the mainly theoretical contributions mentioned above, a lot of empirical work has been carried out; see, e.g., Elsinger, Lehar, and Summer (2006), Cont, Moussa, and Santos (2013), Craig and von Peter (2014), and Langfield, Liu, and Ota (2014) for applications to central bank data from the Austrian, Brazilian, German, and UK banking system, respectively. Furthermore, apart from work based on variations of the Eisenberg and Noe (2001) model, there are also other network-based approaches to systemic risk. We refer to Castiglionesi and Eboli (2015) for the usage of flow network techniques, to Fouque and Sun (2013), Kley, Klüppelberg, and Reichel (2015), and Chong and Klüppelberg (2017) for models relying on stochastic differential equations, and to Amini, Cont, and Minca (2016) and Detering et al. (2017) for contagion analyses on large random graphs. While the analysis in Castiglionesi and Eboli (2015) is again deterministic, the other papers rely on probabilistic approaches that only apply to stylized or limiting networks.

Even though some of the aforementioned works discuss stochastic effects in the context of financial contagion, the full structure of the default probabilities within a given (not necessarily stylized) financial system remains unknown when future returns are random. And this is exactly the problem we want to address in the present paper. For all possible network structures of a financial system subjected to stochastic shocks, we give *a complete characterization of the joint default probability distribution* of the system in a way that the probability of default for a single institution or a group of institutions can be computed efficiently. Moreover, our analysis will also allow for identifying possible channels of financial contagion and measuring the systemic importance of institutions.

Our approach will draw upon methods from the theory of *Bayesian networks*, also known as *directed graphical models*; see, e.g., Lauritzen (1996) and Koller and Friedman (2009). These models are a priori defined on *directed acyclic graphs* (DAGs), but it is widely aknowledged that mutual dependencies and directed cycles are omnipresent in financial systems. Therefore, for the purposes of this paper, we have to extend the theory of Bayesian networks to graphs with directed cycles, and we will do so by interpreting a cyclic model as the margin of an acyclic model on a suitably enlarged graph. This methodology for treating cyclic graphical models or, more generally, systems with feedback cycles is of independent interest and may be useful for other applications as well.

The remaining article is organized as follows. In section 2, we introduce a probabilistic structural default model for interconnected financial institutions. As in Gai and Kapadia (2010), Rogers and Veraart (2013), and Elliott, Golub, and Jackson (2014), our model takes bankruptcy costs into account so that the standard balance sheet equations may no longer have a unique solution. Apart from specific examples, a general characterization of uniqueness of solutions seems to be absent in the literature. So this leads to our first main result, Theorem 2.4, which states that uniqueness holds if and only if the underlying liability network corresponds to a DAG.

If the interfirm liabilities exhibit a DAG structure, we further prove in section 3, Theorem 3.2, that the default variables of the institutions form a Bayesian network on this DAG. In other words, if we fix an institution i and assume knowledge about the state of all its debtors (its *parents* in the language of graph theory), its default probability is independent of all other firms that have not lent (directly or indirectly via other firms) to i (the *nondescendants* of i).

If there are directed lending cycles, which is most likely the case in practice, the previously identified nonuniqueness problem requires further specification of the sets of defaulting and surviving firms. For the two extremal cases, the *mild* and the *strict default rule*, we recover in Theorems 3.10 and 3.12 the Bayesian network structure of the default variables on a suitably enlarged liability graph. So for all financial networks—no matter with or without directed cycles—the joint default distribution of the system can be efficiently structured in terms of a Bayesian network.

In section 4, we discuss the practical consequences and advantages of having a Bayesian network structure. We first reveal how to decide on a graphical basis whether the default of two (or two groups of) firms is stochastically dependent or not. Second, we highlight in a numerical study how established algorithms can be used to compute conditional or unconditional default probabilities, either exactly or approximatively. Finally, we define and analyze two measures of systemic risk that follow from our Bayesian network analysis. In contrast to many already

existing measures, ours are probabilistic in nature and follow from a structural default cascade model. Section 5 concludes and points at further research directions.

2. Default risk model for interconnected firms. We consider N firms labeled $i = 1, \dots, N$ within an economy for a single period. At time $t = 0$, each firm i has some operating assets X_i and some cash K_i that yields a riskless continuous interest at a rate of r_0. The same firm also has some external liabilities F_i that are to be redeemed at time $t = T$, the end of the period, with a continuous interest rate of r_0'. In addition to that, at $t = 0$, each firm i has lent firm j an amount of L_{ij} (possibly 0) which is to be repaid by j at time $t = T$ with a continuous interest rate of r_{ij}. In other words, firm i holds a bond of firm j with face value L_{ij} and a single coupon payment at maturity T. Apart from that, no other financial involvement exists between the N firms.

During the period from $t = 0$ to $t = T$ every company i uses its operating assets X_i to run a firm-specific business. Following the classical work of Merton (1974), this is assumed to evolve according to a geometric Brownian motion with drift μ_i and volatility σ_i^2. The driving Brownian motions $B_1, \dots, B_N$ are considered to be independent. We shall comment on possible relaxations of this assumption in section 5.

In order to simplify the subsequent exposition, the values K_i, F_i, and L_{ij} are assumed to be nonnegative; the asset drifts μ_i as well as the interest rates r_{ij}, r_0, and r_0' may take positive or negative values; and X_i, σ_i, and T are strictly positive. Moreover, for each $i \in \{1, \dots, N\}$ we suppose that

$$(2.1) \qquad K_i < e^{(r_0' - r_0)T} F_i + \sum_{j=1}^{N} e^{(r_{ji} - r_0)T} L_{ji}.$$

Firms that have more cash reserves than the bound on the right-hand side of (2.1) will never experience default at time $t = T$. Even if they do not receive payments from the other institutions and lose all their operating assets, they still have enough money to meet all their obligations. Since we are interested in the default probabilities within the financial network, there is no loss of generality if we exclude such institutions from our analysis.

Let us also point out that we take a "central bank" perspective in this paper. All parameters and quantities introduced above are assumed to be known or have been estimated before. We refer to Anand, Craig, and von Peter (2015), Gandy and Veraart (2016), and Upper (2011) for various methods to estimate interfirm exposures.

Based on the previously described setup, the balance sheet of firm i at time $t = 0$ has the form depicted on the left-hand side of Figure 1. In essence, the situation at time $t = 0$ is a variant of the classical model by Eisenberg and Noe (2001). At the time of maturity, however, ours differs in two important aspects from that of Eisenberg and Noe (2001).

First, the value of firm i's operating assets at $t = T$, which is given by

$$X_i(T) = X_i e^{(\mu_i - \sigma_i^2/2)t + \sigma_i B_i(T)},$$

is *random*, whereas the operating cash flows in Eisenberg and Noe (2001) are assumed to be deterministic.

Second, in contrast to a linear or proportional recovery, we impose a *zero recovery rate* in the default scenario; that is, if a firm is not able to meet all of its obligations, none of

Operating assets X_i	Liabilities - external F_i
Loans $\sum_{j=1}^{N} L_{ij}$	- to other firms $\sum_{j=1}^{N} L_{ji}$
Cash K_i	Equity $E_i(0)$
Assets at $t = 0$	Liabilities and equity at $t = 0$

Operating assets $X_i(T)$	Liabilities - external $\mathrm{e}^{r_0'T}F_i$
Loans $\sum_{j=1}^{N} \mathrm{e}^{r_{ij}T} L_{ij}\mathbb{1}_{\mathcal{S}}(j)$	- to other firms $\sum_{j=1}^{N} \mathrm{e}^{r_{ji}T} L_{ji}$
Cash $\mathrm{e}^{r_0T}K_i$	Equity $E_i(T)$
Assets at $t = T$	Liabilities and equity at $t = T$

Figure 1. *Idealized balance sheet of firm i at $t = 0$ and $t = T$.*

its creditors receives any payment. As Gai and Kapadia (2010) argue, "This assumption is likely to be realisitic in the midst of a crisis: In the immediate aftermath of a default, the recovery rate and the timing of recovery will be highly uncertain and banks' funders are likely to assume the worst-case scenario" (p. 2407). We also refer to Battiston et al. (2012) and Cont, Moussa, and Santos (2013) for a similar reasoning and Memmel, Sachs, and Stein (2012) for empirical evidence for the frequent occurrence of a recovery rate close to zero.

Standard balance sheet considerations lead us to the following definition (we do not distinguish between the terms bankruptcy, default, and insolvency in this paper).

Definition 2.1. *A company defaults (survives) at time $t = T$ if it belongs to the set*

$$(2.2) \qquad \mathcal{D} := \{i \in \{1, \ldots, N\} \colon E_i(T) < 0\} \quad \big(\mathcal{S} := \{i \in \{1, \ldots, N\} \colon E_i(T) \geq 0\}\big),$$

where $E_i(T)$ is the equity of firm i at time T given by

$$(2.3) \qquad E_i(T) := X_i(T) + \sum_{j=1}^{N} \mathrm{e}^{r_{ij}T} L_{ij}\mathbb{1}_{\mathcal{S}}(j) + \mathrm{e}^{r_0T}K_i - \mathrm{e}^{r_0'T}F_i - \sum_{j=1}^{N} \mathrm{e}^{r_{ji}T} L_{ji}.$$

Moreover, we define

$$D_i := \mathbb{1}_{\mathcal{D}}(i), \quad S_i := 1 - D_i = \mathbb{1}_{\mathcal{S}}(i), \quad i = 1, \ldots, N.$$

The resulting balance sheets at time $t = T$ are depicted on the right-hand side of Figure 1. As a consequence of the stochasticity of the asset values $X_i(T)$, the equity values $E_i(T)$, the

sets $\mathcal{D}$ and $\mathcal{S}$ and the variables D_i and S_i are random. The natural question arises whether the knowledge of $X_1(T),\ldots,X_N(T)$ at time T fully determines $\mathcal{D}$, $\mathcal{S}$ and $E_i(T)$ through equations (2.2) and (2.3). In the model of Eisenberg and Noe (2001) the existence and uniqueness of solutions to (2.2) and (2.3) would immediately follow from the strict positivity of all $X_i(T)$. This condition, however, is no longer sufficient for uniqueness under the zero recovery assumption as the following simple counterexample demonstrates.

Example 2.2. Suppose there are two firms in the system owing each other \$10 at maturity with neither cash deposits nor external liabilities. If both firms have realized operating assets amounting to $X_1(T) = X_2(T) = \$5$ at $t = T$, there are two different solutions to equations (2.2) and (2.3):
  (1) $\mathcal{D} = \{1,2\}$, $\mathcal{S} = \emptyset$, and $E_1(T) = E_2(T) = -\$5$ and
  (2) $\mathcal{D} = \emptyset$, $\mathcal{S} = \{1,2\}$, and $E_1(T) = E_2(T) = \$5$.

The emergence of multiple solutions in the previous example highlights the remaining point of indeterminancy in our model: To what extent is clearing between distressed firms or, more precisely, the netting of mutual claims allowed? In the first solution of Example 2.2, netting is forbidden, and both firms default because neither of them is able to initiate the promised payment to the other one. In the second solution, netting is permitted, so both firms abandon their claims and survive with \$5 each. Of course, the degree to which netting of claims occurs depends on many external factors, such as the availability of short-term fundings or the degree of counterparty anonymity. However, as we shall see in section 3, it is always possible to find and characterize the solution to (2.2) and (2.3) with the maximal or the minimal number of surviving firms, corresponding to the situation where netting is always permitted or prohibited.

The nonuniqueness issue as encountered in Example 2.2 is common to cascade models with bankruptcy costs; cf. Rogers and Veraart (2013), Elliott, Golub, and Jackson (2014), and Hurd (2016). A general criterion for uniqueness versus nonuniqueness seems to be missing in the literature and hopeless if it is required to hold ω-wise for all possible outcomes of $X_1(T)(\omega),\ldots,X_N(T)(\omega)$.

However, we will show that it is possible to characterize all models where uniqueness holds *almost surely*. The most efficient way to do so is to represent the interfirm liabilities by a graphical structure. In this article, we only consider *directed graphs* $G = (\mathcal{V},\mathcal{E})$, where $\mathcal{V}$ is some finite *vertex set* and $\mathcal{E} \subseteq \{(i,j) : i,j \in \mathcal{V}, i \neq j\}$ some *edge set* without self-loops. The following graphical representation of the interfirm liabilities plays a crucial role in the default analysis of the firm network described above.

Definition 2.3. *Abbreviating* $[N] := \{1,\ldots,N\}$, *we define the* redemption graph $G = (\mathcal{V},\mathcal{E})$ *as the graph on the vertex set* $\mathcal{V} := [N]$ *with edges* $\mathcal{E} := \{(i,j) \in [N]^2 : L_{ji} > 0\}$.

Thus, there is an edge from i to j precisely when i has borrowed money from j at time $t = 0$ and has to repay the latter at time $t = T$. In graphical terms, this is denoted by $i \to_G j$ or $i \in \mathrm{pa}_G(j)$, and we say that i is a *parent* of j in G. A whole sequence $i_1 \to_G \cdots \to_G i_n$ with $n \geq 2$ is then called a *(directed) path* from i_1 to i_n and a *(directed) cycle* if $i_1 = i_n$. If G contains no directed cycles, G is called a DAG.

For future reference, let us also introduce some further terminology. A vertex i is called an *ancestor* of j, and j a *descendant* of i, in short $i \in \mathrm{an}_G(j)$ and $j \in \mathrm{de}_G(i)$, if there

exists a directed path from i to j in G. Nodes in $\mathrm{nd}_G(i) := \mathcal{V} \setminus (\{i\} \cup \mathrm{de}_G(i))$ are called the *nondescendants* of i. Given an additional sequence $(x_i)_{i \in \mathcal{V}}$, we also write $\mathrm{pa}_G(x_i) := (x_j : j \in \mathrm{pa}_G(i))$ and define $\mathrm{an}_G(x_i)$, $\mathrm{de}_G(x_i)$, and $\mathrm{nd}_G(x_i)$ analogously. If $I \subseteq \mathcal{V}$, we set $\mathrm{pa}_G(I) := \bigcup_{i \in I} \mathrm{pa}_G(i)$, $\mathrm{de}_G(I) := \bigcup_{i \in I} \mathrm{de}_G(i)$, $\mathrm{nd}_G(I) := \bigcap_{i \in I} \mathrm{nd}_G(i)$ and $x_I := (x_i : i \in I)$.

Returning to the question of uniqueness of solutions to (2.2) and (2.3), our first main result asserts that the decision can be made by sole inspection of the redemption graph.

Theorem 2.4. *Let G be the redemption graph underlying the considered financial network.*
(1) *If G is a DAG, there exists a unique solution to (2.2) and (2.3) with probability 1, so $\mathcal{D}$, $\mathcal{S}$, and $E_1(T), \ldots, E_N(T)$ are uniquely determined for almost all realizations of $X_1(T), \ldots, X_N(T)$.*
(2) *If G is not a DAG but contains a directed cycle, there is a strictly positive probability that the realizations $X_1(T), \ldots, X_N(T)$ are such that (2.2) and (2.3) have multiple solutions.*

3. Default probabilities for interconnected firms. As shown in Theorem 2.4, directed cycles in the financial system (as observed in Example 2.2) impose the only hindrance to the uniqueness of solutions to (2.2) and (2.3). So if the redemption graph G is a DAG, the random variables $D_1, \ldots, D_N$ are well defined, and in this case, their joint distribution has a particularly convenient structure.

Definition 3.1. *Given a DAG $G = (\mathcal{V}, \mathcal{E})$, a collection $\{X_i : i \in \mathcal{V}\}$ of random variables taking values in a finite set E is said to form a* Bayesian network *over G if for all $e = (e_i : i \in \mathcal{V}) \in E^{|\mathcal{V}|}$ we have*

$$(3.1) \qquad \mathbb{P}[X_\mathcal{V} = e] = \prod_{i \in \mathcal{V}} \mathbb{P}[X_i = e_i \mid \mathrm{pa}_G(X_i) = \mathrm{pa}_G(e_i)].$$

For a detailed treatment of Bayesian networks we refer to the monographs Lauritzen (1996) and Koller and Friedman (2009). An equivalent characterization is this (see Theorem 3.27 in Lauritzen (1996)): $\{X_i : i \in \mathcal{V}\}$ forms a Bayesian network over G if and only if for every $i \in \mathcal{V}$, the variable X_i is conditionally independent of $X_{\mathrm{nd}_G(i)}$ given $X_{\mathrm{pa}_G(i)}$. Thus, a Bayesian network structure can be understood as a collection of conditional independence statements among variables.

Theorem 3.2. *If the redemption graph G is a DAG, the variables $\{D_1, \ldots, D_N\}$ form a Bayesian network on G with conditional probabilities*

$$(3.2) \quad \mathbb{P}[D_i = 1 \mid \mathrm{pa}_G(D_i)] = \Phi(i, \Sigma_i, T), \quad \mathbb{P}[D_i = 0 \mid \mathrm{pa}_G(D_i)] = 1 - \Phi(i, \Sigma_i, T), \quad i \in [N],$$

where

$$(3.3) \qquad \Sigma_i := \{j \in \mathrm{pa}_G(i) : D_j = 0\},$$

$$(3.4) \qquad \Phi(i, \Sigma_i, T) := \Phi\left(-\frac{\log(Q(i, \Sigma_i)) + (\mu_i - \sigma_i^2/2)T}{\sigma_i \sqrt{T}}\right),$$

$$(3.5) \qquad Q(i, \Sigma_i) := \frac{X_i}{\left(e^{r_0' T} F_i + \sum_{j=1}^N e^{r_{ji} T} L_{ji} - e^{r_0 T} K_i - \sum_{j \in \Sigma_i} e^{r_{ij} T} L_{ij}\right)^+}$$

and Φ is the standard normal distribution function, $x^+ := \max(x,0)$ and $a/0 := \infty$ for $a > 0$, $\log(\infty) := \infty$, and $\Phi(-\infty) := 0$.

In particular, for all $\Sigma \subseteq [N]$ we have that

$$(3.6) \qquad \mathbb{P}[\mathcal{S} = \Sigma, \ \mathcal{D} = [N] \setminus \Sigma] = \prod_{i \in [N] \setminus \Sigma} \Phi(i, \Sigma, T) \prod_{i \in \Sigma} (1 - \Phi(i, \Sigma, T)).$$

Let us give an informal explanation why an acyclic redemption graph leads to a Bayesian network structure of $D_1, \ldots, D_N$. Indeed, every DAG G induces a partial order $\lesssim_G$ on its vertex set via the relation

$$(3.7) \qquad i \lesssim_G j :\Longleftrightarrow i \in \mathrm{an}_G(j) \text{ or } i = j.$$

Therefore, starting with institutions that are minimal with respect to $\lesssim_G$ (i.e., institutions that have not lent money to anybody else), we can determine the values of $D_1, \ldots, D_N$ iteratively along the partial order $\lesssim_G$. The Bayesian network structure follows now from the observation that this iteration procedure is "Markovian" with respect to $\lesssim_G$: For each $i \in [N]$, the variable D_i is independent of $(D_j : j \lesssim_G i)$ given $(D_j : j \in \mathrm{pa}_G(i))$, or equivalently, in order to determine the default behavior of i, it suffices to know whether its debtors are solvent or not. The formal proof of Theorem 3.2 can be found in the Appendix.

When the redemption graph has directed cycles, additional rules to (2.2) and (2.3) have to be set up for $D_1, \ldots, D_N$ to be well defined. These rules may be different for each cycle or vary from firm to firm. The following result, however, holds for all *consistent extensions* of (2.2) and (2.3) (i.e., measurable mappings of the values $X_1(T), \ldots, X_N(T)$ to $D_1, \ldots, D_N$ such that (2.2) and (2.3) are valid). Recall that the *subgraph of $G = (\mathcal{V}, \mathcal{E})$ induced by $I \subseteq \mathcal{V}$* is the graph $G_I := (I, \mathcal{E}_I)$ with $\mathcal{E}_I := I^2 \cap \mathcal{E}$.

Proposition 3.3. *Formula (3.6) is valid if both the induced subgraphs G_Σ and $G_{[N] \setminus \Sigma}$ contain no directed cycles. In the general case, (3.6) still holds with a "$\leq$" sign instead of equality.*

Apart from Proposition 3.3, there is not much more that can be said about a particularly chosen solution to (2.2) and (2.3). However, further structural results can be obtained for the two "extreme" cases, that is, where the set of surviving firms is maximal or minimal, respectively.

Definition 3.4. *The* mild default rule *specifies the sets $\mathcal{S}$ and $\mathcal{D}$ in the following way.*
- *If the operating assets and cash holdings of some firm i are so low that it cannot meet its obligations at time $t = T$ even if it is paid back by all its debtors, then i defaults and is called a* defaulting firm of the first round. *In this case we write $i \in \mathcal{D}_1$. If no firm falls into this category, all firms survive and $\mathcal{D} = \emptyset$ and $\mathcal{S} = [N]$.*
- *For $n \in \{2, \ldots, N\}$ we call $i \in [N] \setminus \bigcup_{m=1}^{n-1} \mathcal{D}_m$ a defaulting firm of the nth round, denoted $i \in \mathcal{D}_n$, if at time $t = T$ firm i is insolvent given it receives all contractual payments except for those from their debtors in $\bigcup_{m=1}^{n-1} \mathcal{D}_m$.*
- *As soon as $\mathcal{D}_n = \emptyset$, we set $\mathcal{D}_{n+1} = \cdots = \mathcal{D}_N = \emptyset$, $\mathcal{D} = \bigcup_{m=1}^{n-1} \mathcal{D}_m$, and $\mathcal{S} = [N] \setminus \mathcal{D}$.*

Definition 3.5. *Under the* strict default rule *the sets $\mathcal{S}$ and $\mathcal{D}$ are determined as follows.*
- *Every firm i whose operating assets and cash holdings at time $t = T$ are large enough to repay its creditors (regardless whether it receives payments from the other firms)*

survives. In this case, we call i a surviving firm of the first round *and write $i \in \mathcal{S}_1$. If no firm meets this criterion, all firms default: $\mathcal{S} = \emptyset$ and $\mathcal{D} = [N]$.*

- *For every $n \in \{2, \ldots, N\}$ we call $i \in [N] \setminus \bigcup_{m=1}^{n-1} \mathcal{S}_m$ a* surviving firm of the nth round, *denoted $i \in \mathcal{S}_n$, if at time $t = T$ its operating assets, its cash holdings, and the payments it receives from firms in $\bigcup_{m=1}^{n-1} \mathcal{S}_m$ are high enough to redeem its debts.*
- *As soon as $\mathcal{S}_n = \emptyset$, we set $\mathcal{S}_{n+1} = \cdots = \mathcal{S}_N = \emptyset$, $\mathcal{S} = \bigcup_{m=1}^{n-1} \mathcal{S}_m$ and $\mathcal{D} = [N] \setminus \mathcal{S}$.*

Both the mild and the strict default rules are variants of the fictitious default algorithm of Eisenberg and Noe (2001) and Rogers and Veraart (2013), the loss cascades described in Cont, Moussa, and Santos (2013), the cascade hierarchies of Elliott, Golub, and Jackson (2014), or the clearing algorithms of Kusnetsov and Veraart (2016).

Proposition 3.6. *We state some immediate consequences of Definitions 3.4 and 3.5.*
1. *The sets $\mathcal{S}$ and $\mathcal{D}$ specified by the mild or the strict default rule are consistent extensions of (2.2) and (2.3).*
2. *For every consistent extension of (2.2) and (2.3), the survival set contains the one prescribed by the strict default rule and is included in the one given by the mild default rule. Similarly, the default set under a consistent extension is always a superset of the default set under the mild default rule and a subset of the default set under the strict default rule.*
3. *Under the mild (strict) default rule, if $\Sigma \subseteq [N]$ and $G_{[N] \setminus \Sigma}$ (G_Σ) has no directed cycles, then formula (3.6) is valid.*

Apart from the special cases discussed in the previous results, the presence of cycles in G prevents us from a concise description of the joint probability distribution of $D_1, \ldots, D_N$ using Bayesian network methods. However, under the mild or the strict default rule, the Bayesian network structure can be recovered if we "blow up" the redemption graph in a suitable way. The idea here is to explicitly incorporate the cascade mechanism behind the mild or the strict default rule into the graphical structure. Instead of a single variable D_i that determines whether firm i defaults or not, we keep track of a whole tuple of variables D_{in}, $n \in [N]$, that depend on whether i defaults in the nth round. In fact, to reduce the number of new variables as far as possible, it suffices to consider this separately for the *strongly connected components* of the redemption graph G. These are the equivalence classes induced by the equivalence relation

$$i \sim_G j :\iff i \lesssim_G j \text{ and } j \lesssim_G i, \quad i, j \in [N],$$

or, in economic terms, these are the maximal subsets of firms where any two of them have a direct or an indirect lender-borrower relationship via intermediate firms.

Definition 3.7. *Let $C_1, \ldots, C_m$ be the strongly connected components in G and N_i the size of the strongly connected component firm i belongs to. The* acyclic augmentation of G *is the graph $\bar{G} := (\bar{\mathcal{V}}, \bar{\mathcal{E}})$ with the vertex set*

$$(3.8) \qquad \bar{\mathcal{V}} := \{(i, n) : i \in [N], n \in [N_i]\}$$

and the edge set

$$\bar{\mathcal{E}} := \big\{ \big((j, N_j), (i, n)\big) : j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i), n = 1, \dots, N_i \big\}$$

$$\cup \bigcup_{i=1}^{N} \big\{ \big((i, n), (i, n+1)\big) : N_i > 1, n = 1, \dots, N_i - 1 \big\}$$

$$\cup \bigcup_{k=1}^{m} \big\{ \big((j, n), (i, n+1)\big) : i, j \in C_k, j \in \mathrm{pa}_G(i), n = 1, \dots, N_i - 1 \big\}$$

$$(3.9) \qquad \cup \bigcup_{k=1}^{m} \big\{ \big((j, n), (i, n+2)\big) : i, j \in C_k, j \in \mathrm{pa}_G(i), N_i > 2, n = 1, \dots, N_i - 2 \big\}.$$

In words, we first create N_i vertex copies in $\bar{G}$ for each firm $i \in [N]$, numbered from $(i, 1)$ to (i, N_i). Then the first type of edges in (3.9) connects two directly linked firms i and j in G that belong to *different* strongly connected components: Namely, if j has to repay i at time T, the *last* copy of j is connected to *each* copy of i. The second type of linkages connects consecutive copies of a fixed firm. The third (resp., fourth) type of edges joins consecutive (resp., next-but-one) copies of two firms that are directly linked to each other in G and belong to the *same* strongly connected component.

Figure 2 shows an example of a cyclic graph and its acyclic augmentation. As we can see, although G has a directed cycle (the vertices $\{1, 2, 3, 4\}$ form a strongly connected component), its augmentation $\bar{G}$ has no cycles. Of course, this holds in general.

Lemma 3.8. *The acyclic augmentation $\bar{G}$ of the redemption graph G is a DAG.*

(a) G　　　　　　　　　　　　　　　　(b) $\bar{G}$
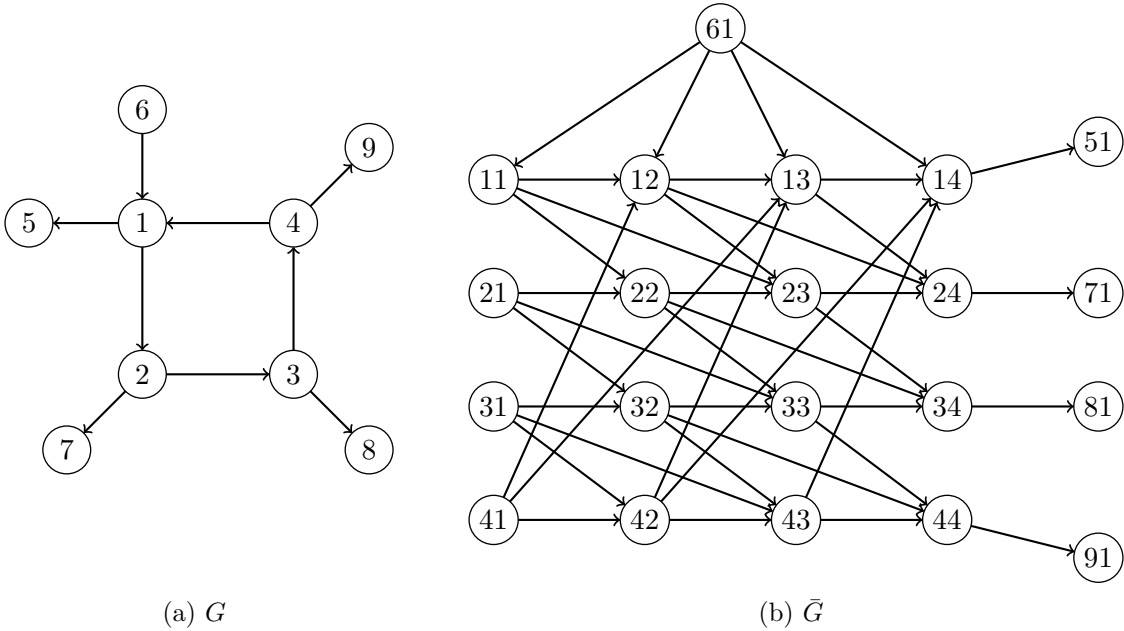
Figure 2. *A cyclic redemption graph G and its acyclic augmentation $\bar{G}$.*

Next we associate random variables to the vertices of $\bar{G}$, in such a way that they contain our variables of interest $D_1, \ldots, D_N$ as subsets and form a Bayesian network on $\bar{G}$. We first assume the mild default rule; for the strict default rule, all definitions and results are analogous and will be provided afterwards.

Definition 3.9. *Let N_i be as in Definition 3.7. For $i \in [N]$ and $n \in [N_i]$, we define $D_{i0} := 0$ and then inductively $D_{in} = 1$ if*

$$X_i(T) + \mathrm{e}^{r_0 T} K_i + \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)} \mathrm{e}^{r_{ij} T} L_{ij} \mathbb{1}_{\{D_{jN_j}=0\}} + \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i)} \mathrm{e}^{r_{ij} T} L_{ij} \mathbb{1}_{\{D_{j,n-1}=0\}}$$

(3.10)

$$- \mathrm{e}^{r_0' T} F_i - \sum_{j=1}^{N} \mathrm{e}^{r_{ji} T} L_{ji} < 0$$

and $D_{in} = 0$ otherwise.

To paraphrase, we first order the strongly connected components of G in such a way that firms in a given component have no borrowings from firms in previous components. Then, starting with those components that have no borrowings from any other component, we assign the value 1 to D_{in} exactly when firm i defaults assuming it only receives payments from surviving firms of previous strongly connected components and from firms j in the same component that have not defaulted up to the $(n-1)$st step, that is, where $D_{j,n-1} = 0$.

Theorem 3.10. *Under the mild default rule the following statements are valid.*
(1) *For all $i \in [N]$ and $n \in [N_i]$ we have $D_{i,n-1} \leq D_{in}$ and $D_i = D_{iN_i}$.*
(2) *The random variables $\{D_{in} : i \in [N],\ n \in [N_i]\}$ form a Bayesian network on $\bar{G}$ with conditional probabilities given by*

$$\mathbb{P}[D_{i1} = 1 \mid \mathrm{pa}_{\bar{G}}(D_{i1})] = \Phi(i, \Sigma_{i1}^{\mathrm{m}}, T),$$

(3.11)

$$\mathbb{P}[D_{in} = 1 \mid \mathrm{pa}_{\bar{G}}(D_{in})] = \begin{cases} 1 & \text{if } D_{i,n-1} = 1, \\ \frac{\Phi(i, \Sigma_{in}^{\mathrm{m}}, T) - \Phi(i, \Sigma_{i,n-1}^{\mathrm{m}}, T)}{1 - \Phi(i, \Sigma_{i,n-1}^{\mathrm{m}}, T)} & \text{if } D_{i,n-1} = 0, \end{cases} \quad n = 2, \ldots, N_i,$$

where

$$\Sigma_{in}^{\mathrm{m}} := \{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i) : D_{jN_j} = 0\} \cup \{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i) : D_{j,n-1} = 0\}, \quad n \in [N_i].$$

Although the variable D_{in} does not explicitly depend on $D_{i,n-2}$ in Definition 3.9, we have to form the next-but-one edges in $\bar{G}$ (i.e., the last of the four types of edges in (3.9)). The reason behind is that on the event $D_{i,n-1} = 0$, the values $D_{j,n-2}$ of parents j belonging to the same strongly connected component as i contain important information about $X_i(T)$. For example, if $D_{j,n-2} = 1$ for many such parents (so i survives in round $n-1$ even though many of its parents have defaulted in round $n-2$), the value $X_i(T)$ has to exceed a higher threshold compared to a situation where $D_{j,n-2} = 0$ for many j. Furthermore, the differences $D_{j,n-1} - D_{j,n-2}$ influence the likelihood of whether $D_{i,n-1} = D_{in} = 0$ or $0 = D_{i,n-1} < D_{in} = 1$. In fact, if $D_{j,n-1} - D_{j,n-2} = 0$ for all parents j, then necessarily the first alternative occurs. By contrast,

the likelihood for the second alternative increases with the frequency that $D_{j,n-1} - D_{j,n-2} = 1$ occurs. These considerations are reflected in formula (3.10) via the dependence on $\Sigma_{i,n-1}^{\mathrm{m}}$, which involves the variables $(D_{j,n-2} : j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i))$.

Next we present analogous statements for the strict default rule, which are simpler to formulate in terms of survival variables.

Definition 3.11. *Under the strict default rule we define for $i \in [N]$ and $n \in [N_i]$ the variables $S_{i0} := 0$, and $S_{in} := 1$ if*

$$X_i(T) + \mathrm{e}^{r_0 T} K_i + \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)} \mathrm{e}^{r_{ij} T} L_{ij} \mathbb{1}_{\{S_{jN_j}=1\}} + \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i)} \mathrm{e}^{r_{ij} T} L_{ij} \mathbb{1}_{\{S_{j,n-1}=1\}}$$

$$(3.12) \qquad - \mathrm{e}^{r_0' T} F_i - \sum_{j=1}^{N} \mathrm{e}^{r_{ji} T} L_{ji} \geq 0$$

and $S_{in} := 0$ otherwise.

Theorem 3.12. *Under the strict default rule the following statements are valid.*
(1) *For all $i \in [N]$ and $n \in [N_i]$ we have $S_{i,n-1} \leq S_{in}$ and $S_i = S_{iN_i}$.*
(2) *The random variables $\{S_{in} : i \in [N],\ n \in [N_i]\}$ form a Bayesian network on $\bar{G}$ with conditional probabilities given by*

$$\mathbb{P}[S_{i1} = 0 \mid \mathrm{pa}_{\bar{G}}(S_{i1})] = \Phi(i, \Sigma_{i1}^{\mathrm{s}}, T),$$

$$(3.13) \qquad \mathbb{P}[S_{in} = 0 \mid \mathrm{pa}_{\bar{G}}(S_{in})] = \begin{cases} 0 & \text{if } S_{i,n-1} = 1, \\ \dfrac{\Phi(i, \Sigma_{in}^{\mathrm{s}}, T)}{\Phi(i, \Sigma_{i,n-1}^{\mathrm{s}}, T)} & \text{if } S_{i,n-1} = 0, \end{cases} \quad n = 2, \ldots, N_i,$$

where

$$\Sigma_{in}^{\mathrm{s}} := \{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i) : S_{jN_j} = 1\} \cup \{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i) : S_{j,n-1} = 1\}, \quad n \in [N_i].$$

The quintessence of this section is that the default variables $D_1, \ldots, D_N$ associated to the financial institutions form a Bayesian network in the case of an acyclic redemption graph and that they can be embedded as margins into a larger random vector with a Bayesian network structure in the cyclic case. But of course, these results are obtained under the hypotheses on the financial system made in section 2. The question of how modifications of these assumptions affect the Bayesian network structure will be addressed in section 5.

4. Implications of the Bayesian network structure. We have proved in Theorems 3.2, 3.10, and 3.12 that the default and survival random variables associated to firms with mutual liabilities form a Bayesian network over the redemption graph or a suitable augmentation thereof. In this section, we elaborate on the benefits of this Bayesian network structure for systemic risk analysis.

4.1. Independence detection. It is important for both the involved institutions and the regulatory authorities to determine whether the default of different firms depends on each other or not and how this relationship is affected by the state of other firms. The underlying Bayesian network structure permits us to reduce this question to a graph-theoretic analysis.

More precisely, given the evidence of D_k, $k \in K$, for some subset $K \subseteq [N]$, one can determine by sole inspection of $\bar{G}$, the augmentation of the redemption graph G from Definition 3.7, whether $\{D_i : i \in I\}$ and $\{D_j : j \in J\}$ for two further subsets $I, J \subseteq [N]$ depend on each other or not.

The key concept here is called d-separation, see p. 48 in Lauritzen (1996):

Definition 4.1. *Let $G = (\mathcal{V}, \mathcal{E})$ be a graph.*
(1) *We call $i_1 \rightleftharpoons_G \cdots \rightleftharpoons_G i_n$ a* chain *between i_1 and i_n if for every $j \in [n-1]$ we have $i_j \rightarrow_G i_{j+1}$ or $i_{j+1} \rightarrow_G i_j$ (or both).*
(2) *For three pairwise disjoint subsets V_0, V_1, V_2 of $\mathcal{V}$, the sets V_1 and V_2 are called d-separated in G given V_0 (or simply* d-separated *if $V_0 = \emptyset$) if every chain $i_1 \rightleftharpoons_G \cdots \rightleftharpoons_G i_n$ in G from some $i_1 \in V_1$ to some $i_n \in V_2$ is* blocked *by V_0; that is, for some $j = 2, \ldots, n-1$ we have*
   - $i_{j-1} \rightarrow_G i_j \rightarrow_G i_{j+1}$ *and $i_j \in V_0$ (type I structure), or*
   - $i_{j-1} \leftarrow_G i_j \leftarrow_G i_{j+1}$ *and $i_j \in V_0$ (type II structure), or*
   - $i_{j-1} \leftarrow_G i_j \rightarrow_G i_{j+1}$ *and $i_j \in V_0$ (type III structure), or*
   - $i_{j-1} \rightarrow_G i_j \leftarrow_G i_{j+1}$ *and $i_j \notin V_0$ and $\mathrm{de}_G(i_j) \cap V_0 = \emptyset$ (type IV structure).*

For example, if G is a DAG, then for every $i \in \mathcal{V}$ the sets $\{i\}$ and $\mathrm{nd}_G(i)$ are d-separated in G given $\mathrm{pa}_G(i)$. If random variables $\{X_i : i \in \mathcal{V}\}$ form a Bayesian network on G, the fundamental result is the following (see Corollary 3.23 together with Proposition 3.25 in Lauritzen (1996)): If V_1 and V_2 are d-separated in G given V_0, then the two collections of random variables $\{X_i : i \in V_1\}$ and $\{X_i : i \in V_2\}$ are conditionally independent given $\{X_i : i \in V_0\}$. Numerical algorithms for checking d-separation are well established and efficient (see section 3.3 of Koller and Friedman (2009)).

Applied to the acyclic augmentation $\bar{G} = (\bar{\mathcal{V}}, \bar{\mathcal{E}})$ of the redemption graph, the d-separation criterion provides a graphical tool for detecting independence between different groups of firms regarding their default behavior. In fact, d-separation is "almost" necessary for independence (see Theorem 7 of Meek (1995)): If V_0, V_1, and V_2 are subsets of $\bar{\mathcal{V}}$, and V_1 and V_2 are not d-separated given V_0, then for almost all volatility parameters σ_i the random variables $\{D_i : i \in V_1\}$ and $\{D_i : i \in V_2\}$ are stochastically dependent given $\{D_i : i \in V_0\}$ if none of the variables D_i, $i \in V_1 \cup V_2$, becomes deterministic given $\{D_i : i \in V_0\}$. The next proposition relates d-separation in G to d-separation in $\bar{G}$.

Proposition 4.2. *Let G be the redemption graph and $\bar{G}$ its acyclic augmentation. Then the following assertions hold for all pairwise disjoints subsets V_1, V_2, and V_3 of $[N]$.*
(1) *V_1 and V_2 are d-separated in G if and only if $\{(i, N_i) : i \in V_1\}$ and $\{(i, N_i) : i \in V_2\}$ are d-separated in $\bar{G}$.*
(2) *If V_1 and V_2 are d-separated in G given V_0, then $\{(i, n) : i \in V_1, \ n \in [N_i]\}$ and $\{(i, n) : i \in V_2, \ n \in [N_i]\}$ are d-separated in $\bar{G}$ given $\{(i, n) : i \in V_0, \ n \in [N_i]\}$.*

4.2. Measures of systemic impact. In the systemic risk literature several measures of systemic importance for interconnected firms have been proposed. While some approaches like in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015), Battiston et al. (2012), and Diebold and Yilmaz (2014) are purely based on the underlying network structure and parameters, others explicitly take the probabilistic nature of contagion into account. Examples for the

latter category include the contagion index of Cont, Moussa, and Santos (2013), the CoVaR measure of Adrian and Brunnermeier (2016), the SES and SRISK measures of Acharya et al. (2016) and Brownlees and Engle (2017), and the general systemic risk measures of Biagini et al. (to appear) and Feinstein, Rudloff, and Weber (2017). Based on the structural cascade model analyzed in sections 2 and 3, we contribute two further probabilisitic measures of systemic risk that are flexible (one can measure the impact of a single firm or a group of firms, unconditionally or conditionally on different stress scenarios), feasible (they can be computed by the methods described in the next subsection), and robust (they depend continuously on the underlying model parameters).

In the following, if not otherwise stated, any consistent extension of (2.2) and (2.3) can be considered, in particular both the mild and the strict default rule.

Definition 4.3. *For two disjoint subsets I and J of $[N]$, the* absolute *and the* relative systemic impact *(ASI and RSI) of I on J are defined as (we use 1 to denote a vector with all entries equal to one)*

$$(4.1) \qquad \mathrm{ASI}(I, J) := \max_{J_0 \subseteq J, \ e \in \{0,1\}^{|J_0|}} \left(\mathbb{P}[D_{J_0} = e \mid D_I = 1] - \mathbb{P}[D_{J_0} = e] \right)$$

$$(4.2) \qquad and \quad \mathrm{RSI}(I, J) := \max_{e \in \{0,1\}^{|J|}} \log_2 \frac{\mathbb{P}[D_J = e \mid D_I = 1]}{\mathbb{P}[D_J = e]}$$

respectively. In (4.2), we set $\log_2 0 = -\infty$, and $0/0 = 1$ in case the denominator, and, as a consequence, also the numerator is zero.

The ASI of I on J equals the *total variation distance* between the distribution of D_J and its conditional distribution given $D_I = 1$. Moreover, the RSI of I on J is the *Rényi divergence of order $+\infty$* between the same pair of distributions. Of course, instead of taking the maximum in (4.1) and (4.2), also L^p-type measures can be investigated. For example, one can define L^1-analogues of the ASI and the RSI measures based on *mutual information* and the *Kullback–Leibler divergence*. We do not go into details at this point but only refer to van Erven and Harremoës (2014) for more information on distances between probability measures.

Turning to the literature on Bayesian networks, we found, to our surprise, only little research where measures of importance for vertices in a network are investigated. Most notably, in Chapter 4 of Pinto (1986) (cf. Pearl (1988), Chapter 6.4), the author proposes quantifying the relevance of a vertex Y for another vertex X by a function depending on the distribution of X and its conditional distribution given Y. They give a version of the ASI measure as an example. Furthermore, Chan and Darwiche (2005) employ a symmetric version of the RSI measure in the context of sensitivity analysis for Bayesian networks.

Next, we collect some basic properties of the ASI and RSI measures.

Proposition 4.4. *The following assertions hold for disjoint sets $I, J \subseteq [N]$.*
(1) $\mathrm{ASI}(I, J) \in [0, 1)$ *and* $\mathrm{RSI}(I, J) \in [0, \infty)$.
(2) *If $J_1 \subseteq J$, then $\mathrm{ASI}(I, J_1) \leq \mathrm{ASI}(I, J)$ and $\mathrm{RSI}(I, J_1) \leq \mathrm{RSI}(I, J)$.*

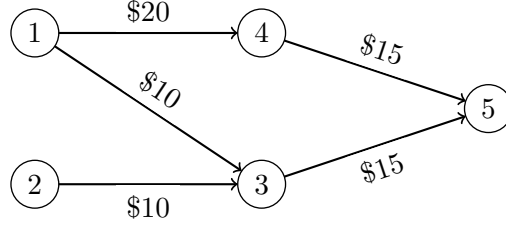**Figure 3.** *Redemption graph for the firm network of Example* 4.5.

(3) *An explicit formula for* (4.1) *is*

$$\text{ASI}(I, J) = \sum_{e \in \{0,1\}^{|J|}} \left(\mathbb{P}[D_J = e \mid D_I = 1] - \mathbb{P}[D_J = e] \right)^{+}$$

$$(4.3) \qquad\qquad = \frac{1}{2} \sum_{e \in \{0,1\}^{|J|}} \left| \mathbb{P}[D_J = e \mid D_I = 1] - \mathbb{P}[D_J = e] \right|.$$

(4) *In contrast to* (4.1), *there is no need to consider subsets of J in* (4.2) *because*

$$(4.4) \qquad\qquad \text{RSI}(I, J) = \max_{J_0 \subseteq J, \ e \in \{0,1\}^{|J_0|}} \log_2 \frac{\mathbb{P}[D_{J_0} = e \mid D_I = 1]}{\mathbb{P}[D_{J_0} = e]}.$$

(5) *Under the strict default rule, if $I = I_1 \cup I_2$ and I_2 and J are d-separated in G given I_1, then* $\text{ASI}(I, J) = \text{ASI}(I_1, J)$ *and* $\text{RSI}(I, J) = \text{RSI}(I_1, J)$.

One would also expect $\text{ASI}(I_1, J) \leq \text{ASI}(I, J)$ and $\text{RSI}(I_1, J) \leq \text{RSI}(I, J)$ to hold for $I_1 \subseteq I$ since the knowledge of default of a larger set of firms ought to be worse, or at least not better, for the survival of another firm in the network. However, a simple counterexample demonstrates that this intuition is wrong.

Example 4.5. We consider a network of $N = 5$ institutions with redemption graph as in Figure 3 and zero interest rates. Further suppose that $F_3 = F_4 = K_3 = K_4 = \$0$; that is, both firm 3 and firm 4 have neither cash nor external liabilities at time $t = T$. Moreover, we make the following assumptions on the firms' operative performance:

- $\mu_1 = \mu_2$ are very high, and $\sigma_1 = \sigma_2$ are strictly positive but small.
- μ_3 and μ_4 are strongly negative, and σ_3 and σ_4 are small in comparison.
- $X_1 = X_2$ such that $X_1(T)$ and $X_2(T)$ are independent with the same distribution.

As a consequence of the first hypothesis, the probabilities of default for firms 1 and 2 are very small but strictly positive. Moreover, by the second assumption, without repayment from firms 1 and 2, firms 3 and 4 have almost no chance to survive and pay back firm 5. However, as soon as firm 1 survives, also firm 4 survives, and as soon as both firm 1 and firm 2 are solvent, also firm 3 is so. Hence we obtain the following:

- $\mathbb{P}[D_4 = 1] \approx 0$ because firm 1 has a very high chance of survival such that also firm 4 is most likely to survive.
- $\mathbb{P}[D_4 = 1 \mid D_3 = 1] \approx 1/2$ because conditional on the default of firm 3, the event of exactly one failure among 1 and 2 is almost one (the probability of both firms surviving is zero because otherwise 3 would not have defaulted, and the probability of a joint

default of 1 and 2 is still comparatively small). Since only the failure of firm 1 triggers the default of 4, and $X_1(T)$ and $X_2(T)$ are identically distributed, the conditional default probability of 4 given the default of 3 is around $1/2$.

- $\mathbb{P}[D_4 = 1 \mid D_2 = D_3 = 1] \approx 0$ because the failure of 2 already explains the failure of 3, so the latter gives us almost no additional information about 1 (and therefore 4).

As a result, $\mathrm{ASI}(\{2,3\},\{4\}) \le \mathrm{ASI}(\{3\},\{4\})$ and $\mathrm{RSI}(\{2,3\},\{4\}) \le \mathrm{RSI}(\{3\},\{4\})$.

4.3. Computation of default probabilities. In addition to finding the dependence relations among interconnected networks, another central objective of risk management is to determine the default probabilities of firms or the impact of a firm's default on other firms, measured, for example, in terms of the ASI and RSI of the previous subsection. Known as the *probabilistic inference problem* in the artificial intelligence community, this topic has been extensively studied in the literature of Bayesian networks, where various exact or approximative algorithms have been developed for this purpose (see Koller and Friedman (2009) for an in-depth treatment on this subject). In order to illustrate the results and concepts of the previous sections, we demonstrate in a numerical example how these inference algorithms can be applied to our setting. All numerical computations in the following are made using the *junction tree algorithm* within the open-source Bayes Net Toolbox for Matlab (Murphy, 2001).

We consider a financial network with $N = 100$ banks in total where the redemption graph is depicted in Figure 4. The network has a core-periphery structure in the sense that the first five banks form a complete network where all possible lending relationships are present, and each of these core banks has further borrowed money from 19 different periphery banks. We refer to the papers Elsinger, Lehar, and Summer (2006), Cont, Moussa, and Santos (2013), Craig and von Peter (2014), and Langfield, Liu, and Ota (2014) for empirical evidence of
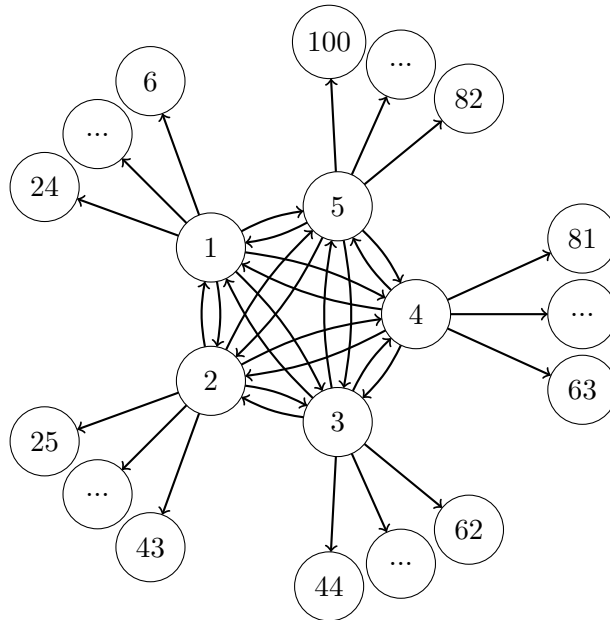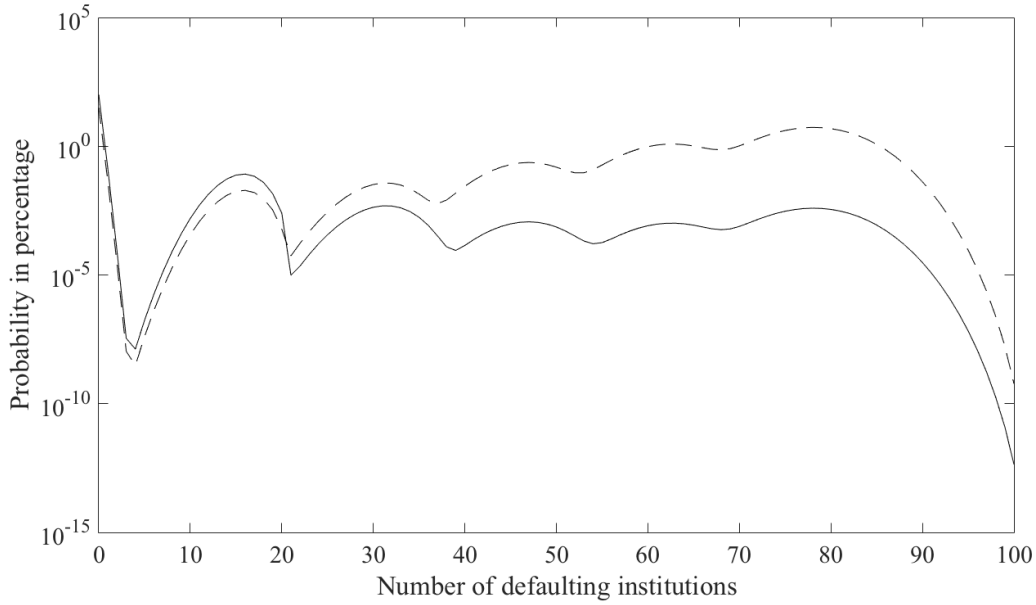

**Figure 4.** *A core-periphery structured redemption graph.*

Figure 5. *Default probabilities under the mild default rule (solid line) and the strict default rule (dashed line).*

such structures in real banking networks. The time horizon in our example is $T = 1$ year, all interest rates are zero, and we assume statistical equality among all core banks and among all periphery banks:

$$\mu_C = 0.1, \qquad \sigma_C = 0.2, \qquad X_C = 2000, \qquad F_C = 500, \qquad L_{CC} = 1600/(N-1) = 400,$$
$$\mu_P = 0.05, \qquad \sigma_P = 0.1, \qquad X_P = 80, \qquad F_P = 90, \qquad L_{PC} = 35 \text{ (if } C \rightarrow P).$$

The subscript C (P) here represents an arbitrary core (periphery) banks. Figure 5 displays the probability distribution of the number of defaults in the resulting financial systems under the mild and the strict default rule. While the shape of the default distribution is the same for both the mild and the strict default rule, the actual probabilities are very different. No default occurs under the mild rule with a probability of 99.39%, but this probability decreases to 31.15% when the strict rule applies. This discrepancy is due to the fact that a core bank only has a 11.13% chance to survive without the repayment of the other core banks but a 99.90% chance if it receives full repayment. As a consequence, in the considered example, the possibility of netting mutual claims is decisive, whether the contagion cascade becomes unleashed or not.

From Figure 5, we can observe several prominent qualitative features of the default distribution. First of all, the distribution is multimodal, corresponding to the number of core banks in default. Indeed, conditional on the default of a core bank, each periphery bank with lendings to this core bank defaults (independently) with a probability of 76.66%. Thus, whenever a core bank defaults, it triggers the default of $19 \times 76.66\% = 14.57$ periphery banks on average, which explains why for both the mild and the strict rule the default distribution has modes at distances of 15 or 16 (one core bank plus 14 or 15 periphery banks).
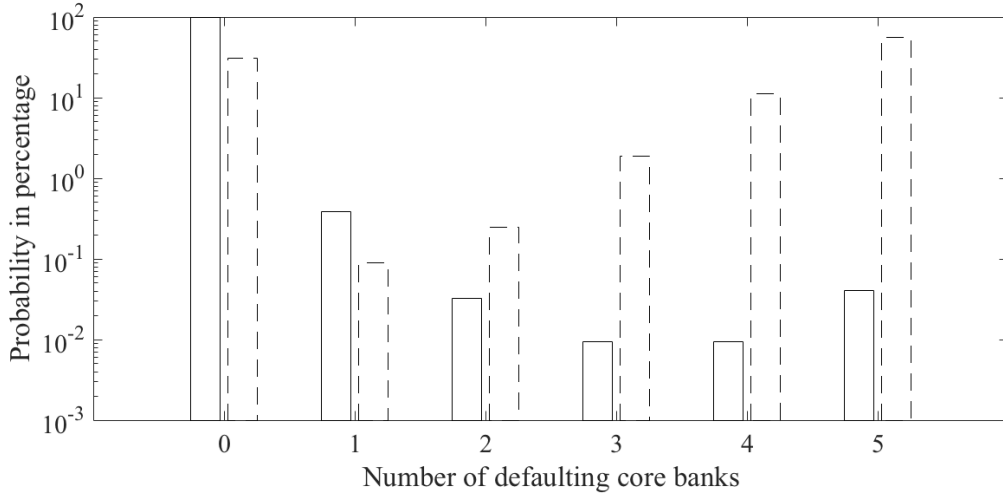
Figure 6. *Distribution of the number of defaulting core banks under the mild rule (solid bars) and the strict rule (dashed bars).*

Furthermore, we observe that also the probabilities at the modes are not monotone. This is caused by contagious effects among the core banks. When enough core banks default, it is rather likely that even more core banks default: The probability function of the number of defaulting core banks is not monotone as Figure 6 reveals.

Last but not least, the probabilities beyond the last mode with 78 defaults fall into the large deviation regime and decrease exponentially fast.

The measures of systemic impact introduced in the previous subsection can also be calculated in our example. For simplicity, we only report figures for the mild default rule; the computations under the strict default rule are completely analogous. To this end, let C and C$'$ be two arbitrary core banks and P and P$'$ be two periphery banks that are creditors of C and C$'$, respectively. For the absolute systemic impact of C we obtain

$$\text{ASI}(\{C\}, \{C'\}) = \mathbb{P}[D_{C'} = 1 \mid D_C = 1] - \mathbb{P}[D_{C'} = 1] = 36.21\% - 0.14\% = 36.07\%,$$
$$\text{ASI}(\{C\}, \{P\}) = \mathbb{P}[D_P = 1 \mid D_C = 1] - \mathbb{P}[D_P = 1] = 76.66\% - 0.11\% = 76.55\%,$$
$$\text{ASI}(\{C\}, \{P'\}) = \mathbb{P}[D_{P'} = 1 \mid D_C = 1] - \mathbb{P}[D_{P'} = 1] = 27.76\% - 0.11\% = 27.65\%,$$

which translates to the relative systemic impact of C as

$$\text{RSI}(\{C\}, \{C'\}) = 7.97, \quad \text{RSI}(\{C\}, \{P\}) = 9.42, \quad \text{RSI}(\{C\}, \{P'\}) = 7.96.$$

Concering the ASI and RSI of periphery banks, we fix another periphery bank P$''$ that is, like P, also a creditor of C. Then we obtain

$$\text{ASI}(\{P\}, \{C\}) = \mathbb{P}[D_C = 1 \mid D_P = 1] - \mathbb{P}[D_C = 1] = 98.79\% - 0.14\% = 98.65\%,$$
$$\text{ASI}(\{P\}, \{C'\}) = \mathbb{P}[D_{C'} = 1 \mid D_P = 1] - \mathbb{P}[D_{C'} = 1] = 35.78\% - 0.14\% = 35.63\%,$$
$$\text{ASI}(\{P\}, \{P'\}) = \mathbb{P}[D_{P'} = 1 \mid D_P = 1] - \mathbb{P}[D_{P'} = 1] = 27.43\% - 0.11\% = 27.32\%,$$
$$\text{ASI}(\{P\}, \{P''\}) = \mathbb{P}[D_{P''} = 1 \mid D_P = 1] - \mathbb{P}[D_{P''} = 1] = 75.74\% - 0.11\% = 75.63\%,$$

and

$$\text{RSI}(\{P\}, \{C\}) = 9.42, \quad \text{RSI}(\{P\}, \{C'\}) = 7.96, \quad \text{RSI}(\{P\}, \{P'\}) = 7.94,$$
$$\text{RSI}(\{P\}, \{P''\}) = 9.40.$$

As both the ASI and RSI reveal, even banks that are not direct counterparties may be systemically important for each other. For instance, if a core bank defaults, the default probability of a periphery bank increases by 27.65% (or by a factor of $2^{7.96} \approx 248$) even if it has no lendings to that defaulted core bank. Another observation is that also periphery banks have a high systemic impact although they have no borrowings from other banks in the system. The reason behind is simply that periphery banks that are repaid by their core bank counterparty survive with a probability larger than 99.99%. So if a periphery bank defaults, this is very likely caused by the default of its core bank counterparty.

5. Conclusion and outlook. In this paper we analyzed a financial contagion model in which the joint default probability distribution of the system can be characterized in terms of a Bayesian network on a graph determined by the interfirm liabilities. We further explained how this graphical structure can be employed to detect systemic dependencies within the network, to define measures of systemic importance for institutions, and to compute their conditional or unconditional default probabilities. Since the involved methodologies apply to all possible networks, this work provides a useful device to analyze and monitor the systemic risk in financial systems.

Naturally, we did impose various conditions on the financial network, so let us comment on possible generalizations. First of all, regarding probabilistic assumptions, the lognormal hypothesis on the evolution of the operating firm assets is merely for simplicity. The Bayesian network structure of the default variables is preserved if we assume another distribution for the asset returns. Only the conditional probabilities stated in Theorems 3.2, 3.10, and 3.12 have to be adapted.

The situation is different if we allow for dependent asset returns of the financial institutions. Of course, if one is only interested in computing default probabilities using Monte Carlo techniques, this is still feasible as long as the joint asset return distribution can be simulated from. However, the Bayesian network structure of the default variables $D_1, \ldots, D_N$ will be lost in general because observations of nondescendant institutions (with respect to the redemption graph) may reveal additional information about the driving noise. But let us suppose that each institution i is driven by a finite number of common market noises $W_1, \ldots, W_K$ plus its own idiosyncratic noise B_i, assuming that the B_i's are still independent and also independent of the W_j's. Then conditional on the W_j's, the variables $D_1, \ldots, D_N$ will have a Bayesian network structure. Hence, our approach permits default risk analysis under various market situations, which is useful for stress testing, for instance.

Other generalizations may concern the assumptions on the financial system itself and the way institutions are interconnected. In our model, all mutual exposures are loans with two possible values at the end of the period (depending on whether the counterparty fails or not). It would be more realistic to include other types of exposures, for example, exposures with continuously varying values like asset cross-holdings (cf. Gouriéroux, Héam, and Monfort (2012) and Elliott, Golub, and Jackson (2014)), exposures with different seniorities, exposures

with a nonzero recovery rate in the case of counterparty default, or exposures with different maturity dates. These considerations clearly go beyond the scope of the current paper and are left to future research.

Appendix A. Proofs.

Proof of Theorem 2.4. (1) If G is a DAG, the relation $\lesssim_G$ defined in (3.7) induces a partial order on $\{1, \ldots, N\}$ with a nonempty set $\mathcal{V}_1$ of minimal nodes with respect to $\lesssim_G$, which precisely contains those firms that have not lent money to any other firm. For such a firm $i \in \mathcal{V}_1$, the equation for $E_i(T)$ in (2.3) is explicit: It does not depend on any other $j \neq i$ because $L_{ij} = 0$. Next, we consider the minimal nodes in $\mathcal{V} \setminus \mathcal{V}_1$ with respect to $\lesssim_{G_{\mathcal{V} \setminus \mathcal{V}_1}}$, that is, those firms that have only lent money to firms in $\mathcal{V}_1$. Their values $E_i(T)$ are given by an expression involving $E_j(T)$ only for $j \in \mathcal{V}_1$. Iterating this at most N times, we obtain an explicit formula for all $E_i(T)$ in terms of $X_1(T), \ldots, X_N(T)$ and the model parameters.

(2) Now we suppose that G has a directed cycle, say, involving the firms $1, \ldots, N_0$ in that order. By assumption (2.1), for each firm $i \in \{N_0 + 1, \ldots, N\}$ there exists a range $[x_i^{(1)}, x_i^{(2)})$ with $0 = x_i^{(1)} < x_i^{(2)}$ such that i necessarily defaults whenever $X_i(T)$ falls into this interval. For the N_0 firms within the cycle we first define L_0 to be the smallest amount a firm in $\{1, \ldots, N_0\}$ has to repay to its creditor within the cycle. An inspection of (2.3) reveals that there exists a rectangle $\prod_{i=1}^{N_0}[x_i^{(1)}, x_i^{(2)})$ with $x_i^{(1)} < x_i^{(2)}$ and the following property: When $(X_1(T), \ldots, X_{N_0}(T))$ lies in that rectangle, then each firm in $[N_0]$ is able to repay the external liabilities and all its creditors except the one in the cycle, but the remaining asset value is larger than the debt owed to this cycle creditor minus L_0. As a result, if $(X_1(T), \ldots, X_N(T)) \in \prod_{i=1}^{N}[x_i^{(1)}, x_i^{(2)})$, which is an event with strictly positive probability, then one solution would be that all firms default, and another would be that the firms in the cycle survive while all other firms default. ∎

Proof of Theorem 3.2. This theorem is a special case of Theorem 3.12. ∎

Proof of Proposition 3.3. That (3.6) holds with "≤" is immediate: If $\mathcal{S} = \Sigma$, equation (2.3) becomes an explicit expression in terms of $X_i(T)$ and the model parameters for each $i \in [N]$. The independence of $X_1(T), \ldots, X_N(T)$ together with their lognormal distributions yields (3.6) with "≤." In order to prove equality when both G_Σ and $G_{[N] \setminus \Sigma}$ are DAGs, we consider the situation where $X_1(T), \ldots, X_N(T)$ satisfy (2.2) and (2.3) with $\mathcal{S} = \Sigma$. We have to show that all other choices $\mathcal{S} = \Sigma'$ with $\Sigma' \neq \Sigma$ would violate (2.2) and (2.3). Since G_Σ is a DAG, the relation $\lesssim_{G_\Sigma}$ induces a partial order on Σ. The minimal nodes in Σ with respect to that partial order, denoted by Σ_1, correspond to firms that survive without receiving any payment from other firms in the network at time T. So they must belong to Σ' whenever $\mathcal{S} = \Sigma'$ is supposed to satisfy (2.2) and (2.3). Next, we consider the minimal nodes in $\Sigma \setminus \Sigma_1$. Again, they must belong to Σ' because they survive regardless of the solvency of all other firms except for the minimal nodes, which, however, are already known to survive. Inductively, this proves that $\Sigma \subseteq \Sigma'$ for every potential Σ'. Similarly, because $G_{[N] \setminus \Sigma}$ is a DAG, the maximal elements in $[N] \setminus \Sigma$ default regardless of whether the other firms in $[N] \setminus \Sigma$ default or not, so they cannot belong to Σ'. An analogous induction argument yields $[N] \setminus \Sigma \subseteq [N] \setminus \Sigma'$, so $\Sigma = \Sigma'$. ∎

Proof of Proposition 3.6. We assume the strict default rule; for the mild default rule, the proof is completely analogous.

(1) It is obvious that the strict default rule maps $X_1(T), \ldots, X_N(T)$ in a measurable way to $S_1, \ldots, S_N$. So we are left to prove that for $i \in [N]$ the term in (2.3) is nonnegative if and only if $i \in \mathcal{S}$. By Definition 3.5, the expression (2.3) for $i \in \mathcal{S}_n$ is already nonnegative when we take the indicator $\mathbb{1}_{\mathcal{S}_{n-1}}$ (for $n = 1$ we set $\mathcal{S}_0 := \emptyset$), which is always less than or equal to the original indicator $\mathbb{1}_{\mathcal{S}}$. Since $\mathcal{S} = \mathcal{S}_1 \cup \cdots \cup \mathcal{S}_N$, this proves that (2.3) is nonnegative for all $i \in \mathcal{S}$. By contrast, for all $i \in \mathcal{D}$ equation (2.3) must take a strictly negative value because otherwise i belongs to $\mathcal{S}_1 \cup \cdots \cup \mathcal{S}_{N_0+1}$ according to the strict default rule where N_0 is the smallest number in $[N]$ such that $\mathcal{S} = \mathcal{S}_1 \cup \cdots \cup \mathcal{S}_{N_0}$, which is obviously a contradiction.

(2) Denote by $\mathcal{S}_1, \ldots, \mathcal{S}_N$ the surviving firms of the nth round under the strict default rule. As in (1), a simple induction argument shows that $\mathcal{S}_n \subseteq \mathcal{S}$ holds for all $n \in [N]$ whenever $\mathcal{S}$ and $\mathcal{D}$ satisfy (2.2) and (2.3).

(3) We have already shown in Proposition 3.3 that (3.6) always holds with "$\leq$." For the reverse inequality, suppose that $X_1(T), \ldots, X_N(T)$ satisfy (2.2) and (2.3) with Σ instead of $\mathcal{S}$. Under the strict default rule, the proof of the same proposition reveals that, because of the DAG structure of G_Σ, we necessarily have $\Sigma \subseteq \mathcal{S}$, so part (2) of the current proposition completes the proof. ∎

Proof of Lemma 3.8. Suppose that $\bar{G}$ does have a cycle, say, of the form

$$(i_1, n_1) \to_{\bar{G}} \cdots \to_{\bar{G}} (i_k, n_k) \to_{\bar{G}} (i_1, n_1).$$

Then at least one of the edges must belong to the first of the four collections of edges in (3.9) because otherwise we would have $n_1 < n_2 < \cdots < n_k < n_1$, which is absurd. But if the cycle has an edge of the form $((i_l, N_{i_l}), (i_{l+1}, n_{l+1}))$ with $i_l \in \mathrm{pa}_G(i_{l+1}) \cap \mathrm{nd}_G(i_{l+1})$, there cannot exist a path in $\bar{G}$ that connects (i_{l+1}, n_{l+1}) to (i_l, N_{i_l}) because, by (3.9), this would give rise to a path in G from i_{l+1} to i_l and contradict the fact that $i_l \in \mathrm{nd}_G(i_{l+1})$. However, this in turn produces a contradiction to the assumption that $(i_1, n_1), \ldots, (i_k, n_k)$ form a cycle in $\bar{G}$. ∎

Proof of Theorem 3.10. (1) The first statement follows immediately from the definition in (3.10). For the second statement we number the strongly connected components of G in such a way that no vertex in C_k has parents in $C_{k+1}, \ldots, C_m$. By induction on k we may assume that $D_i = D_{iN_i}$ holds under the mild default rule for all $i \in C_1 \cup \cdots \cup C_{k-1}$. We must prove that $D_i = D_{iN_i}$ is true for $i \in C_k$. We first show that $D_i = 1$ holds as soon as $D_{in} = 1$ for some $n \in [N_i]$. To this end, we proceed with another induction on n and assume that we already know that $D_{j,n-1} = 1$ implies $D_j = 1$ for all $j \in C_k$. Then firms $i \in C_k$ with $D_{in} = 1$ cannot repay their debts if they only receive payments from $j \in \mathrm{pa}_G(i) \setminus C_k$ with $D_{jN_j} = 0$ and $j \in \mathrm{pa}_G(i) \cap C_k$ with $D_{j,n-1} = 0$. Indeed, all other firms default under the mild default rule: Firms $j \in \mathrm{pa}_G(i) \setminus C_k$ with $D_{jN_j} = 1$ by the first induction hypothesis and firms $j \in \mathrm{pa}_G(i) \cap C_k$ with $D_{j,n-1} = 1$ by the second. Thus, also i must default according to the mild default rule. For the converse direction that $D_i = 1$ implies $D_{iN_i} = 1$ for all $i \in C_k$, we recall that $\mathcal{D} = \bigcup_{n=1}^N \mathcal{D}_n$, where $\mathcal{D}_n$ contains the defaulting firms of the nth round. By induction on $n \in [N]$, let us assume that those $i \in C_k$ that belong to $\mathcal{D}_1 \cup \cdots \cup \mathcal{D}_{n-1}$ satisfy $D_{iN_i} = 1$. Then every $i \in \mathcal{D}_n$ has a negative equity value at time T if it does not receive money back from $j \in \mathrm{pa}_G(i) \cap (\mathcal{D}_1 \cup \cdots \cup \mathcal{D}_{n-1})$. If such a firm j belongs to $\mathrm{pa}_G(i) \setminus C_k$,

we have $D_{jN_j} = 1$ by the first induction hypothesis, while if it belongs to $\mathrm{pa}_G(i) \cap C_k$, we have $D_{jN_j} = 1$ by the second hypothesis, and consequently $D_{j,N_j-1} = 1$ (because the event that $D_{j,N_j-1} = 0$ and $D_{jN_j} = 1$ happens only if j defaults in a later round than all firms in C_k, which, however, is not possible because $j \in \mathcal{D}_1 \cup \cdots \cup \mathcal{D}_{n-1}$ and $i \in \mathcal{D}_n$). Therefore, we conclude that $D_{iN_i} = 1$, which ends the proof of $D_i = D_{iN_i}$ for $i \in C_k$.

(2)   If $\{e_{in}: i \in [N],\ n \in [N_i]\}$ is a given sequence taking values 0 or 1, then, because of (1), the probability that $D_{in} = e_{in}$ for all i and n can only be larger than zero if $e_{in} \leq e_{i,n+1}$ for all $i \in [N]$ with $N_i > 1$ and $n \in [N_i - 1]$. In this case, we define for $i \in [N]$ the number

$$n_i := \begin{cases} \min\{n \in [N_i]\colon e_{in} = 1\} & \text{if } e_{iN_i} = 1 \\ N_i + 1 & \text{if } e_{iN_i} = 0 \end{cases}$$

and observe that $D_{in} = e_{in}$ holds for all $i \in [N]$ and $n \in [N_i]$ if and only if (with $e_{i0} := 0$ and the abbreviation $L_i(T) := \mathrm{e}^{r_0'T} F_i + \sum_{j=1}^{N} \mathrm{e}^{r_{ji}T} L_{ji} - \mathrm{e}^{r_0 T} K_i$)

- for all $i \in [N]$ with $n_i = 1$ we have

$$X_i(T) < L_i(T) - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{jN_j}=0\}},$$

- for all $i \in [N]$ with $n_i = N_i + 1$ we have

$$X_i(T) \geq L_i(T) - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{jN_j}=0\}} - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{j,N_i-1}=0\}},$$

- and for all $i \in [N]$ with $2 \leq n_i \leq N_i$ we have

$$X_i(T) < L_i(T) - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{jN_j}=0\}} - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{j,n_i-1}=0\}}$$

and

$$X_i(T) \geq L_i(T) - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{jN_j}=0\}} - \sum_{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{j,n_i-2}=0\}}.$$

Thus, using the independence of $X_1(T), \ldots, X_N(T)$ and their lognormal distribution, we obtain on the one hand

$$\mathbb{P}[D_{in} = e_{in}\ \forall i \in [N],\ n \in [N_i]]$$
$$= \prod_{i:\, n_i=1} \Phi(i, \bar{\Sigma}^{\mathrm{m}}_{in_i}, T) \prod_{i:\, 2 \leq n_i \leq N_i} (\Phi(i, \bar{\Sigma}^{\mathrm{m}}_{in_i}, T) - \Phi(i, \bar{\Sigma}^{\mathrm{m}}_{i,n_i-1}, T))$$

(A.1)
$$\times \prod_{i:\, n_i=N_i+1} (1 - \Phi(i, \bar{\Sigma}^{\mathrm{m}}_{i,n_i-1}, T)),$$

where $\bar{\Sigma}^{\mathrm{m}}_{in} := \{j \in \mathrm{pa}_G(i) \cap \mathrm{nd}_G(i)\colon e_{jN_j} = 0\} \cup \{j \in \mathrm{pa}_G(i) \cap \mathrm{de}_G(i)\colon e_{j,n-1} = 0\}$.

On the other hand, if we assume for a moment that formula (3.13) is valid, we deduce that for all $i \in [N]$ with $n_i = 1$ we have

$$\prod_{n=1}^{N_i} \mathbb{P}[D_{in} = e_{in} \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})] = \mathbb{P}[D_{i1} = 1 \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})]$$

$$= \Phi(i, \bar{\Sigma}_{i1}^{\mathrm{m}}, T),$$

that for all $i \in [N]$ with $n_i = N_i + 1$ we have

$$\prod_{n=1}^{N_i} \mathbb{P}[D_{in} = e_{in} \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})] = \prod_{n=1}^{N_i} \mathbb{P}[D_{in} = 0 \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})]$$

$$= (1 - \Phi(i, \bar{\Sigma}_{i1}^{\mathrm{m}}, T)) \prod_{n=2}^{N_i} \frac{1 - \Phi(i, \bar{\Sigma}_{in}^{\mathrm{m}}, T)}{1 - \Phi(i, \bar{\Sigma}_{i,n-1}^{\mathrm{m}}, T)}$$

$$= 1 - \Phi(i, \bar{\Sigma}_{iN_i}^{\mathrm{m}}, T),$$

and that for all other $i \in [N]$ with $2 \le n_i \le N_i$ we have

$$\prod_{n=1}^{N_i} \mathbb{P}[D_{in} = e_{in} \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})]$$

$$= \mathbb{P}[D_{in_i} = 1 \mid \mathrm{pa}_{\bar{G}}(D_{in_i}) = \mathrm{pa}_{\bar{G}}(e_{in})] \prod_{n=1}^{n_i-1} \mathbb{P}[D_{in} = 0 \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})]$$

$$= \frac{\Phi(i, \bar{\Sigma}_{in_i}^{\mathrm{m}}, T) - \Phi(i, \bar{\Sigma}_{i,n_i-1}^{\mathrm{m}}, T)}{1 - \Phi(i, \bar{\Sigma}_{i,n_i-1}^{\mathrm{m}}, T)} (1 - \Phi(i, \bar{\Sigma}_{i1}^{\mathrm{m}}, T)) \prod_{n=2}^{n_i-1} \frac{1 - \Phi(i, \bar{\Sigma}_{in}^{\mathrm{m}}, T)}{1 - \Phi(i, \bar{\Sigma}_{i,n-1}^{\mathrm{m}}, T)}$$

$$= \Phi(i, \bar{\Sigma}_{in_i}^{\mathrm{m}}, T) - \Phi(i, \bar{\Sigma}_{i,n_i-1}^{\mathrm{m}}, T).$$

Comparing with (A.1), we conclude that

$$\mathbb{P}[D_{in} = e_{in} \; \forall i \in [N], \; n \in [N_i]] = \prod_{i=1}^{N} \prod_{n=1}^{N_i} \mathbb{P}[D_{in} = e_{in} \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})],$$

which proves that the random variables $\{D_{in} \colon i \in [N], n \in [N_i]\}$ form a Bayesian network on $\bar{G}$.

It remains to demonstrate the correctness of the conditional probabilities stated in (3.13). We only carry out the proof for $N_i > 1$ and $n > 1$; the case $n = 1$ is simpler and can be achieved by a straightforward modification of the following arguments. It is evident by the monotonicity statement in (1) that $e_{i,n-1} = 1$ immediately implies

$$\mathbb{P}[D_{in} = 1 \mid \mathrm{pa}_{\bar{G}}(D_{in}) = \mathrm{pa}_{\bar{G}}(e_{in})] = 1,$$

so only the case $e_{i,n-1} = e_{i,n-2} = \cdots = e_{i1} = 0$ needs to be considered further. Conditional on $D_{i,n-1} = 0$, we can now represent the parents of D_{in} as

$$\mathrm{pa}_{\bar{G}}(D_{in}) = F_n(X_j(T) \colon j \in \mathrm{an}_G(i) \setminus \{i\})$$

with some function F_n whose explicit form can be derived from (3.10). Thus, abbreviating

$$M_{in} := L_i(T) - \sum_{j \in \text{pa}_G(i) \cap \text{nd}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{jN_j}=0\}} - \sum_{j \in \text{pa}_G(i) \cap \text{de}_G(i)} \mathrm{e}^{r_{ij}T} L_{ij} \mathbb{1}_{\{e_{j,n-1}=0\}},$$

we obtain

$$\begin{aligned}
&\mathbb{P}[D_{in} = 1 \mid \text{pa}_{\bar{G}}(D_{in}) = \text{pa}_{\bar{G}}(e_{in})] \\
&\quad = \mathbb{P}[X_i(T) < M_{in} \mid X_i(T) \geq M_{i,n-1}, \; F_n(X_j(T)\colon j \in \text{an}_G(i) \setminus \{i\}) = \text{pa}_{\bar{G}}(e_{in})] \\
&\quad = \mathbb{P}[X_i(T) < M_{in} \mid X_i(T) \geq M_{i,n-1}]
\end{aligned}$$

by the independence of $X_1(T), \ldots, X_N(T)$. This is exactly what formula (3.11) states. ∎

Proof of Theorem 3.12. The proof is completely analogous to that of Theorem 3.10. ∎

Proof of Proposition 4.2. (1)   We first assume that V_1 and V_2 are d-separated in G and pick a chain between some (i_1, N_{i_1}) and (i_k, N_{i_k}) with $i_1 \in V_1$ and $i_k \in V_2$. Observing that $(i,n) \to_{\bar{G}} (i',n')$ implies $i = i'$ or $i \to_G i'$, it follows that the chain under consideration cannot be a path (otherwise also i_1 and i_k would be connected through a path in G, which contradicts d-separation). It can neither be of the form

$$(i_1, N_{i_1}) \leftarrow_{\bar{G}} \cdots \leftarrow_{\bar{G}} (i_j, n_j) \to_{\bar{G}} \cdots \to_{\bar{G}} (i_k, N_{i_k}),$$

as this would entail a chain in G of the same form, which would again violate the d-separation assumption. As a result, the chain in $\bar{G}$ must contain a type IV structure and is therefore blocked. The other direction follows similarly, and we omit the details.

(2)   Every chain $(i_1, n_1) \rightleftharpoons_{\bar{G}} \cdots \rightleftharpoons_{\bar{G}} (i_k, n_k)$ with $i_1 \in V_1$ and $i_k \in V_2$ induces a chain $i_1 \rightleftharpoons_G \cdots \rightleftharpoons_G i_k$ in G, where consecutive copies of a vertex are understood as merged to a single vertex. By assumption, this chain in G must be blocked by V_0 and therefore contain at least one blocking structure of type I–IV. If a blocking structure of type I, II, or III is present, say, with middle vertex i_j, then $i_j \in V_0$ and the corresponding sequence in the original chain in $\bar{G}$ must contain a structure of the same type with some middle vertex (i_j, n) and some $n \in [N_{i_j}]$, hence blocking the chain in $\bar{G}$. Similarly, for the last remaining case, if the chain in G has a blocking structure of type IV, say, again with middle node i_j, then V_0 contains neither i_j nor any of its descendants. Since this immediately implies that $\{(i,n)\colon i \in V_0, \; n \in [N_i]\}$ and $\{(i,n)\colon i = i_j \text{ or } i \in \text{de}_G(i_j), \; n \in [N_i]\}$ are disjoint, and the original chain in $\bar{G}$ necessarily has a type IV structure with some middle vertex (i_j, n) and some $n \in [N_{i_j}]$, it is blocked by $\{(i,n)\colon i \in V_0, \; n \in [N_i]\}$. ∎

Proof of Proposition 4.4. (1) is obviously true. In (2) the statement for the absolute systemic impact follows from the definition, while it follows from (4) for the relative systemic impact. (3) is a well-known result (see Proposition 4.2 and the following remark in Levin, Peres, and Wilmer (2009)). Next, (4) is a consequence of an elementary inequality: If a/b and c/d are bounded by M, then also $(a+c)/(b+d)$. Therefore, the maximum in (4.4) is attained on J. Finally, for (5) we notice that if I_2 and J are d-separated in G given I_1, then $\{(i, N_i)\colon i \in I_2\}$ and $\{(j, N_j)\colon j \in J\}$ are d-separated in $\bar{G}$ given $\{(i,n)\colon i \in I_1, \; n \in [N_i]\}$. Since observing $S_{I_1} = 0$ means $S_{in} = 0$ for all $i \in I_1$ and $n \in [N_i]$, the claim follows from the fact that d-separation entails independence in Bayesian networks. ∎

REFERENCES

D. ACEMOGLU, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2015), *Systemic risk and stability in financial networks*, Amer. Econ. Rev., 105, pp. 564–608.

V. V. ACHARYA, L.H. PEDERSEN, T. PHILIPPON, AND M. RICHARDSON (2016), *Measuring systemic risk*, Rev. Financ. Stud., 30, pp. 2–47.

T. ADRIAN AND M. K. BRUNNERMEIER (2016), *CoVaR*, Amer. Econ. Rev., 106, pp. 1705–1741.

F. ALLEN AND D. GALE (2000), *Financial contagion*, J. Polit. Econ., 108, pp. 1–33.

H. AMINI, R. CONT, AND A. MINCA (2016), *Resilience to contagion in financial networks*, Math. Finance, 26, pp. 329–365.

K. ANAND, B. CRAIG, AND G. VON PETER (2015), *Filling in the blanks: Network structure and interbank contagion*, Quant. Finance, 15, pp. 625–636.

S. BATTISTON, M. PULIGA, R. KAUSHIK, P. TASCA, AND G. CALDARELLI (2012), *DebtRank: Too central to fail? Financial networks, the FED and systemic risk*, Sci. Rep., 2, pp. 1–6.

F. BIAGINI, J.-P. FOUQUE, M. FRITTELLI, AND T. MEYER-BRANDIS, *A unified approach to systemic risk measures via acceptance sets*, Math. Finance, to appear.

C. BROWNLEES AND R. F. ENGLE (2017), *SRISK: A conditional capital shortfall measure of systemic risk*, Rev. Financ. Stud., 30, pp. 48–79.

F. CASTIGLIONESI AND M. EBOLI (2015), *Liquidity Flows in Interbank Networks*, preprint, at https://sites.google.com/site/fabiocasti2310/research-1.

H. CHAN AND A. DARWICHE (2005), *A distance measure for bounding probabilistic belief change*, Internat. J. Approx. Reason., 38, pp. 149–174.

C. CHONG AND C. KLÜPPELBERG (2017), *Partial Mean Field Limits in Heterogeneous Networks*, preprint, http://arxiv.org/abs/1507.01905.

R. CONT, A. MOUSSA, AND E. B. SANTOS (2013), *Network structure and systemic risk in banking systems*, in Handbook on Systemic Risk, J.-P. Fouque and J. A. Langsam, eds., Cambridge University Press, Cambridge, pp. 327–367.

B. CRAIG AND G. VON PETER (2014), *Interbank tiering and money center banks*, J. Financ. Intermed., 23, pp. 322–347.

N. DETERING, T. MEYER-BRANDIS, K. PANAGIOTOU, AND D. RITTER (2017), *Managing Default Contagion in Inhomogeneous Financial Networks*, preprint, http://arxiv.org/abs/1610.09542.

F. X. DIEBOLD AND K. YILMAZ (2014), *On the network topology of variance decompositions: Measuring the connectedness of financial firms*, J. Econom., 182, pp. 119–134.

L. EISENBERG AND T. H. NOE (2001), *Systemic risk in financial systems*, Manag. Sci., 47, pp. 236–249.

M. ELLIOTT, B. GOLUB, AND M. O. JACKSON (2011), *Financial networks and contagion*, Amer. Econ. Rev., 104, pp. 3115–3153.

H. ELSINGER, A. LEHAR, AND M. SUMMER (2006), *Risk assessment for banking systems*, Manag. Sci., 52, pp. 1301–1314.

Z. FEINSTEIN, B. RUDLOFF, AND S. WEBER (2017), *Measures of systemic risk*, SIAM J. Financial Math., 8, pp. 672–708.

J.-P. FOUQUE AND L.-H. SUN (2013), *Systemic risk illustrated*, in Handbook on Systemic Risk, J.-P. Fouque and J. A. Langsam, eds., Cambridge University Press, Cambridge, pp. 444–452.

X. FREIXAS, B. M. PARIGI, AND J.-C. ROCHET (2000), *Systemic risk, interbank relations, and liquidity provision by the central bank*, J. Money Credit Bank., 32, pp. 611–638.

P. GAI AND S. KAPADIA (2010), *Contagion in financial networks*, Proc. R. Soc. A, 466, pp. 2401–2423.

A. GANDY AND L. A. M. VERAART (2016), *A Bayesian methodology for systemic risk assessment in financial networks*, Manag. Sci., 63, pp. 4428–4446.

P. GLASSERMAN AND H. P. YOUNG (2015), *How likely is contagion in financial networks?*, J. Bank. Finance, 50, pp. 383–399.

C. Gouriéroux, J.-C. Héam, and A. Monfort (2012), *Bilateral exposures and systemic solvency risk*, Can. J. Econ., 45, pp. 1273–1309.

T. R. Hurd (2016), *Contagion! Systemic Risk in Financial Networks*, Springer, Cham.

O. Kley, C. Klüppelberg, and L. Reichel (2015), *Systemic risk through contagion in a core–periphery structured banking network*, in Advances in Mathematics of Finance, A. Palczewski and Ł. Stettner, eds., Banach Center Publications, Warsaw, pp. 133–149.

D. Koller and N. Friedman (2009), *Probabilistic Graphical Models*, MIT Press, Cambridge, MA.

M. Kusnetsov and L. A. M. Veraart (2016), *Interbank Clearing in Financial Networks with Multiple Maturities*, preprint, https://ssrn.com/abstract=2854733.

S. Langfield, Z. Liu, and T. Ota (2014), *Mapping the UK interbank system*, J. Bank. Finance, 45, pp. 288–303.

S. L. Lauritzen (1996), *Graphical Models*, Clarendon Press, Oxford.

D. A. Levin, Y. Peres, and E. L. Wilmer (2009), *Markov Chains and Mixing Times*, American Mathematical Society, Providence, RI.

C. Meek (1995), *Strong completeness and faithfulness in Bayesian networks*, Uncertainty in Artificial Intelligence, in Proceedings of the Eleventh Conference (1995), P. Besnard, ed., Morgan Kaufmann, San Francisco, CA, pp. 411–418.

C. Memmel, A. Sachs, and I. Stein (2012), *Contagion in the interbank market with stochastic loss given default*, Int. J. Cent. Bank., 8, pp. 177–206.

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

K. P. Murphy (2001), *The Bayes Net Toolbox for MATLAB*, in E. J. Wegman, A. Braverman, A. Goodman, and P. Smyth, eds., *Computing Science and Statistics*, vol. 33, Interface Foundation of North America, Fairfax Station, VA, pp. 331–350.

J. Pearl (1988), *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann, San Francisco, CA.

J. A. Pinto (1986), *Relevance Based Propagation in Bayesian Networks*, Master's thesis, University of California, Berkeley, ftp://ftp.cs.ucla.edu/tech-report/198_-reports/860098.pdf.

L. C. G. Rogers and L. A. M. Veraart (2013), *Failure and rescue in an interbank network*, Manag. Sci., 59, pp. 882–898.

C. Upper (2011), *Simulation methods to assess the danger of contagion in interbank markets*, J. Financ. Stab., 7, pp. 111–125.

T. van Erven and P. Harremoës (2014), *Rényi divergence and Kullback-Leibler divergence*, IEEE Trans. Inform. Theory, 60, pp. 3797–3820.