entropy

MDPI

Article

Systemic Risk in a Multilayer Financial Network: A Criticality-Based Perspective

Ning Xi †, Yamin Dou † and Yazheng Dang *

Department of Systems Science, Business School, University of Shanghai for Science and Technology, Shanghai 200093, China; nxi@usst.edu.cn (N.X.); dym1182022@163.com (Y.D.)

* Correspondence: jgdyz@163.com

† These authors contributed equally to this work.

Abstract

Banks and firms are connected through multiple credit relationships that can transmit financial shocks across sectors, amplify default contagion, and ultimately lead to system-wide collapse. Understanding systemic risk therefore requires not only examining default contagion mechanisms but also characterizing system-wide collapse. To this end, we construct a three-layer financial network model and examine its response to exogenous shocks. Under the baseline configuration, numerical simulations reveal pronounced bimodal distributions of bankruptcy fractions, indicating that the sharp increase in average bankruptcy fractions mainly reflects a shift in the relative frequencies of limited-contagion and system-wide collapse states. Susceptibility and finite-size analyses provide further evidence of phase-transition-like behavior. Building on these findings, we characterize systemic risk using complementary measures of collapse likelihood, shock tolerance, and financial damage, measured by system-wide collapse probability, collapse thresholds, and layer-specific loss ratios, respectively. Matched counterfactual experiments show that two-layer configurations without endogenous interfirm credit propagation have higher system-wide collapse thresholds than the three-layer model under both recovery-matching conventions, although the increase is much larger under recovery-rate matching than under recovery-amount matching. Parameter experiments further show how balance-sheet composition, interfirm link formation, and credit exposure are associated with shock tolerance.

Keywords: system-wide collapse; criticality; multilayer financial network; default contagion; systemic risk; interfirm trade credit

Check for updates

Academic Editors: Gaogao Dong and Duxin Chen

Received: 10 August 2026
Revised: 30 September 2026
Accepted: 15 October 2026
Published: 19 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The 2008 global financial crisis precipitated severe financial turmoil and a deep global recession, shifting the focus of academic and institutional discourse toward understanding and mitigating systemic risk [1]. In the years that followed, however, the global financial system was repeatedly tested by major shocks, underscoring the persistence of these vulnerabilities. The COVID-19 pandemic, for example, inflicted profound and widespread damage on global finance [2]: lockdown measures and supply chain disruptions produced acute economic stagnation and dramatic market volatility, renewing concerns about systemic risk. More recently, the failure of Silicon Valley Bank and the emergency rescue of Credit Suisse exposed critical structural weaknesses within the financial system, triggering contagion effects that reverberated through depositors, borrowers, cryptocurrency markets, and interconnected banks [3–5]. Although these banking stress events did not escalate

into a full-blown global crisis comparable to 2008, they caused substantial disruptions to financial markets and the broader economy. These episodes highlight the importance of understanding how interconnected financial relationships transmit and amplify distress, potentially leading to system-wide collapse.

Early research on financial networks and systemic risk focused primarily on interbank networks, in which banks are connected through direct credit exposures. Allen and Gale were among the first to introduce network structure into the study of systemic risk [6]. By analyzing a simple four-bank network, they found that under their model assumptions, a complete network is more resilient than an incomplete one. Building on similar simplified frameworks, subsequent studies have confirmed that network structure significantly affects systemic risk in banking systems [7,8]. The development of complex network science has provided a broad set of analytical tools for systemic risk research, enabling theoretical modeling, empirical characterization, and quantitative risk assessment of financial networks. Theoretically, researchers have developed multi-agent models based on network theory, comparing systemic risk across interbank networks with different statistical properties [9–11]. Empirically, analyses using real bank data have shown that interbank networks often display small-world or scale-free topological features [12–15]. Additionally, researchers have conducted quantitative assessments of systemic risk in specific national interbank networks. For example, Upper and Worms estimated that the failure of a single bank in Germany could lead to a breakdown affecting up to 15% of the system's assets [16], while Degryse and Nguyen found that the Belgian banking system's evolution from a complete structure to a multiple-money-center network structure helped reduce contagion risk [17].

These studies have primarily examined single-layer interbank networks, in which only one type of financial relationship is modeled at a time. However, in reality, banks are connected by multiple types of financial relationships, each of which can be represented as a distinct network layer. For instance, Langfield et al. distinguished interbank exposure and interbank funding networks as separate layers within a multilayer interbank network, using financial instruments extracted from UK regulatory data [18]. Risk propagation in multilayer networks can exhibit non-additive effects, whereby the combined contagion effects across layers may substantially exceed the sum of individual-layer effects [19], a finding supported by empirical evidence [20]. Moreover, Bargigli et al. demonstrated that aggregating different network layers can lead to biased estimates of systemic risk [21]. These findings indicate that multilayer network models can better capture distinct channels of risk transmission and reduce biases arising from the aggregation of different network layers. A comprehensive review of financial networks, encompassing both single-layer and multilayer frameworks, is provided by Bardoscia et al. [22].

Although multilayer frameworks have improved the modeling of interbank relationships, the role of firms in systemic risk remains comparatively underexplored. In modern economies, firms are tightly connected to banks through credit relationships, so shocks affecting firms can be transmitted to banks and further amplified through interbank linkages. Using Brazilian data, empirical studies have constructed multilayer networks that include both bank–firm and interbank relationships, showing that firm distress can propagate through the banking network and highlighting the importance of incorporating firms into systemic risk estimation [23,24]. Lux proposed a random bipartite model of bank–firm credit relationships that reproduces key statistical regularities observed in loan data across multiple countries. Simulation results indicate that the bank–firm lending channel can play a more important role than the interbank lending channel in default propagation [25]. Building on this model, Gao introduced bank asset portfolios and examined how systemic risk evolves under different shock types, including firm defaults, asset devaluation,

and bank failures [26]. Nevertheless, direct interfirm credit relationships remain insufficiently incorporated in most bank–firm network studies. In reality, firms are embedded in large-scale production networks, acting simultaneously as customers and suppliers. Since suppliers frequently extend trade credit to their customers, production linkages are commonly accompanied by interfirm credit exposures, creating an additional channel through which credit shocks can propagate. A growing body of research has documented important statistical and dynamic features of production and interfirm networks [27–37]. Notably, an empirical study based on Uruguayan data shows that neglecting interfirm credit exposures can lead to a substantial underestimation of systemic risk [38].

In this context, Li and Sui constructed an endogenous financial network model incorporating interbank, bank–firm, and interfirm relationships. By introducing different partner selection mechanisms, they examined how network evolution influences default contagion [39]. Their work represents an important step toward incorporating interfirm relationships into multilayer models of default contagion. Subsequent studies have further extended multilayer bank–firm modeling frameworks. For example, Jin et al. developed a dynamic multilayer financial network model incorporating interbank lending, bank–firm credit relationships, and cross-shareholding among firms, and examined how network structure and attack strategies influence systemic stability [40]. Similarly, Lei et al. constructed a multilayer bank–firm network integrating interbank lending, bank–firm credit, interfirm trade credit, guarantee relationships, and cross-holdings and analyzed how different types of shocks propagate through these interconnected channels [41]. Taken together, these studies have examined how network structure, credit relationships, and shock mechanisms affect default contagion in multilayer financial networks. However, the statistical characteristics of system-wide collapse remain less well understood in these models. A sharp increase in average bankruptcy fractions alone cannot distinguish a gradual shift in the distribution toward higher values from a change in the relative frequencies of distinct states, nor does it establish how system-wide collapse behavior depends on system size. These questions motivate a criticality-based analysis of system-wide collapse.

To address these questions, we adopt a criticality-based perspective to analyze systemic risk in multilayer financial networks. Building on the Lux model, we construct a three-layer financial network comprising interbank lending, bank–firm lending, and interfirm trade credit within a unified balance-sheet framework. Treating external shock intensity as a control parameter, we examine bankruptcy-fraction distributions, susceptibility, and finite-size behavior to characterize system-wide collapse. We then use three complementary measures to characterize systemic risk: system-wide collapse probability to quantify collapse likelihood at a given shock intensity; the system-wide collapse threshold to characterize shock tolerance, supplemented by sector-specific collapse thresholds to distinguish bank and firm responses; and layer-specific loss ratios to capture the severity of financial damage.

Within this framework, we investigate the contribution of endogenous interfirm credit propagation and the relationships between selected balance-sheet and interfirm-layer parameters and shock tolerance. Matched counterfactual experiments show that removing endogenous interfirm credit propagation raises the system-wide collapse threshold under both recovery-matching conventions, although the magnitude of the increase depends on how recovery resources are matched. We further examine how bank capitalization, firm cash holdings, interfirm link formation, and credit exposure are associated with shock tolerance and sectoral responses.

The remainder of this paper is organized as follows. Section 2 presents the multilayer financial network model and default contagion mechanisms. Section 3 examines system-wide collapse behavior, system-wide collapse probabilities, system-wide and sector-specific

collapse thresholds, and layer-specific loss ratios, followed by finite-size analysis, matched counterfactual comparisons of endogenous interfirm credit propagation, balance-sheet and interfirm-layer parameter experiments, and robustness checks. Section 4 summarizes the main findings and discusses the study's limitations.

2. Model

We model a financial system consisting of banks and firms that are connected through three types of credit relationships, each represented as a distinct network layer. Banks provide loans to firms, forming the bank–firm network; banks engage in interbank lending, forming the interbank network; and firms transact with other firms through trade credit, forming the interfirm network. These financial activities are captured in the balance sheets of banks and firms. By leveraging the accounting identity embedded in the balance sheets, the three networks are structurally coupled, jointly forming a complex multilayer financial network.

2.1. Bank–Firm Network

Empirical studies have identified several stylized characteristics of bank–firm networks. Axtell's analysis of U.S. firms shows that firm size follows a power-law distribution with an exponent close to one [42]. Similarly, Bremus et al. analyzed cross-country data and found that the right tail of the bank size distribution exhibits a power-law pattern [43]. De Masi and Gallegati, based on their study of the Italian market, drew two key conclusions regarding the bank–firm network. First, banks have a much higher average degree than firms. Second, connectivity is heterogeneous and exhibits clear size-related patterns on both sides: bank degree is strongly associated with bank size, while large firms tend to be financed by multiple banks [44].

Motivated by these empirical findings, we adopt the network-generation algorithm proposed by Lux to construct the bank–firm layer [25]. The statistical properties of the networks generated under our baseline parameter settings are evaluated separately in Section 3.1.1. Let N_b denote the number of banks, N_f the number of firms, and λ_f the average degree of firms in the bank–firm network. The average degree of banks is then given by $\lambda_b = \lambda_f N_f / N_b$. For the banking sector, we assume that the total assets A_i^b ($i = 1, 2, \ldots, N_b$) of bank i follow a truncated Pareto distribution with the probability density function:

$$f(A_i^b) = \frac{\alpha (A_b^{LB})^{\alpha} (A_i^b)^{-\alpha-1}}{1 - \left(A_b^{LB}/A_b^{UB}\right)^{\alpha}}, \tag{1}$$

where A_b^{LB} and A_b^{UB} represent the lower and upper bounds of bank total assets, respectively, and α is the Pareto exponent. The total assets of each bank are generated by sampling from this distribution. In line with the empirical findings of De Masi and Gallegati [44], we assume that the expected degree $\widetilde{\lambda}_i^b$ of bank i is proportional to its total assets A_i^b, such that $\widetilde{\lambda}_i^b = \lambda_b A_i^b / \bar{A}_b$, where $\bar{A}_b$ denotes the sample mean of bank total assets. Furthermore, for simplicity, we assume that each bank allocates a fixed proportion θ_{bf} of its total assets to total loans to firms, as follows:

$$L_i^{bf} = \theta_{bf} A_i^b, \tag{2}$$

where θ_{bf} is identical across banks. Since firm assets often serve as an important basis for bank lending, firms' total bank loans tend to increase with their total assets. Consequently, the distribution of total bank loans across firms is expected to exhibit a similar shape to the distribution of firms' total assets, which follows a power-law distribution. Accordingly, for the firm sector, we assume that the expected total bank loans $\widetilde{L}_j^{fb}$ ($j = 1, 2, \ldots, N_f$) of firm j follow a distribution analogous to Equation (1), with the same Pareto exponent

α, but with modified bounds. Specifically, the lower and upper bounds are defined as $L_{fb}^{LB} = \theta_{bf} A_b^{LB} N_b / N_f$ and $L_{fb}^{UB} = \theta_{bf} A_b^{UB} N_b / N_f$, respectively. These bounds specify the range of expected total bank loans for each firm. The expected total bank loans for each firm are generated by sampling from this distribution. We further assume that the expected degree $\widetilde{\lambda}_j^f$ of firm j is proportional to its expected total bank loans: $\widetilde{\lambda}_j^f = \lambda_f \widetilde{L}_j^{fb} / \bar{\widetilde{L}}^{fb}$, where $\bar{\widetilde{L}}^{fb}$ denotes the sample mean of firms' expected total bank loans.

By construction, the sums of expected degrees on the bank and firm sides are equal: $\sum_i \widetilde{\lambda}_i^b = \lambda_b N_b = \lambda_f N_f = \sum_j \widetilde{\lambda}_j^f$. We therefore set the target number of links to $\lambda_f N_f$. The network is then constructed using weighted random sampling. Specifically, at each step, a bank and a firm are selected independently, with probabilities proportional to their respective expected degrees. A link is established only if the selected pair is not already connected. This iterative process continues until the target number of links is reached. To further ensure that each entity has at least one link, each unconnected bank is connected to a firm selected with probability proportional to the firm's expected degree. Similarly, each unconnected firm is connected to a bank selected with probability proportional to the bank's expected degree. In the resulting network, all links are directed from banks to firms, indicating loan provision from banks to firms. This structure is represented by a directed adjacency matrix $M^{bf} = [m_{ij}^{bf}]$ of dimension $N_b \times N_f$, where $m_{ij}^{bf} = 1$ if a directed link from bank i to firm j exists, and 0 otherwise. We assume that each bank distributes its total loans to firms equally across all its connected firms. Accordingly, the loan from bank i to firm j is given by

$$L_{ij}^{bf} = \frac{L_i^{bf} m_{ij}^{bf}}{\lambda_i^b}, \tag{3}$$

where $\lambda_i^b = \sum_{j=1}^{N_f} m_{ij}^{bf}$ represents the degree of bank i. It follows that the total bank loans of firm j are calculated as

$$L_j^{fb} = \sum_{i=1}^{N_b} L_{ij}^{bf}. \tag{4}$$

The resulting total bank loans of firms are determined by the generated network and the equal-allocation rule and may differ from their sampled expected values.

2.2. Interbank Network

Empirical studies have examined the degree distributions and mixing patterns of interbank networks [12,13]. Motivated by these studies, Montagna and Lux proposed a Monte Carlo framework for generating interbank networks with asset-dependent link formation [45]. This algorithm has been widely applied in studies on systemic risk in the banking sector [25,26]. We adopt this framework to construct the interbank layer. Its realized degree distributions and mixing patterns under our baseline settings are examined in Section 3.1.1. The probability of a directed link from bank i to bank j—representing an interbank loan extended by bank i to bank j—is given by

$$p_{ij}^{bb} = d_b \left(\frac{A_i^b}{A_b^{\max}} \right)^{\alpha_1} \left(\frac{A_j^b}{A_b^{\max}} \right)^{\alpha_2}, \tag{5}$$

where α_1, α_2, and d_b are parameters; A_i^b and A_j^b denote the total assets of bank i and bank j, respectively; and $A_b^{\max}$ represents the maximum total assets among all banks. Subsequently, the directed adjacency matrix M^{bb} of the interbank network is constructed, where diagonal entries are set to zero, and off-diagonal entries are determined by

$$m_{ij}^{bb} = \begin{cases} 1, & \text{with probability } p_{ij}^{bb} \\ 0, & \text{with probability } 1 - p_{ij}^{bb} \end{cases}. \tag{6}$$

To ensure that each bank has at least one outgoing link, we assign an outgoing link from each bank with zero out-degree to another bank chosen with probability proportional to its total assets. The adjacency matrix M^{bb} is updated accordingly.

After establishing the network topology, we define the weights of the credit relationships. For a given bank, the total loans it extends to other banks are referred to as interbank lending, while the total loans it receives from other banks are referred to as interbank borrowing. To simplify the model, we assume that each bank allocates a fixed proportion θ_{bb} of its total assets to interbank lending:

$$L_i^{IBL} = \theta_{bb} A_i^b, \tag{7}$$

where θ_{bb} is identical across banks. Furthermore, we assume that the interbank lending of bank i is allocated among borrowing banks in proportion to their total assets. Thus, the interbank loan from bank i to bank j is given by

$$L_{ij}^{bb} = L_i^{IBL} \frac{A_j^b m_{ij}^{bb}}{\sum_{h=1}^{N_b} A_h^b m_{ih}^{bb}}. \tag{8}$$

Accordingly, the interbank borrowing of bank i can be expressed as

$$L_i^{IBB} = \sum_{k=1}^{N_b} L_{ki}^{bb}. \tag{9}$$

2.3. Interfirm Network

Empirical studies have investigated heterogeneous connectivity and size-related linking patterns in interfirm networks [46–48]. Motivated by these findings, we adopt the asset-dependent network-generation framework of Montagna and Lux to construct the interfirm layer [45]. The properties of the resulting network are evaluated in Section 3.1.1. The construction process begins by determining the total assets of each firm. For simplicity, we assume that firms' total bank loans are proportional to their total assets. Given firm j's total bank loans L_j^{fb} in Equation (4), its total assets can then be inferred as

$$A_j^f = \frac{L_j^{fb}}{\mu_{fb}}, \tag{10}$$

where μ_{fb} represents the share of a firm's total bank loans in its total assets and is assumed to be identical across firms. Next, the probability of a directed link from firm i to firm j is defined as

$$p_{ij}^{ff} = d_f \left(\frac{A_i^f}{A_f^{\max}} \right)^{\alpha_1} \left(\frac{A_j^f}{A_f^{\max}} \right)^{\alpha_2}, \tag{11}$$

where α_1, α_2, and d_f are parameters; A_i^f and A_j^f denote the total assets of firm i and firm j, respectively; and $A_f^{\max}$ represents the maximum total assets among all firms. Here, a directed link from firm i to firm j indicates that firm i is the supplier and firm j is the customer, with firm i extending trade credit to firm j. For simplicity, α_1 and α_2 are set equal to those in Equation (5). Based on this probability, the directed adjacency matrix M^{ff} of the interfirm network is constructed with diagonal entries equal to 0 and off-diagonal entries determined by

$$m_{ij}^{ff} = \begin{cases} 1, & \text{with probability } p_{ij}^{ff} \\ 0, & \text{with probability } 1 - p_{ij}^{ff} \end{cases}. \tag{12}$$

To ensure that each firm has at least one incoming link, any firm with zero in-degree receives a directed link from another firm selected with probability proportional to its total assets. The adjacency matrix M^{ff} is updated accordingly.

In real-world economies, transactions between firms are often conducted through trade credit. Under such arrangements, suppliers deliver goods or services first, and customers settle payments at a later, pre-agreed date. This practice essentially constitutes a form of short-term financing, recorded as accounts receivable on suppliers' balance sheets and as accounts payable on customers' balance sheets. For simplicity, we assume that firm j, acting as a customer, has accounts payable equal to a fixed proportion μ_{ff} of its total assets. Accordingly, firm j's accounts payable are given by

$$L_j^{AP} = \mu_{ff} A_j^f, \tag{13}$$

where μ_{ff} is identical across firms. Moreover, we assume that firm j allocates its accounts payable among its suppliers in proportion to their total assets. Therefore, the trade credit extended by firm i to firm j is determined as follows:

$$L_{ij}^{ff} = L_j^{AP} \frac{A_i^f m_{ij}^{ff}}{\sum_{h=1}^{N_f} A_h^f m_{hj}^{ff}}. \tag{14}$$

Consequently, firm i's accounts receivable are obtained as the sum of all such trade credits extended to its customers:

$$L_i^{AR} = \sum_{j=1}^{N_f} L_{ij}^{ff}. \tag{15}$$

2.4. Balance Sheets of Banks and Firms

The balance sheet presents the financial position of banks and firms through a set of financial variables and is constructed in accordance with the fundamental accounting identity that total assets equal total liabilities. Its structure comprises distinct categories for assets and liabilities. Based on the model setup, we simplify the conventional balance sheet while ensuring that it still accurately captures the key financial characteristics of banks and firms.

The balance sheet of bank i is shown in Table 1. On the asset side, the total assets A_i^b of bank i comprise total loans to firms L_i^{bf}, interbank lending L_i^{IBL}, and cash C_i^b, as expressed by

$$A_i^b = L_i^{bf} + L_i^{\mathrm{IBL}} + C_i^b. \tag{16}$$

On the liability side, the total liabilities L_i^b of bank i consist of equity E_i^b, interbank borrowing L_i^{IBB}, and deposits D_i^b, as given by

$$L_i^b = E_i^b + L_i^{\mathrm{IBB}} + D_i^b. \tag{17}$$

As required by the accounting identity, the total assets and total liabilities of bank i must be equal, yielding

$$A_i^b = L_i^b. \tag{18}$$

The variables are specified as follows. The total assets A_i^b are obtained by sampling from the truncated Pareto distribution defined in Equation (1). The total loans to firms L_i^{bf} are

computed using Equation (2). The interbank lending L_i^{IBL} is derived from Equation (7). Cash C_i^b is determined from Equation (16). According to Equation (18), the total liabilities L_i^b are equal to the total assets A_i^b. For simplicity, we assume that the share of equity in total assets, denoted by ω_b, is identical across banks. The equity is therefore expressed as

$$E_i^b = \omega_b A_i^b. \tag{19}$$

The interbank borrowing L_i^{IBB} is obtained from Equation (9). The deposits D_i^b are determined from Equation (17). Therefore, once the three network layers are established, the balance sheets of all banks can be generated to represent their financial positions.

Table 1. Balance sheet of bank i.

Total Assets A_i^b	**Total Liabilities L_i^b**
Total loans to firms L_i^{bf}	Equity E_i^b
Interbank lending L_i^{IBL}	Interbank borrowing L_i^{IBB}
Cash C_i^b	Deposits D_i^b

The balance sheet of firm j is shown in Table 2. On the asset side, the total assets A_j^f of firm j comprise accounts receivable L_j^{AR}, fixed assets K_j^f, and cash C_j^f, as defined by

$$A_j^f = L_j^{AR} + K_j^f + C_j^f. \tag{20}$$

On the liability side, the total liabilities L_j^f of firm j consist of equity E_j^f, accounts payable L_j^{AP}, and total bank loans L_j^{fb}, as formulated in

$$L_j^f = E_j^f + L_j^{AP} + L_j^{fb}. \tag{21}$$

In accordance with the accounting identity, the total assets and total liabilities of firm j must be equal, leading to

$$A_j^f = L_j^f. \tag{22}$$

The variables are specified as follows. The total assets A_j^f are obtained from Equation (10). The accounts receivable L_j^{AR} are given by Equation (15). Let μ_c denote the share of cash in total assets. Cash is then calculated as

$$C_j^f = \mu_c A_j^f, \tag{23}$$

where μ_c is identical across firms. The fixed assets K_j^f can be determined from Equation (20). Based on Equation (22), the total liabilities L_j^f are equal to the total assets A_j^f. The accounts payable L_j^{AP} are specified by Equation (13). The total bank loans L_j^{fb} are computed using Equation (4). The equity E_j^f is then derived from Equation (21). Therefore, after constructing the three network layers, the balance sheets of all firms can be generated to represent their financial positions. To ensure balance-sheet feasibility, candidate financial configurations with negative bank deposits, firm fixed assets, or firm equity are rejected, and the network and associated balance sheets are regenerated until these quantities are nonnegative. Sections 2.1–2.4 describe the initial configuration of the financial system, including the three network layers and the balance sheets of banks and firms. For simplicity, all variables in this part are presented without explicit time indices and refer to their initial values.

Table 2. Balance sheet of firm j.

Total Assets A_j^f	**Total Liabilities** L_j^f
Accounts receivable L_j^{AR}	Equity E_j^f
Fixed assets K_j^f	Accounts payable L_j^{AP}
Cash C_j^f	Total bank loans L_j^{fb}

2.5. Default Contagion Mechanism

We model the contagion dynamics as a discrete-time iterative process. Hereafter, the variables introduced in the preceding sections evolve dynamically as the contagion process unfolds. Within each iteration, events such as liquidations, recoveries, repayments, and defaults occur sequentially, and the affected variables are updated instantaneously. These updated values then serve as inputs for subsequent calculations within the same iteration. Consequently, certain balance-sheet items—such as banks' cash assets, total loans to firms, and firms' accounts receivable and accounts payable—may undergo multiple updates within a single iteration as different events occur. To maintain conciseness, variables are presented without time indices unless a specific temporal reference is required to avoid ambiguity. These variables without time indices denote their most up-to-date values at the time of calculation. For reproducibility, the complete pseudocode for the shock, default-contagion, liquidation, and default-loss accounting procedures is provided in Section S1 of the Supplementary Materials (Algorithms S1 and S2).

Within each iteration, the financial positions of banks and firms are updated sequentially as events occur, following the rules below. In the first iteration, a random permutation of firms is generated, and firms are selected in that order until their cumulative pre-shock assets first reach or exceed a fraction ε of the total pre-shock assets of the firm sector. The realized shock share may therefore slightly exceed the target value ε and is recorded for each simulation run. The selected firms are then subjected to exogenous shocks sequentially in the same order. Exogenous shocks are imposed exclusively on the firm sector, allowing us to examine how firm-originated disturbances propagate through the credit linkages of the multilayer financial network and subsequently affect the banking sector. For each selected firm, fixed assets and cash are set to zero on its balance sheet, and its equity is adjusted accordingly to maintain the accounting identity. If a shocked firm's equity becomes negative (i.e., $E_j^f < 0$), it is deemed bankrupt and liquidated immediately, before the next shocked firm is processed, so that firms shocked later may face balance sheets already altered by these liquidations. Any bank or firm outside the initially shocked set that becomes bankrupt as a result of these liquidations is liquidated in the next iteration. Thereafter, bankruptcies generated in one iteration are liquidated in the next, and the process continues until no new bankrupt entities emerge. In each iteration after the first, banks are processed before firms, and entities within each sector are processed in ascending order of their indices. Because balance-sheet variables are updated immediately after each liquidation, earlier liquidations may affect entities processed later in the same iteration. We therefore also examine the sensitivity of the results to alternative liquidation priorities (Section 3.6.2). The liquidation process involves bankrupt entities collecting available funds and repaying their debts according to a specific priority order. The following subsections describe in detail the liquidation mechanisms for bankrupt firms and banks.

For clarity, we distinguish two failure modes. Bankruptcy triggered by negative equity is referred to as insolvency, whereas bankruptcy triggered by the exhaustion of available cash when meeting repayment obligations is referred to as illiquidity. An entity may exhibit both failure modes.

2.5.1. Liquidation of Bankrupt Firms

The liquidation funds of a bankrupt firm are sourced from three channels: recovered accounts receivable, the gross liquidation value of fixed assets, and remaining cash assets. The liquidation process unfolds as follows. First, the bankrupt firm recovers outstanding accounts receivable from its debtor firms. To maintain model tractability, we assume that a debtor firm can only settle its debts using its cash assets and cannot generate additional funds by collecting its accounts receivable or liquidating fixed assets. This assumption prevents recursive debt recovery loops across firms, ensuring that liquidation cascades remain finite and well-defined within each iteration. If a debtor firm has insufficient cash assets, it transfers all its available cash assets and is then declared bankrupt due to illiquidity, entering liquidation in the next iteration. The accounts receivable recovered by bankrupt firm j are given by

$$R_j^{AR} = \sum_{k=1}^{N_f} \min\left(L_{jk}^{ff}, C_k^f\right), \tag{24}$$

where L_{jk}^{ff} and C_k^f, defined in Sections 2.3 and 2.4, represent the amount owed by firm k to firm j and the cash assets of firm k, respectively. Whenever a claim is processed during liquidation, the corresponding claim–liability relationship is closed after the resulting payment, which may be zero; any unrecovered amount is written off immediately and is not carried forward for recovery in subsequent iterations. Next, the bankrupt firm liquidates its fixed assets. It then pools the recovered receivables R_j^{AR}, the gross liquidation value of fixed assets K_j^f, and its cash assets C_j^f to form a liquidation fund. To account for liquidation-related costs, such as administrative fees, employee wages, taxes, and liquidation discounts on fixed assets, we introduce the firm liquidation recovery rate δ_f, where $0 < \delta_f < 1$. The actual funds available for debt repayment are then computed as

$$F_j^f = \delta_f\left(R_j^{AR} + K_j^f + C_j^f\right). \tag{25}$$

Debt repayment follows a strict priority order: bank loans are repaid first, followed by interfirm debts (i.e., accounts payable). This priority order reflects the common institutional arrangement in which bank loans generally have higher repayment priority than trade credit. If the available funds F_j^f are sufficient to cover all bank loans, they are repaid in full; otherwise, repayments are allocated on a pro rata basis according to each bank's claim. The amount repaid to creditor bank i is

$$R_{ij}^{fb} = \begin{cases} L_{ij}^{bf}, & F_j^f \geq L_j^{fb} \\ \dfrac{L_{ij}^{bf}}{L_j^{fb}} F_j^f, & F_j^f < L_j^{fb} \end{cases}, \tag{26}$$

where L_{ij}^{bf} and L_j^{fb}, defined in Section 2.1, represent the loan from bank i to firm j and the total bank loans of firm j, respectively. Upon receiving repayment, the creditor bank increases its cash assets by the repayment amount and removes the entire corresponding loan from its total loans to firms. The shortfall between the original loan amount and the repayment amount is deducted from the creditor bank's equity as a loss to maintain the accounting identity. If the creditor bank's equity becomes negative as a result, it is declared bankrupt due to insolvency and enters liquidation in the next iteration. The remaining funds after repayment to banks are calculated as

$$F_j^{fr} = \max\left(F_j^f - L_j^{fb}, 0\right). \tag{27}$$

These funds are subsequently used to settle interfirm debts (i.e., the accounts payable of the bankrupt firm). If sufficient, these debts are paid in full; otherwise, payments are allocated on a pro rata basis according to each creditor firm's claim. The amount paid to creditor firm j' is

$$R_{j'j}^{ff} = \begin{cases} L_{j'j}^{ff}, & F_j^{fr} \geq L_j^{AP} \\ \dfrac{L_{j'j}^{ff}}{L_j^{AP}} F_j^{fr}, & F_j^{fr} < L_j^{AP} \end{cases}, \tag{28}$$

where $L_{j'j}^{ff}$ and L_j^{AP}, defined in Section 2.3, represent the amount owed by firm j to firm j' and the accounts payable of firm j, respectively. The payment to a creditor firm may be partial or even zero, depending on the bankrupt firm's remaining funds. Upon receiving payment, the creditor firm increases its cash assets by the payment amount and removes the entire corresponding receivable from its accounts receivable. The shortfall between the original receivable amount and the payment amount is deducted from the creditor firm's equity as a loss to maintain the accounting identity. If the creditor firm's equity becomes negative as a result, it is declared bankrupt due to insolvency and enters liquidation in the next iteration.

2.5.2. Liquidation of Bankrupt Banks

The liquidation funds of a bankrupt bank are sourced from three channels: recovered interbank lending, recovered loans to firms, and remaining cash assets. The liquidation process unfolds as follows. First, the bankrupt bank recovers its interbank lending. As before, we assume that a debtor bank can only repay its interbank borrowing using its cash assets and cannot generate additional funds by recalling its interbank lending or loans to firms. If a debtor bank has insufficient cash assets, it transfers all its remaining cash assets and is declared bankrupt due to illiquidity, entering liquidation in the next iteration. The interbank lending recovered by bankrupt bank i is given by

$$R_i^{\mathrm{IBL}} = \sum_{k=1}^{N_b} \min\left(L_{ik}^{bb}, C_k^b\right), \tag{29}$$

where L_{ik}^{bb} represents the interbank loan from bank i to bank k (as defined in Section 2.2), and C_k^b represents the cash assets of bank k (as defined in Section 2.4). Next, the bankrupt bank recovers its loans to firms. Similarly, a debtor firm can only repay its bank loan using its cash assets and cannot generate additional funds by collecting its accounts receivable or liquidating fixed assets. If a debtor firm has insufficient cash assets, it transfers all its remaining cash assets and is then declared bankrupt due to illiquidity, entering liquidation in the next iteration. The loans to firms recovered by bankrupt bank i are given by

$$R_i^{bf} = \sum_{k=1}^{N_f} \min\left(L_{ik}^{bf}, C_k^f\right), \tag{30}$$

where L_{ik}^{bf} represents the loan from bank i to firm k (as defined in Section 2.1), and C_k^f represents the cash assets of firm k (as defined in Section 2.4). As in the liquidation of bankrupt firms, each processed claim–liability relationship is closed after settlement, and any unrecovered amount is written off immediately rather than carried forward to subsequent iterations. Bankrupt bank i then pools the recovered interbank lending R_i^{IBL}, the recovered loans to firms R_i^{bf}, and its own cash assets C_i^b to form a liquidation fund. To

account for liquidation-related costs, we introduce the bank liquidation recovery rate δ_b, where $0 < \delta_b < 1$. The actual funds available for debt repayment are then computed as

$$F_i^b = \delta_b \left(R_i^{\mathrm{IBL}} + R_i^{bf} + C_i^b \right). \tag{31}$$

The allocation of these liquidation funds follows a strict priority order. The liquidation funds are first used to repay deposit liabilities. This priority order reflects the common institutional arrangement in which deposit liabilities are given priority over interbank claims. The remaining funds after deposit repayment are denoted by F_i^{br} and calculated as

$$F_i^{br} = \max\left(F_i^b - D_i^b, 0 \right), \tag{32}$$

where D_i^b represents the deposit liabilities of bank i (as defined in Section 2.4). These remaining funds are then used to settle interbank borrowing. If F_i^{br} is sufficient, interbank borrowing is repaid in full; otherwise, repayments are allocated on a pro rata basis according to each creditor bank's claim. The amount repaid to creditor bank i' is calculated as

$$R_{i'i}^{bb} = \begin{cases} L_{i'i}^{bb}, & F_i^{br} \geq L_i^{\mathrm{IBB}} \\ \dfrac{L_{i'i}^{bb}}{L_i^{\mathrm{IBB}}} F_i^{br}, & F_i^{br} < L_i^{\mathrm{IBB}} \end{cases}, \tag{33}$$

where $L_{i'i}^{bb}$ denotes the interbank loan from bank i' to bank i, and L_i^{IBB} represents the interbank borrowing of bank i, both as defined in Section 2.2. Upon receiving repayment, which may be zero if the liquidation funds are insufficient to cover deposit liabilities, the creditor bank increases its cash assets by the repayment amount and removes the entire corresponding interbank loan from its interbank lending. The shortfall between the original loan amount and the repayment amount is deducted from the creditor bank's equity as a loss to maintain the accounting identity. If the creditor bank's equity becomes negative as a result, it is declared bankrupt due to insolvency and enters liquidation in the next iteration.

2.5.3. Quantification of Default Losses

Default loss refers to the reduction in an entity's asset value resulting from counterparties' failure to meet financial obligations, manifested as the unrecovered portion of the entity's financial claims. In the financial system, such losses erode balance sheets, potentially precipitating cascading bankruptcies and risk propagation. To characterize the realized severity and distribution of losses during default contagion, we quantify default losses separately in the bank–firm, interbank, and interfirm layers.

In the bank–firm network, default losses occur when a bankrupt creditor bank fails to fully recover its loans to firms, or when a bankrupt debtor firm cannot fully repay its bank loans. For a bankrupt creditor bank i, the recovered amount of loans to firms R_i^{bf} is given by Equation (30). The corresponding default loss is the unrecovered portion:

$$\mathcal{L}_i^{bf} = L_i^{bf} - R_i^{bf}. \tag{34}$$

For a bankrupt debtor firm j, the available liquidation funds F_j^f are determined by Equation (25). Since these funds are first used to repay bank loans L_j^{fb} (up to the full amount owed), the default loss to creditor banks is the unpaid portion:

$$\mathcal{L}_j^{fb} = \max\left(L_j^{fb} - F_j^f, 0 \right). \tag{35}$$

Let $\mathcal{B}_t$ and $\mathcal{F}_t$ denote the sets of banks and firms liquidated during iteration t, respectively. The total default loss in the bank–firm network during iteration t is

$$\mathcal{L}_{bf}(t) = \sum_{i \in \mathcal{B}_t} \mathcal{L}_i^{bf} + \sum_{j \in \mathcal{F}_t} \mathcal{L}_j^{fb}. \tag{36}$$

To quantify loss severity in the bank–firm layer, we introduce the cumulative default loss ratio ℓ_{bf}. Since default losses in this layer arise directly from bank loans to firms, the initial total amount of such loans provides a natural normalization baseline. This ratio is thus defined as the total default losses during the contagion process normalized by this baseline:

$$\ell_{bf} = \frac{\sum_{t=1}^{T} \mathcal{L}_{bf}(t)}{\sum_{i=1}^{N_b} L_i^{bf}(0)}, \tag{37}$$

where $L_i^{bf}(0)$ denotes the initial loan amount extended by bank i to firms. The contagion process terminates at iteration T when no unliquidated bankrupt entities remain.

In the interbank network, default losses arise when a bankrupt creditor bank fails to fully recover its interbank lending, or when a bankrupt debtor bank cannot fully repay its interbank borrowing. For a bankrupt creditor bank i, the recovered interbank lending R_i^{IBL} is given by Equation (29). The resulting default loss is

$$\mathcal{L}_i^{\mathrm{IBL}} = L_i^{\mathrm{IBL}} - R_i^{\mathrm{IBL}}. \tag{38}$$

Simultaneously, bank i also acts as a debtor in the interbank network. As a bankrupt debtor bank, bank i uses the remaining liquidation funds after repaying deposit liabilities, denoted by F_i^{br} (see Equation (32)), to repay its interbank borrowing. The default loss imposed on creditor banks is given by the repayment shortfall:

$$\mathcal{L}_i^{\mathrm{IBB}} = \max\left(L_i^{\mathrm{IBB}} - F_i^{br}, 0\right). \tag{39}$$

The total default loss in the interbank network during iteration t is

$$\mathcal{L}_{bb}(t) = \sum_{i \in \mathcal{B}_t} \left(\mathcal{L}_i^{\mathrm{IBL}} + \mathcal{L}_i^{\mathrm{IBB}}\right). \tag{40}$$

To quantify loss severity in the interbank layer, we introduce the cumulative default loss ratio ℓ_{bb}. Since default losses in this layer arise directly from interbank lending, this ratio normalizes the total default losses during the contagion process by the initial total amount of such lending:

$$\ell_{bb} = \frac{\sum_{t=1}^{T} \mathcal{L}_{bb}(t)}{\sum_{i=1}^{N_b} L_i^{\mathrm{IBL}}(0)}, \tag{41}$$

where $L_i^{\mathrm{IBL}}(0)$ represents the initial value of bank i's interbank lending.

In the interfirm network, default losses occur when a bankrupt creditor firm fails to fully recover its accounts receivable, or when a bankrupt debtor firm cannot fully settle its accounts payable. For a bankrupt creditor firm j, the recovered accounts receivable R_j^{AR} is given by Equation (24). The associated default loss is

$$\mathcal{L}_j^{AR} = L_j^{AR} - R_j^{AR}. \tag{42}$$

Likewise, firm j also serves as a debtor in the interfirm network. As a bankrupt debtor firm, firm j uses the remaining liquidation funds after repaying bank loans, denoted by F_j^{fr} (see

Equation (27)), to settle its accounts payable. The default loss imposed on creditor firms is given by the repayment shortfall:

$$\mathcal{L}_j^{AP} = \max\left(L_j^{AP} - F_j^{fr}, 0\right). \tag{43}$$

The total default loss in the interfirm network during iteration t is

$$\mathcal{L}_{ff}(t) = \sum_{j \in \mathcal{F}_t} \left(\mathcal{L}_j^{AR} + \mathcal{L}_j^{AP}\right). \tag{44}$$

To quantify loss severity in the interfirm layer, we introduce the cumulative default loss ratio ℓ_{ff}. Since default losses in this layer arise directly from accounts receivable, this ratio normalizes the total default losses during the contagion process by the initial total amount of accounts receivable:

$$\ell_{ff} = \frac{\sum_{t=1}^{T} \mathcal{L}_{ff}(t)}{\sum_{j=1}^{N_f} L_j^{AR}(0)}, \tag{45}$$

where $L_j^{AR}(0)$ denotes the initial accounts receivable of firm j. These layer-specific loss ratios provide ex-post measures of default-loss severity in each network layer. The interfirm loss ratio ℓ_{ff} and the interbank loss ratio ℓ_{bb} characterize losses arising within the interfirm and interbank credit layers, respectively, whereas the bank–firm loss ratio ℓ_{bf} characterizes losses arising through bank–firm credit relationships. Because these ratios are recorded for every realization, they can be evaluated conditional on the realized system state rather than only through unconditional averages; this distinction is developed in Section 2.6.1.

2.6. Characterization of System-Wide Collapse and Systemic Risk Measures

In nonlinear interconnected networks, system-wide collapse is an important manifestation of systemic risk. To better understand this behavior in the multilayer financial network, we examine bankruptcy fractions, susceptibility, and finite-size behavior from complementary perspectives. We further quantify systemic risk along three dimensions—collapse likelihood, shock tolerance, and financial damage—using system-wide collapse probability, collapse thresholds, and layer-specific loss ratios, respectively. The following subsections define these quantities and explain how they are estimated.

2.6.1. Bankruptcy Fractions and Collapse Probabilities

For each simulation realization at a given shock intensity ε, we define the cumulative bankruptcy fractions of the bank and firm sectors as

$$\rho_b = \frac{n_b}{N_b}, \qquad \rho_f = \frac{n_f}{N_f}, \tag{46}$$

where n_b and n_f denote the numbers of banks and firms, respectively, that become bankrupt at any point during the contagion process. Initial exogenous shocks are applied exclusively to firms, and n_f includes firms rendered bankrupt directly by these shocks. Banks can become bankrupt only through subsequent endogenous contagion.

To distinguish sector-wide collapse from partial default contagion, we define the sectoral collapse indicators,

$$I_b = \mathbf{1}(\rho_b \geq q), \qquad I_f = \mathbf{1}(\rho_f \geq q), \tag{47}$$

where $\mathbf{1}(\cdot)$ is the indicator function. Setting q close to unity defines sector-wide collapse as the bankruptcy of nearly all entities in the corresponding sector. Unless otherwise stated, we set $q = 0.99$. System-wide collapse requires both sectors to collapse:

$$I_{\text{sys}} = I_b I_f. \tag{48}$$

The pair (I_b, I_f) identifies four mutually exclusive collapse configurations: neither sector $(0,0)$, bank-only $(1,0)$, firm-only $(0,1)$, and joint collapse $(1,1)$. Their probabilities are denoted by P_{00}, P_{10}, P_{01}, and P_{11}, respectively.

For each shock intensity, collapse probabilities are estimated from R independent realizations:

$$P_b(\varepsilon) = \frac{1}{R}\sum_{r=1}^{R} I_b^{(r)}, \qquad P_f(\varepsilon) = \frac{1}{R}\sum_{r=1}^{R} I_f^{(r)}, \tag{49}$$

and

$$P_{\text{sys}}(\varepsilon) = \frac{1}{R}\sum_{r=1}^{R} I_{\text{sys}}^{(r)} = P_{11}(\varepsilon). \tag{50}$$

The sectoral collapse probabilities satisfy $P_b = P_{10} + P_{11}$ and $P_f = P_{01} + P_{11}$. Bankruptcy fractions measure the extent of bankruptcy within individual realizations, whereas collapse probabilities measure how frequently the specified collapse criterion is met across realizations at a given shock intensity.

However, collapse probabilities do not capture the magnitude of the associated financial damage. The layer-specific default loss ratios defined in Section 2.5.3 therefore provide a complementary measure of financial severity. Because these ratios are recorded for each realization, we calculate their means conditional on whether system-wide collapse occurs, allowing loss severity to be compared directly between the two groups. Conditional means are reported only when the corresponding state is observed in the simulation sample.

2.6.2. Collapse Thresholds and Sectoral Differences

To quantify the shock intensity associated with a specified collapse probability, we define the collapse threshold ε_p^X for $X \in \{b, f, \text{sys}\}$ as

$$\varepsilon_p^X = \inf\left\{\varepsilon : P_X(\varepsilon) \geq \frac{p}{100}\right\}, \qquad p \in \{10, 50, 90\}. \tag{51}$$

The collapse threshold is estimated at the first crossing of the specified collapse-probability level, using linear interpolation between adjacent points on the shock-intensity grid. If the specified collapse-probability level is not reached within the shock-intensity range examined in the corresponding experiment, the collapse threshold is left undefined rather than extrapolated.

We designate $\varepsilon_{50}^{\text{sys}}$ as the system-wide collapse threshold. A higher $\varepsilon_{50}^{\text{sys}}$ indicates that a stronger external shock is needed for the system-wide collapse probability to reach 0.5, implying greater system-wide shock tolerance. Thus, $\varepsilon_{50}^{\text{sys}}$ serves as our primary measure of system-wide shock tolerance. We also report $\varepsilon_{10}^{\text{sys}}$ and $\varepsilon_{90}^{\text{sys}}$. The transition width of the system-wide collapse probability is defined as the shock-intensity interval between its first crossings of 0.1 and 0.9:

$$W_{10\text{–}90} = \varepsilon_{90}^{\text{sys}} - \varepsilon_{10}^{\text{sys}}. \tag{52}$$

A smaller $W_{10\text{–}90}$ indicates a sharper transition, with the system-wide collapse probability rising from 0.1 to 0.9 over a narrower range of shock intensity.

To compare the shock tolerance of the bank and firm sectors, we define the difference between their collapse thresholds as

$$S_{\text{sector}} = \varepsilon_{50}^{b} - \varepsilon_{50}^{f}. \tag{53}$$

A positive value indicates that the banking sector reaches a collapse probability of 0.5 at a higher shock intensity than the firm sector, while a negative value indicates the reverse.

2.6.3. Susceptibility and Finite-Size Analysis

As the number of financial entities increases, fluctuations in bankruptcy fractions tend to decrease when bankruptcies are independent, because individual-level randomness is averaged out. Correlations among bankruptcies, however, may weaken this averaging effect. To quantify fluctuations in bankruptcy fractions at a given shock intensity, we define the susceptibilities of the bank and firm sectors as

$$\chi_b = N_b\Big(\langle \rho_b^2 \rangle - \langle \rho_b \rangle^2\Big), \qquad \chi_f = N_f\Big(\langle \rho_f^2 \rangle - \langle \rho_f \rangle^2\Big), \tag{54}$$

where $\langle \cdot \rangle$ denotes the ensemble average. Under an independent-bankruptcy benchmark, the variance of the bankruptcy fraction typically decreases in proportion to $1/N_x$, leaving χ_x approximately constant as the sector grows. Examining how susceptibility changes with system size therefore helps assess whether, in larger systems, fluctuations in bankruptcy fractions diminish as expected under this benchmark or remain substantial.

For each sector $x \in \{b, f\}$, we identify the susceptibility peak $\chi_x^{\max}$ as the maximum observed on the shock-intensity grid and its location $\varepsilon_{\chi,x}$ as the corresponding shock intensity. To characterize the width of the susceptibility peak, we use the full width at half maximum (FWHM), defined as the difference between the shock intensities at which susceptibility reaches half its peak value, $\chi_x^{\max}/2$, on either side of the peak. These crossing points are estimated by linear interpolation between adjacent shock-intensity grid points whose susceptibility values bracket the half-maximum level. A smaller FWHM indicates that the peak is confined to a narrower shock-intensity interval and is therefore sharper. If both crossings are not observed within the examined range, the FWHM is left undefined rather than extrapolated. We compare the peak heights, locations, and FWHMs across system sizes for the bank and firm sectors separately in Section 3.2.

3. Numerical Results and Criticality Analysis

This section presents the numerical analysis of the multilayer financial network under exogenous shocks. We begin with the baseline configuration, characterizing its network structure and the resulting bankruptcy fractions, susceptibility, collapse probabilities, and layer-specific financial losses. We then examine how these features change with system size through finite-size analysis. Next, matched two- and three-layer experiments assess the contribution of endogenous interfirm credit propagation. We subsequently analyze how selected balance-sheet and interfirm-network parameters affect shock tolerance and sectoral responses, followed by robustness and numerical-sensitivity analyses.

3.1. Baseline Network and Shock Response

3.1.1. Baseline Parameter Configuration and Network Characteristics

The baseline system consists of 200 banks and 3000 firms connected through three directed credit layers: interbank lending, bank–firm lending, and interfirm trade credit. The procedures for network generation and balance-sheet construction are described in Sections 2.1–2.4. The baseline parameter configuration is summarized in Table 3.

The baseline values are used as stylized reference settings rather than as empirically calibrated estimates.

Table 3. Baseline parameter configuration of the multilayer financial network.

Parameter	Description	Value
N_b	Number of banks	200
N_f	Number of firms	3000
α	Pareto exponent	1.2
A_b^{LB}	Lower bound of bank total assets	100
A_b^{UB}	Upper bound of bank total assets	2000
λ_f	Average degree of firms in the bank–firm network	2
α_1	Parameter governing link formation (common to interbank and interfirm networks)	0.25
α_2	Parameter governing link formation (common to interbank and interfirm networks)	1
d_b	Parameter governing interbank link formation	0.5
d_f	Parameter governing interfirm link formation	0.5
ω_b	Share of equity in total assets (banks)	0.1
θ_{bf}	Share of loans to firms in total assets (banks)	0.5
θ_{bb}	Share of interbank lending in total assets (banks)	0.2
μ_{fb}	Share of bank loans in total assets (firms)	0.3
μ_{ff}	Share of accounts payable in total assets (firms)	0.4
μ_c	Share of cash in total assets (firms)	0.1
δ_b	Bank liquidation recovery rate	0.7
δ_f	Firm liquidation recovery rate	0.5

To characterize the topology generated by the baseline configuration, we analyze 150 independently generated pre-shock networks, separate from the realizations used in the contagion experiments. All links are directed from creditors to debtors. In the bank–firm layer, a bank's out-degree is the number of firms it lends to, whereas a firm's in-degree is the number of banks it borrows from. Table 4 reports the resulting network statistics.

Table 4. Pre-shock characteristics of the baseline networks. Values are means across 150 independently generated networks.

Layer	Mean Out-Degree	Mean in-Degree	Density	Out–in Degree Correlation	Out-Degree–Asset Correlation	In-Degree–Asset Correlation
BF	33.815	2.254	0.01127	−0.0205	0.984	0.993
BB	9.512	9.512	0.04780	−0.0292	0.480	0.953
FF	94.092	94.092	0.03137	−0.0031	0.835	0.995

Notes: BF denotes the bank–firm layer, BB the interbank layer, and FF the interfirm layer. In the BF layer, out-degree refers to banks and in-degree to firms. The out–in degree correlation is the Pearson correlation between creditors' out-degree and debtors' in-degree across directed links within each layer. Degree–asset correlations are node-level Pearson correlations between the corresponding degree and pre-shock total assets.

In the bank–firm layer, the realized mean out-degree of banks is 33.815, and the mean in-degree of firms is 2.254, compared with the initial targets of 30 and 2, respectively. The higher realized values reflect the addition of links to ensure that every bank and firm has at least one link. Degree–asset correlations are positive in all three layers, consistent with the asset-dependent network construction. In the interbank and interfirm layers, these correlations are stronger for in-degree than for out-degree. The out–in degree correlations are close to zero in all three layers, indicating that pronounced disassortative mixing is not reproduced under the baseline configuration. The degree distributions differ both across layers and between out-degree and in-degree within each layer. Figure S1 in Section S2 of the Supplementary Materials presents their complementary cumulative distributions. These descriptive diagnostics do not establish a power-law form for the generated degree distributions. Accordingly, the networks are treated as stylized financial configurations rather than empirically calibrated replicas.

For the baseline shock experiment, the prescribed shock intensity ranges from $\varepsilon = 0.120$ to 0.220 in increments of 0.001, with 500 independent realizations at each grid point. A new network and its associated balance sheets are generated for each realization, together with a new selection of shocked firms.

3.1.2. Bimodal Bankruptcy Fractions and Susceptibility

We first examine the ensemble-averaged cumulative bankruptcy fractions of banks and firms as functions of shock intensity ε. As shown in Figure 1a, both sectors exhibit a pronounced nonlinear response. At $\varepsilon = 0.120$, the mean bankruptcy fractions are approximately 0.00025 for banks and 0.157 for firms. Both increase rapidly over a relatively narrow shock-intensity range and reach unity at $\varepsilon = 0.160$. However, the ensemble averages alone do not reveal how bankruptcy fractions are distributed across individual realizations.

To determine where bankruptcy fractions exhibit the strongest fluctuations, we examine susceptibility. Since $\chi_x = N_x \operatorname{Var}(\rho_x)$ for $x \in \{b, f\}$, susceptibility provides a size-scaled measure of the variance across realizations. As shown in Figure 1b, the bank and firm susceptibilities both reach their observed maxima at $\varepsilon = 0.144$, with $\chi_b^{\max} = 49.935$ and $\chi_f^{\max} = 448.784$. We therefore examine the realization-level distributions at this shock intensity.

Figure 1c,d reveal two distinct groups of realizations. In 255 of the 500 realizations, both sectors experience complete collapse, with bankruptcy fractions equal to unity. The remaining 245 realizations exhibit limited contagion, with mean bankruptcy fractions of 0.00045 for banks and 0.227 for firms. The ensemble means at this shock intensity therefore result from averaging over these two distinct groups rather than from realizations concentrated at intermediate bankruptcy fractions.

The bimodal distributions raise a further question: whether the large fluctuations identified by susceptibility arise mainly from variation within each group or from the separation between them. To quantify these two sources of variation, let H denote the system-wide collapse state and L its complement. For a given shock intensity, we define p_H as the probability of H and $\mu_{H,x}$ and $\mu_{L,x}$ as the conditional mean bankruptcy fractions of sector $x \in \{b, f\}$ in states H and L, respectively. The variance of the bankruptcy fraction can be decomposed as

$$\operatorname{Var}(\rho_x) = p_H \operatorname{Var}(\rho_x \mid H) + (1 - p_H) \operatorname{Var}(\rho_x \mid L) + p_H(1 - p_H)(\mu_{H,x} - \mu_{L,x})^2. \quad (55)$$

The first two terms represent the within-state contribution, while the final term represents the between-state contribution. At $\varepsilon = 0.144$, the between-state contribution accounts for 99.9996% of the variance in bank bankruptcy fractions and 99.8526% of the variance in firm bankruptcy fractions. Thus, the large fluctuations in bankruptcy fractions at this shock intensity arise almost entirely from the separation between the two states rather than from substantial variation within either state.

This decomposition further clarifies the origin of the susceptibility peaks. For two well-separated states, the between-state contribution contains the factor $p_H(1 - p_H)$, which is maximized when the two states are equally probable. At $\varepsilon = 0.144$, the estimated probability of system-wide collapse is $p_H = 0.510$, placing this factor very close to its theoretical maximum. The observed susceptibility maxima therefore occur at a shock intensity where the limited-contagion and system-wide-collapse states coexist with comparable frequencies and remain strongly separated.

Overall, the analysis shows that at the observed susceptibility peak, the intermediate ensemble averages conceal a pronounced two-state structure consisting of limited contagion and system-wide collapse. Having identified this structure, Section 3.1.3 next examines

how the frequency of system-wide collapse changes with shock intensity, together with the associated shock-tolerance measures and financial losses.

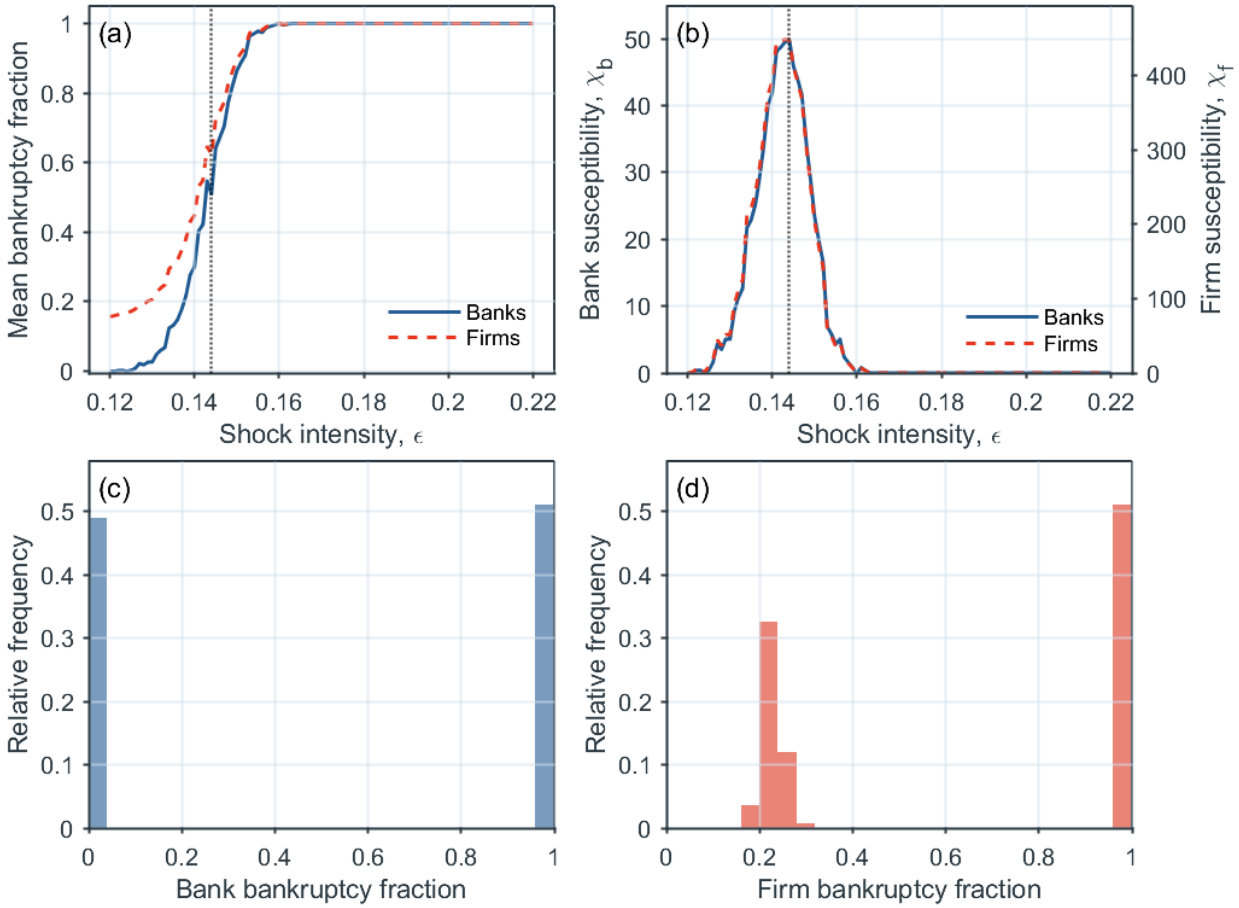


Figure 1. Baseline bankruptcy fractions and susceptibility. (**a**) Ensemble-averaged bankruptcy fractions of banks and firms as functions of shock intensity. (**b**) Bank and firm susceptibilities, shown against the left and right vertical axes, respectively. The dotted vertical lines in panels (**a**,**b**) mark $\varepsilon = 0.144$, where both susceptibilities reach their maxima on the simulated shock-intensity grid and at which the distributions in panels (**c**,**d**) are evaluated. (**c**) Distribution of bank bankruptcy fractions at $\varepsilon = 0.144$. (**d**) Distribution of firm bankruptcy fractions at the same shock intensity.

3.1.3. Collapse Likelihood, Shock Tolerance, and Financial Damage

The two-state pattern identified in Section 3.1.2 motivates a quantitative characterization of the baseline system along three complementary dimensions: collapse likelihood, shock tolerance, and financial damage. Collapse likelihood is measured by the estimated collapse probabilities, shock tolerance by the collapse thresholds, and financial damage by the layer-specific default loss ratios. We first examine how the probability of system-wide collapse changes with shock intensity. As defined in Section 2.6, system-wide collapse requires the bankruptcy fractions of both banks and firms to reach or exceed $q = 0.99$. Figure 2a shows that the estimated probability of system-wide collapse, $P_{\mathrm{sys}}(\varepsilon)$, increases rapidly over a narrow shock-intensity range. It rises from 0.300 at $\varepsilon = 0.140$ to 0.864 at $\varepsilon = 0.150$. At $\varepsilon = 0.160$, all 500 realizations satisfy the system-wide collapse criterion, giving an estimated collapse probability of unity. The bank, firm, and system-wide collapse probability curves nearly coincide. Across all 50,500 baseline realizations, only two exhibit sectorally asymmetric collapse: firms meet the collapse criterion while banks do not. No bank-only collapse is observed. Thus, in almost every baseline realization, either both sectors meet the collapse criterion or neither does.

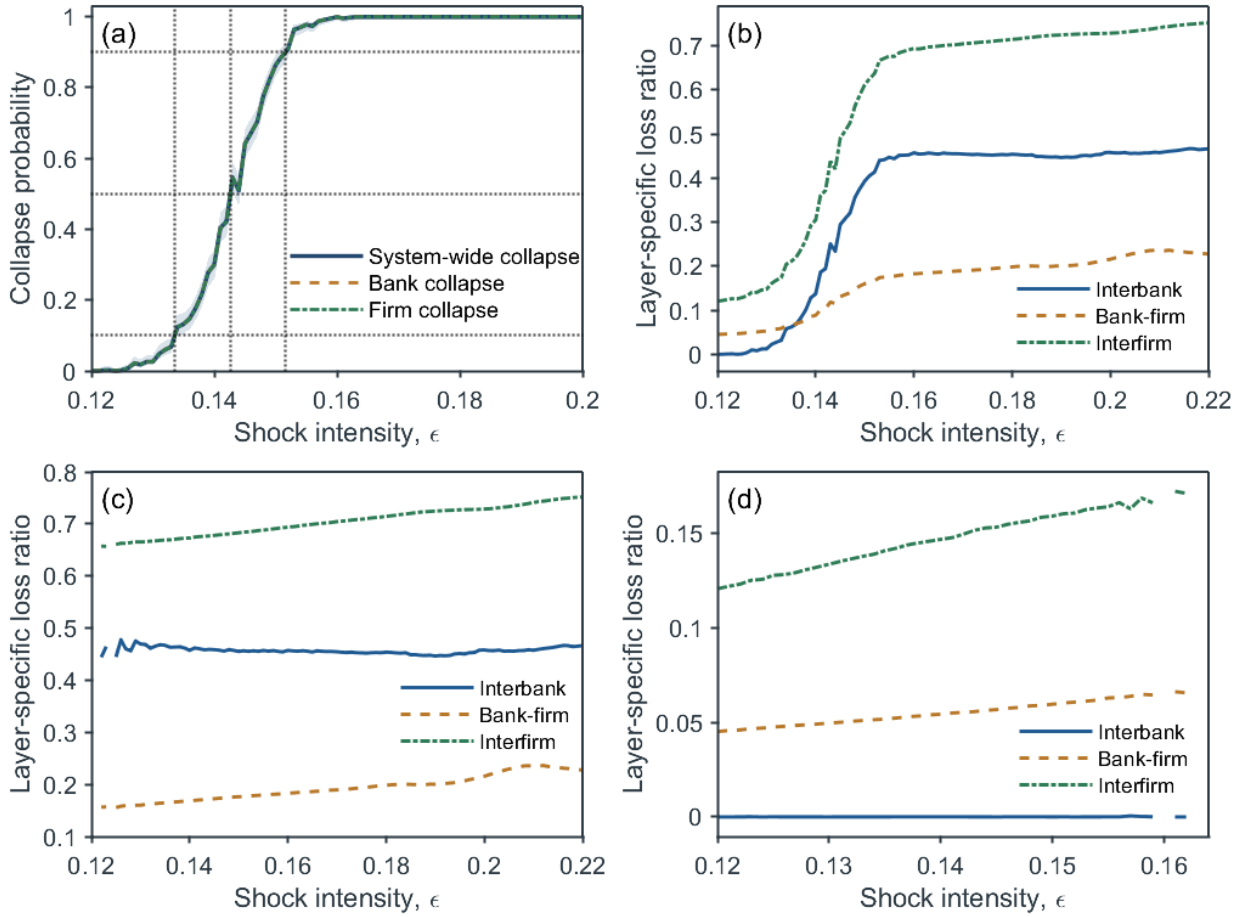


Figure 2. Baseline collapse probabilities and default losses. (**a**) Estimated bank, firm, and system-wide collapse probabilities as functions of shock intensity. The shaded region represents pointwise 95% Wilson confidence intervals for the system-wide collapse probability. The horizontal reference lines mark probabilities of 0.1, 0.5, and 0.9, and the vertical lines indicate the system-wide collapse thresholds $\epsilon_{10}^{\mathrm{sys}}$, $\epsilon_{50}^{\mathrm{sys}}$, and $\epsilon_{90}^{\mathrm{sys}}$. (**b**) Unconditional mean default loss ratios in the interbank (BB), bank–firm (BF), and interfirm (FF) layers. (**c**,**d**) Mean default loss ratios conditional on system-wide collapse and on the absence of system-wide collapse, respectively. Conditional means are shown only at shock intensities where the corresponding state is observed among the 500 realizations.

Table 5 reports the estimated collapse thresholds and their bootstrap confidence intervals. These thresholds quantify shock tolerance by indicating the shock intensity required for the collapse probability to reach a specified level; higher thresholds therefore correspond to greater shock tolerance. The system-wide collapse threshold is $\epsilon_{50}^{\mathrm{sys}} = 0.142629$, with a 95% confidence interval of $[0.142358, 0.142956]$. The collapse thresholds $\epsilon_{10}^{\mathrm{sys}}$ and $\epsilon_{90}^{\mathrm{sys}}$ are 0.133571 and 0.151600, respectively, giving $W_{10\text{–}90} = 0.018029$. Thus, the baseline system reaches a 50% probability of system-wide collapse at a shock intensity of approximately 0.143, while the transition from a 10% to a 90% system-wide collapse probability occurs over a shock-intensity interval of about 0.018. The estimated collapse thresholds for banks, firms, and the system as a whole coincide at the probability levels of 0.1, 0.5, and 0.9, giving $S_{\mathrm{sector}} = 0$. This coincidence indicates no sectoral difference in shock tolerance as measured by the collapse thresholds under the baseline configuration.

Collapse probability characterizes collapse likelihood but does not quantify the severity of the associated financial damage. We therefore use the layer-specific default loss ratios to measure financial damage across the three credit layers. Figure 2b reports the unconditional mean default loss ratios in the three credit layers, while panels (c) and (d) show the

mean loss ratios conditional on whether system-wide collapse occurs. At $\varepsilon = 0.144$, the mean interbank loss ratio is 0.4593 among system-wide-collapse realizations, compared with approximately 0.00003 among realizations without system-wide collapse. The corresponding ratios are 0.1717 and 0.0564 in the bank–firm layer and 0.6770 and 0.1529 in the interfirm layer. The near absence of interbank losses in realizations without system-wide collapse is consistent with the very low bank bankruptcy fraction reported in Section 3.1.2. However, bank–firm and interfirm losses remain present even when system-wide collapse does not occur, indicating that financial damage is not confined to realizations that meet the system-wide collapse criterion.

Table 5. Baseline system-wide collapse thresholds and transition width, with 95% bootstrap confidence intervals. Confidence intervals are obtained from 1000 nonparametric bootstrap resamples.

Measure	Estimate	95% CI
$\varepsilon_{10}^{\mathrm{sys}}$	0.133571	[0.133222, 0.134125]
$\varepsilon_{50}^{\mathrm{sys}}$	0.142629	[0.142358, 0.142956]
$\varepsilon_{90}^{\mathrm{sys}}$	0.151600	[0.150611, 0.152211]
$W_{10\text{–}90}$	0.018029	[0.016809, 0.018767]

Loss severity can also differ between shock intensities with the same estimated system-wide collapse probability. At both $\varepsilon = 0.160$ and 0.220, all 500 realizations satisfy the system-wide collapse criterion. Nevertheless, the mean bank–firm loss ratio increases from 0.1829 to 0.2275 between these two shock intensities, while the mean interfirm loss ratio increases from 0.6938 to 0.7520. The mean interbank loss ratio changes only slightly over the same interval. Thus, the estimated system-wide collapse probability alone does not capture differences in financial damage once it has saturated at unity.

Together, these results show that collapse likelihood, shock tolerance, and financial damage capture distinct aspects of the baseline system's response. We next examine how collapse behavior and susceptibility vary with system size.

3.2. Finite-Size Dependence of Collapse Behavior

The baseline analysis reveals a pronounced two-state structure, a rapid change in system-wide collapse probability, and sharp susceptibility peaks. We now examine how these features change with system size. Five sizes are considered, with $N_b \in \{150, 200, 250, 300, 400\}$ and $N_f = 15N_b$. To reduce the confounding effect of changes in network connectivity, the interbank and interfirm link-formation parameters, d_b and d_f, are calibrated separately at each size to maintain approximately constant realized mean degrees. A pilot scan with 20 realizations per shock intensity is used to identify the relevant shock range for each size. The subsequent formal experiments use 500 independent realizations per shock intensity, with a grid spacing of 0.001. The formal shock ranges, ordered by increasing N_b, are $[0.117, 0.168]$, $[0.116, 0.167]$, $[0.115, 0.165]$, $[0.116, 0.166]$, and $[0.117, 0.167]$, respectively.

Table 6 summarizes the calibrated network parameters and finite-size results. The target mean degrees are 9.339 for the interbank layer and 94.136 for the interfirm layer. Across the five sizes, the measured mean degrees range from 9.336 to 9.437 and from 93.563 to 95.042, respectively. Their maximum absolute deviations from the targets are approximately 1.05% and 0.96%. The calibration therefore substantially reduces variation in mean degree across system sizes, although the realized connectivity is not identical.

Table 6. Finite-size results under the degree-controlled protocol. The number of firms is set to $N_f = 15N_b$. $\langle k_{bb} \rangle$ and $\langle k_{ff} \rangle$ denote the realized mean degrees of the interbank and interfirm layers, respectively, measured from pre-shock network realizations. Susceptibility peaks are the maxima observed on the simulated shock grid, and the system-wide collapse threshold is estimated by linear interpolation. The $N_b = 200$ results are from an independent finite-size experiment and are not substituted for the baseline results in Section 3.1.

N_b	d_b	d_f	$\langle k_{bb} \rangle$	$\langle k_{ff} \rangle$	$\varepsilon_{50}^{\mathrm{sys}}$	$\chi_b^{\max}$	$\chi_f^{\max}$	W_{10-90}
150	0.6315	0.6303	9.354	93.739	0.144034	37.464	341.317	0.019096
200	0.5000	0.5000	9.437	94.681	0.142895	49.923	452.196	0.016640
250	0.4104	0.4108	9.378	93.563	0.142273	62.292	565.469	0.014673
300	0.3464	0.3551	9.336	95.042	0.142250	74.332	675.893	0.015590
400	0.2671	0.2758	9.348	93.570	0.141560	99.681	905.816	0.014455

Figure 3a shows that the estimated system-wide collapse probabilities rise rapidly with shock intensity at all five system sizes. The system-wide collapse threshold $\varepsilon_{50}^{\mathrm{sys}}$ changes only modestly, decreasing from 0.144034 at $N_b = 150$ to 0.141560 at $N_b = 400$. By contrast, the transition width W_{10-90} decreases overall from 0.019096 to 0.014455, as shown in Figure 3b. The decrease is not strictly monotonic: W_{10-90} increases from 0.014673 at $N_b = 250$ to 0.015590 at $N_b = 300$. Thus, the system-wide collapse-probability transition shifts only slightly in shock intensity but generally occurs over a narrower shock-intensity range as system size increases.

The susceptibility peak heights exhibit a stronger dependence on system size. As shown in Figure 3c, the banking-sector peak increases from 37.464 at $N_b = 150$ to 99.681 at $N_b = 400$, while the firm-sector peak increases from 341.317 at $N_f = 2250$ to 905.816 at $N_f = 6000$. To quantify this dependence, we fit the empirical relation

$$\chi_s^{\max}(N_s) \propto N_s^{\kappa_s}, \qquad s \in \{b, f\}. \tag{56}$$

Log–log regressions of the observed grid-based peak heights yield $\kappa_b = 0.9954$ and $\kappa_f = 0.9945$, with coefficients of determination exceeding 0.9999. Their respective 95% bootstrap confidence intervals are $[0.9910, 1.0064]$ and $[0.9818, 1.0070]$. Thus, within the five sizes examined, the susceptibility peaks grow approximately linearly with the size of the corresponding sector.

The susceptibility peaks also become narrower as system size increases. Figure 3d shows that the banking-sector FWHM decreases from 0.01553 at $N_b = 150$ to 0.01092 at $N_b = 400$, while the firm-sector FWHM decreases from 0.01556 to 0.01098. The shock intensities at which the susceptibility peaks occur also change only slightly, remaining within $\varepsilon \in [0.142, 0.144]$ for banks and $\varepsilon \in [0.141, 0.144]$ for firms. Hence, increasing system size is associated primarily with higher and narrower susceptibility peaks rather than a substantial shift in their positions.

The two-state structure identified in Section 3.1.2 also persists across the finite-size experiments. At the sector-specific susceptibility peaks, the between-state component of the variance accounts for more than 99.999% of banking-sector variance and approximately 99.86–99.89% of firm-sector variance. Using the two-state notation introduced in Section 3.1.2, the between-state contribution to susceptibility is

$$\chi_{s,\mathrm{between}} = N_s p_H (1 - p_H)(\mu_{H,s} - \mu_{L,s})^2. \tag{57}$$

This expression shows that susceptibility can grow approximately proportionally to N_s when the two states remain well separated and neither state becomes rare near the peak. The dominance of the between-state variance component in the simulations therefore provides an explanation for the observed near-linear growth of the susceptibility peaks.

Overall, as system size increases, the location of the system-wide collapse-probability transition changes only modestly, its width generally narrows, and the susceptibility peaks become higher and narrower. Together with the persistence of the two-state structure, these results provide consistent finite-size evidence of phase-transition-like collapse behavior. The following section uses matched two- and three-layer models to examine the contribution of endogenous interfirm credit propagation to collapse behavior.

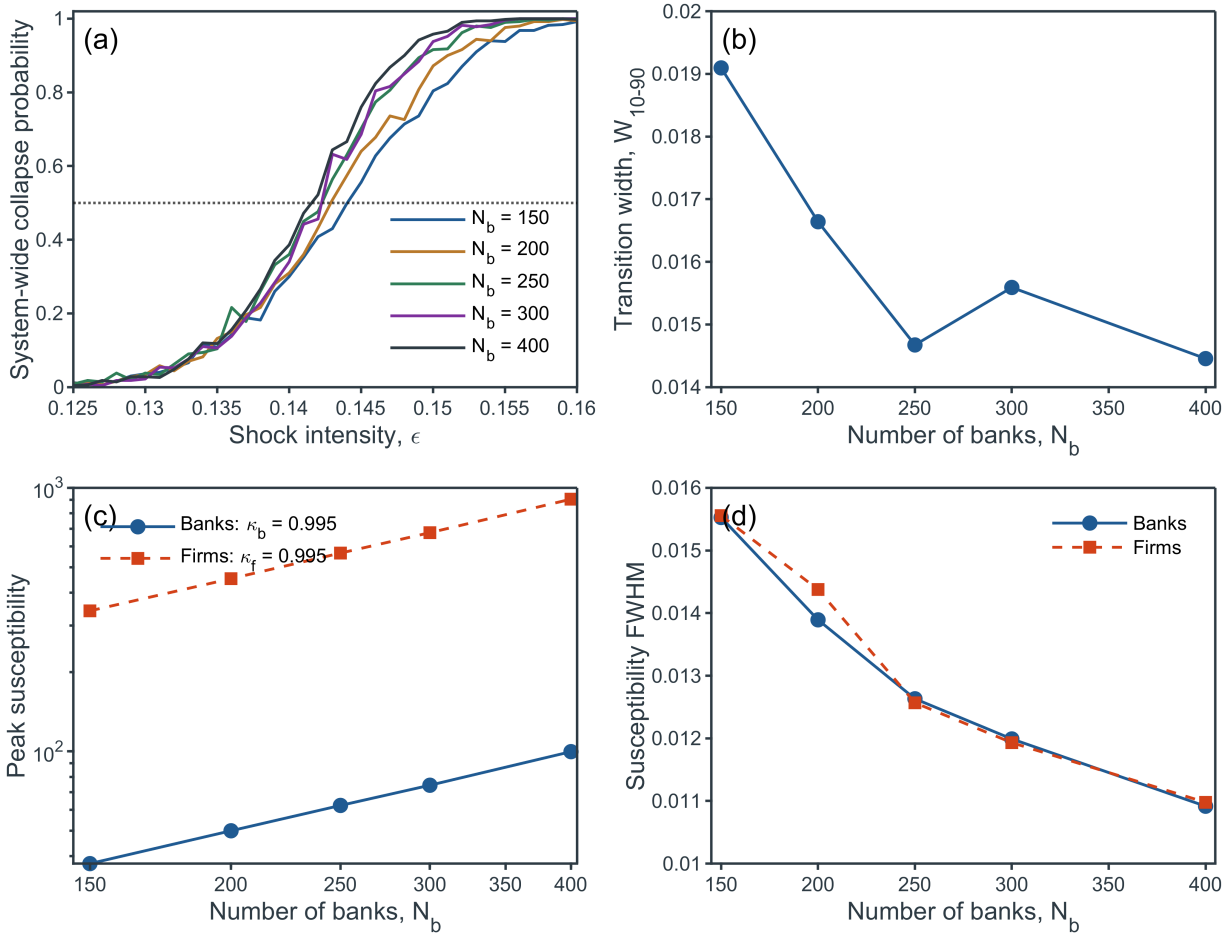


Figure 3. Degree-controlled finite-size analysis of the three-layer system, with $N_f = 15N_b$. (**a**) Estimated system-wide collapse probabilities for five system sizes. (**b**) Transition width $W_{10–90}$, defined as the shock-intensity interval between the first crossings of 0.1 and 0.9 by the system-wide collapse probability. (**c**) Observed susceptibility peak heights as functions of N_b on logarithmic axes; because $N_f = 15N_b$, this constant rescaling does not change the fitted firm-sector exponent. The legend reports the exponents κ_b and κ_f obtained from log–log regression. (**d**) Full widths at half maximum (FWHM) of the bank and firm susceptibility peaks. Susceptibility peaks are identified on the observed shock grid, while $W_{10–90}$ and the FWHM values are calculated using linear interpolation where required. All results are obtained from the independent degree-controlled finite-size experiments.

3.3. Contribution of Endogenous Interfirm Credit Propagation: Matched Model Comparisons

The preceding analysis characterizes collapse behavior in the three-layer system but does not isolate the contribution of endogenous interfirm credit propagation. To examine this contribution, we construct two-layer counterfactuals that retain the interbank and bank–firm networks while replacing endogenous interfirm settlements with external recovery arrangements. Accounts receivable and accounts payable arising from interfirm trade credit remain on firms' initial balance sheets, but their settlement no longer transmits payments or losses between firms. When a firm goes bankrupt, its accounts receivable are recovered

externally rather than through repayments from other firms. Within each simulation pair, the three-layer and two-layer models share the same initial network realization, balance sheets, shocked firms, and realized shock intensity.

Removing endogenous interfirm settlement also changes the recovery of accounts receivable during firm liquidation. To reduce the confounding effect of this change, we consider two recovery-matching methods. Under both methods, a bankrupt firm externally recovers a fixed fraction of the face value of its accounts receivable. This external recovery rate is common to all firms and realizations at a given shock intensity but may vary across shock intensities. The two methods differ in the aggregate recovery quantity they seek to reproduce.

Under recovery-rate matching, the external recovery rate is chosen to match the aggregate recovery rate of the three-layer model. At each shock intensity ε, we calculate the total recovered amount of accounts receivable, $R_{3L}(\varepsilon)$, and their total face value, $F_{3L}(\varepsilon)$, across the 500 three-layer realizations. The external recovery rate is then

$$r_{\text{rate}}(\varepsilon) = \frac{R_{3L}(\varepsilon)}{F_{3L}(\varepsilon)}. \tag{58}$$

The resulting rate is applied in the two-layer model.

Recovery-amount matching instead targets the average recovered amount per bankrupt firm in the three-layer model. At each shock intensity, the target is

$$A_{3L}(\varepsilon) = \frac{R_{3L}(\varepsilon)}{N_{3L}(\varepsilon)}, \tag{59}$$

where $N_{3L}(\varepsilon)$ is the total number of bankrupt firms across the 500 three-layer realizations. We then use the recovery-rate-matched two-layer model to calculate the average face value of accounts receivable per bankrupt firm:

$$B_{2L,\text{rate}}(\varepsilon) = \frac{F_{2L,\text{rate}}(\varepsilon)}{N_{2L,\text{rate}}(\varepsilon)}, \tag{60}$$

where $F_{2L,\text{rate}}(\varepsilon)$ is the total face value of accounts receivable held by bankrupt firms, and $N_{2L,\text{rate}}(\varepsilon)$ is the total number of bankrupt firms, both aggregated across the 500 recovery-rate-matched two-layer realizations. The external recovery rate is then set to

$$r_{\text{amount}}(\varepsilon) = \min\left\{1, \frac{A_{3L}(\varepsilon)}{B_{2L,\text{rate}}(\varepsilon)}\right\}. \tag{61}$$

The two-layer model is subsequently rerun using this rate and the same firm-level liquidation rule as under recovery-rate matching. Because changing the recovery rate can alter subsequent bankruptcies, the achieved average recovered amount need not exactly equal the three-layer target.

Figure 4 compares the resulting system-wide collapse probabilities. At $\varepsilon = 0.144$, the estimated system-wide collapse probabilities are 0.510 in the three-layer model, 0 under recovery-rate matching, and 0.076 under recovery-amount matching, corresponding to 255, 0, and 38 system-wide collapse events, respectively, out of 500 realizations per model. The recovery-rate-matched curve is shifted farther toward higher shock intensities than the recovery-amount-matched curve.

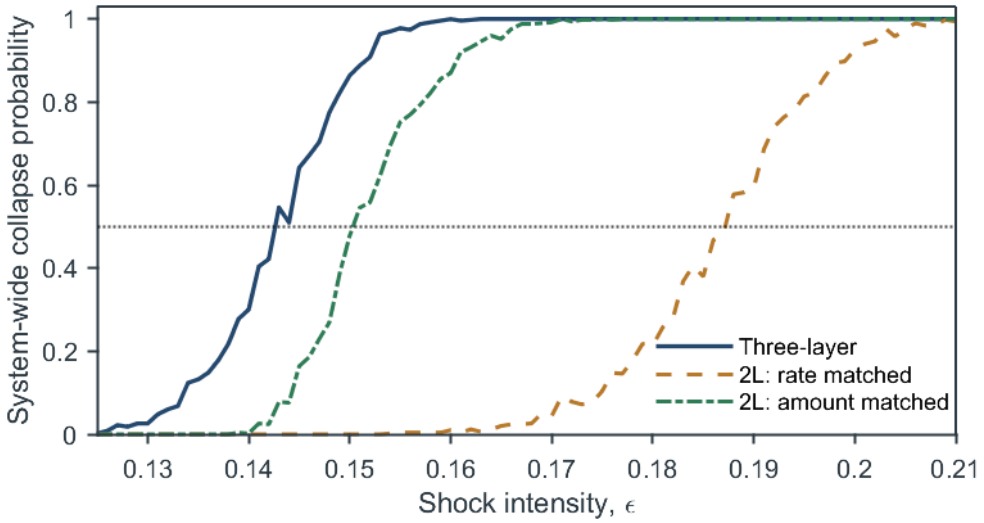


Figure 4. Matched comparisons between the three-layer model and two two-layer counterfactuals. The curves show estimated system-wide collapse probabilities for the three-layer model and the two two-layer counterfactuals under recovery-rate and recovery-amount matching, using paired initial conditions and shocks. The black dotted horizontal line indicates a system-wide collapse probability of 0.5; the corresponding system-wide collapse thresholds are reported in Table 7.

Table 7 reports the system-wide collapse thresholds and transition widths for the three models. We define the collapse-threshold difference for each matching method as

$$\Delta\varepsilon_{50}^{m} = \varepsilon_{50}^{2L,m} - \varepsilon_{50}^{3L}, \qquad m \in \{\text{rate}, \text{amount}\}, \tag{62}$$

where the thresholds refer to system-wide collapse. The system-wide collapse threshold is 0.142629 in the three-layer model, compared with 0.187170 under recovery-rate matching and 0.150324 under recovery-amount matching. The corresponding threshold differences are 0.044541 and 0.007694, with 95% paired-bootstrap confidence intervals of $[0.043204, 0.044975]$ and $[0.006999, 0.008248]$, respectively. Thus, both matching methods yield higher system-wide collapse thresholds in the two-layer counterfactuals, although the magnitude of the increase differs substantially between the two methods.

Table 7. Matched three-layer and two-layer comparisons. The system-wide collapse threshold $\varepsilon_{50}^{\text{sys}}$ and transition width $W_{10\text{–}90}$ are estimated from the observed system-wide collapse-probability curves using linear interpolation. Differences are defined as the two-layer value minus the three-layer value and are calculated before rounding.

Model	$\varepsilon_{50}^{\text{sys}}$	$W_{10\text{–}90}$	$\Delta\varepsilon_{50}$	$\Delta W_{10\text{–}90}$
Three-layer baseline	0.142629	0.018029	–	–
Two-layer, rate-matched	0.187170	0.024214	+0.044541	+0.006186
Two-layer, amount-matched	0.150324	0.016327	+0.007694	−0.001701

The two-state structure identified in the three-layer model also persists in both two-layer counterfactuals. At the sector-specific susceptibility peaks, the between-state component accounts for more than 99.6% of the variance in both sectors under either matching method. Thus, removing endogenous interfirm credit propagation does not eliminate the pronounced separation between the two states.

The difference in threshold shifts between the two matching methods reflects the fact that they preserve different aspects of recovery in the three-layer model. Recovery-rate matching uses a single unconditional aggregate recovery rate at each shock intensity and

therefore does not preserve the variation in recovery rates associated with different realized bankruptcy fractions in the three-layer model. At $\varepsilon = 0.144$, the aggregate recovery rates in the three-layer model are approximately 0.535 in state H and 0.992 in state L, whereas the unconditional rate used for matching is 0.642. Recovery-amount matching instead targets the average recovered amount per bankrupt firm. Across the 16 shock-grid points from $\varepsilon = 0.145$ to 0.160, the maximum absolute relative deviation from the three-layer recovery-amount target is 4.05%, and the median absolute relative deviation is 1.62%. These differences in what is preserved by the two matching methods are consistent with the substantially different threshold shifts observed in the two counterfactuals.

The effect on transition width differs between the two matching methods. Under recovery-rate matching, $W_{10\text{–}90}$ increases from 0.018029 in the three-layer model to 0.024214 in the two-layer model, giving a difference of 0.006186 with a 95% paired-bootstrap confidence interval of $[0.004270, 0.010962]$. Under recovery-amount matching, the transition width decreases to 0.016327, giving a difference of -0.001701 with a corresponding confidence interval of $[-0.002625, -0.000118]$. Thus, unlike the system-wide collapse threshold, the direction of the change in transition width depends on the recovery-matching method.

These comparisons characterize differences between the specified model arrangements rather than a pure topological effect of the interfirm layer. Although the initial conditions and shocks are paired, replacing endogenous interfirm settlements also changes subsequent repayment and recovery dynamics. Neither matching method reproduces the full realization-level distribution of recoveries and financial flows. The paired-bootstrap confidence intervals treat the calibrated external recovery rates as fixed and therefore do not include the additional uncertainty associated with estimating these rates.

Overall, removing endogenous interfirm credit propagation raises the system-wide collapse threshold under both recovery-matching methods, although the magnitude of the increase depends substantially on how recovery is matched. The direction of the change in transition width is also matching-dependent, whereas the two-state structure persists in both two-layer counterfactuals. The following section examines how changes in balance-sheet composition affect collapse behavior in the three-layer system.

3.4. Effects of Balance-Sheet Composition

We next examine how two aspects of balance-sheet composition affect shock tolerance: the share of equity in total assets for banks, ω_b, and the share of cash in total assets for firms, μ_c. Bank equity provides a buffer against losses, whereas firm cash supports the fulfillment of immediate repayment obligations. Figure 5 summarizes the effects of these two parameters on system-wide and sector-specific collapse thresholds.

3.4.1. Bank Capitalization

Bank capitalization is measured by ω_b, the share of equity in total assets. We vary ω_b over seven equally spaced values from 0 to 0.20, holding the remaining model parameters at their baseline settings. Under the accounting identities in Section 2.4, increasing ω_b raises bank equity while reducing residual deposit liabilities for a given asset size. The experiment therefore captures a change in both equity buffers and liability composition. The configuration $\omega_b = 0$ is a degenerate boundary case and is excluded from the following threshold comparisons.

Figure 5a shows that the system-wide collapse threshold increases with bank capitalization. It rises from 0.070386 at $\omega_b = 0.0333$ to 0.142700 at the baseline value of 0.10 and reaches 0.361278 at $\omega_b = 0.20$. The increase is uneven, with relatively little change between $\omega_b = 0.0667$ and 0.10 and larger increases at higher capitalization levels. Thus, higher bank

capitalization is associated with greater system-wide shock tolerance under the specified balance-sheet adjustments.

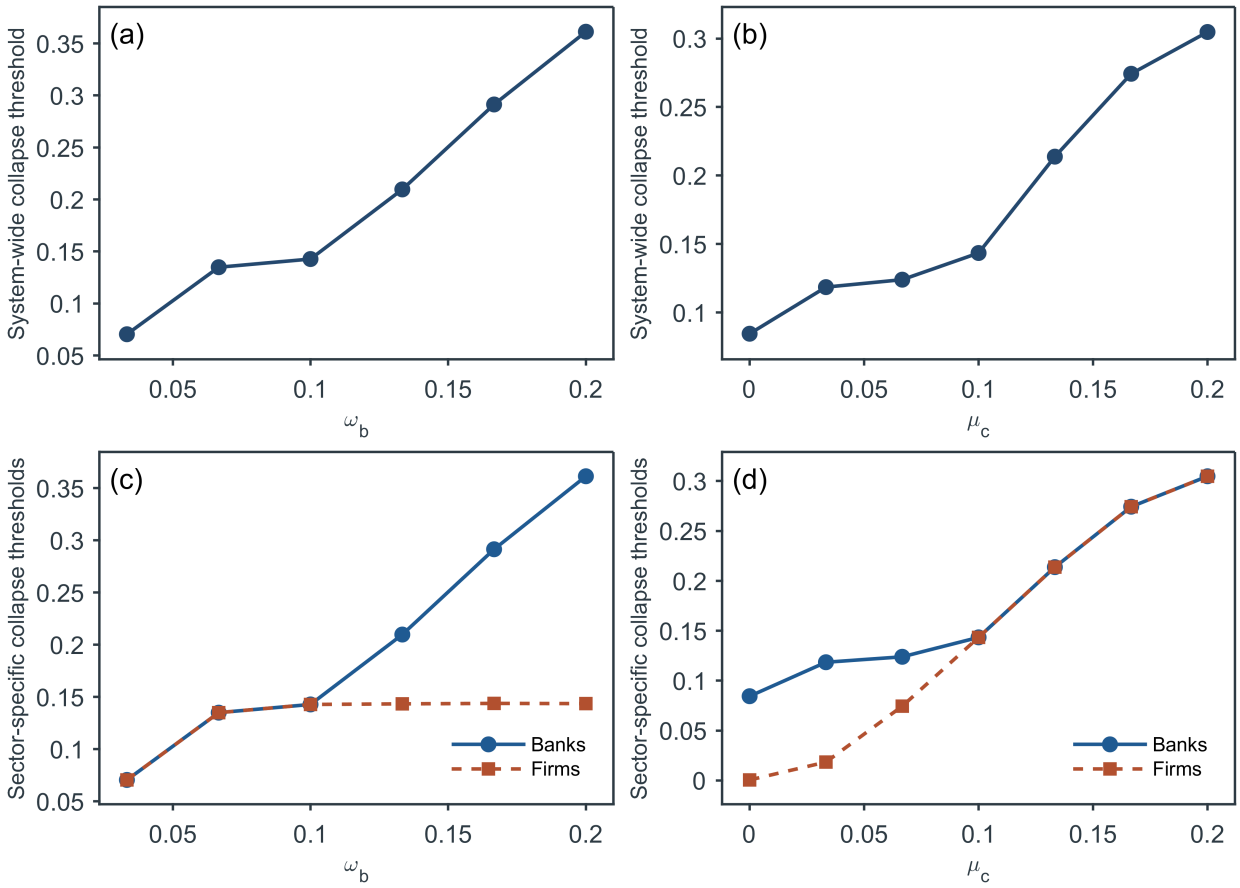


Figure 5. Effects of balance-sheet composition on collapse thresholds. (**a**,**b**) System-wide collapse thresholds under changes in the bank equity share ω_b and firm cash share μ_c, respectively. Circles show point estimates; uncapped vertical lines indicate 95% bootstrap confidence intervals. The intervals are very narrow relative to the vertical-axis scale and may be difficult to distinguish from the circles. (**c**,**d**) Bank- and firm-sector collapse thresholds under changes in ω_b and μ_c, respectively. The degenerate configuration $\omega_b = 0$ is excluded from panels (**a**,**c**) and from quantitative threshold comparisons. At $\mu_c = 0$, the firm-sector threshold is estimated by interpolation between the first two shock-grid points and is therefore treated as a boundary estimate.

The sector-specific thresholds show that this increase in system-wide shock tolerance does not reflect comparable changes in both sectors. As shown in Figure 5c, the bank and firm thresholds coincide for the nondegenerate configurations with $\omega_b \leq 0.10$. At $\omega_b = 0.1333$, however, the bank threshold rises to 0.209645 while the firm threshold remains at 0.143273, giving $S_{\text{sector}} = 0.066372$. At $\omega_b = 0.20$, the separation reaches 0.217761, with a bank threshold of 0.361278 and a firm threshold of 0.143517. The system-wide threshold coincides with the bank threshold throughout the investigated configurations. Hence, at higher capitalization levels, the increase in system-wide shock tolerance primarily reflects a higher banking-sector threshold rather than a comparable increase in firm-sector shock tolerance.

3.4.2. Firm Liquidity

Firm liquidity is measured by μ_c, the share of cash in total assets. We vary μ_c over seven equally spaced values from 0 to 0.20, holding the remaining model parameters at their baseline settings. Under the accounting identities in Section 2.4, increasing μ_c raises firms' initial cash holdings while reducing their residual fixed assets. The experiment therefore

represents a reallocation of firm assets toward cash rather than an isolated increase in liquidity with other balance-sheet items unchanged.

Figure 5b shows that the system-wide collapse threshold increases across the investigated liquidity range. It rises from 0.084400 at $\mu_c = 0$ to 0.143355 at the baseline value of 0.10 and reaches 0.304800 at $\mu_c = 0.20$. The increase is not uniform: the threshold changes only from 0.118444 to 0.123889 between $\mu_c = 0.0333$ and 0.0667, followed by larger increases over parts of the higher-liquidity range. Thus, greater initial cash holdings are associated with higher system-wide shock tolerance under the specified asset reallocation.

The sector-specific thresholds reveal a different pattern from that observed for bank capitalization. As shown in Figure 5d, at $\mu_c = 0.0333$, the firm-sector threshold is 0.018400, substantially below the banking-sector threshold of 0.118444, giving $S_{\text{sector}} = 0.100044$. The separation decreases to 0.049444 at $\mu_c = 0.0667$, and from the baseline value of $\mu_c = 0.10$ onward, the estimated thresholds of the two sectors coincide. At $\mu_c = 0$, the firm-sector threshold is a boundary estimate obtained by interpolation between the first two shock-grid points and is excluded from the quantitative comparison of sectoral threshold separation. Thus, low firm liquidity is associated with a substantially lower firm-sector threshold, and this separation disappears from $\mu_c = 0.10$ onward.

Overall, both higher bank capitalization and higher firm cash shares are associated with greater system-wide shock tolerance, but their sectoral effects differ. At high capitalization levels, the banking-sector threshold continues to increase while the firm-sector threshold changes little. At low firm liquidity, by contrast, the firm-sector threshold lies substantially below the banking-sector threshold, with the two thresholds coinciding from $\mu_c = 0.10$ onward.

3.5. Effects of Interfirm Link Formation and Credit Exposure

We next examine two aspects of the interfirm credit layer: link formation, governed by the parameter d_f, and credit exposure, governed by μ_{ff}, the share of accounts payable in firms' total assets. We vary these parameters separately to examine their effects on system-wide and sector-specific shock tolerance. Figure 6 summarizes the resulting collapse thresholds.

3.5.1. Interfirm Link Formation

The parameter d_f governs interfirm link formation. We vary d_f over seven values from 0.25 to 1.00, holding the remaining model parameters at their baseline settings. Increasing d_f produces a denser interfirm network and changes counterparties and the allocation of interfirm exposures, thereby affecting firms' accounts receivable and residual fixed assets. The mean in-degree and mean out-degree of the interfirm network are identical and increase from 47.05 at $d_f = 0.25$ to 94.47 at the baseline value of 0.50, reaching 189.33 at $d_f = 1$. The measured mean degrees in the interbank and bank–firm layers remain approximately stable. Accordingly, this experiment should be interpreted as a change in the interfirm network configuration rather than as isolating a pure topological effect.

Figure 6a shows that the system-wide collapse threshold increases monotonically as d_f increases. It rises from 0.116714 at $d_f = 0.25$ to 0.143545 at the baseline value of 0.50 and reaches 0.166556 at $d_f = 1$. The increments in the estimated threshold become progressively smaller as d_f increases. Thus, a higher d_f, which produces a denser interfirm network, is associated with greater system-wide shock tolerance under the specified network-generation and balance-sheet rules.

The sector-specific thresholds remain closely aligned. As shown in Figure 6c, at $d_f = 0.25$, the banking-sector threshold is 0.116714, compared with 0.113904 for the firm sector, giving $S_{\text{sector}} = 0.002810$. The estimated thresholds coincide at the remaining six parameter

values. Thus, unlike the pronounced sectoral separation observed for some balance-sheet configurations in Section 3.4, changes in d_f shift the bank- and firm-sector thresholds toward higher shock intensities without producing substantial separation between them.

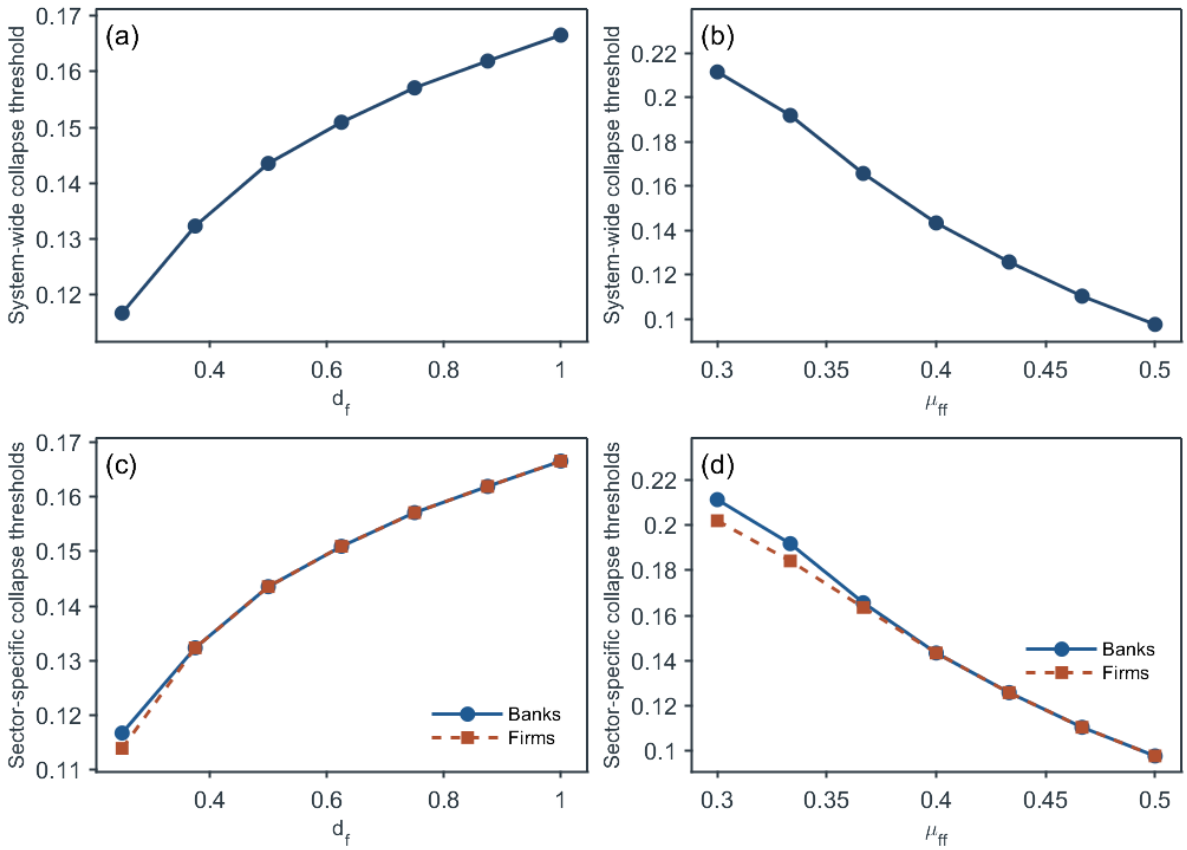


Figure 6. Effects of interfirm link formation and credit exposure on collapse thresholds. (**a**,**b**) System-wide collapse thresholds under changes in the interfirm link-formation parameter d_f and the share of accounts payable in firms' total assets μ_{ff}, respectively. Circles show point estimates; uncapped vertical lines indicate 95% bootstrap confidence intervals. The intervals are very narrow relative to the vertical-axis scale and may be difficult to distinguish from the circles. (**c**,**d**) Bank- and firm-sector collapse thresholds under changes in d_f and μ_{ff}, respectively.

3.5.2. Interfirm Credit Exposure

The parameter μ_{ff} specifies the share of accounts payable in firms' total assets. We vary μ_{ff} over seven values from 0.30 to 0.50, holding the remaining model parameters at their baseline settings. Increasing μ_{ff} raises firms' accounts payable and, consequently, the corresponding accounts receivable of their suppliers. Under the accounting identities in Section 2.4, higher accounts payable reduce residual firm equity, while higher accounts receivable reduce residual fixed assets. The experiment therefore captures a combined change in trade credit exposure and balance-sheet composition.

In contrast to the d_f experiment, increasing trade credit exposure shifts the system-wide collapse threshold toward lower shock intensities. As shown in Figure 6b, the threshold decreases monotonically from 0.211394 at $\mu_{ff} = 0.30$ to 0.143333 at the baseline value of 0.40 and reaches 0.097600 at $\mu_{ff} = 0.50$. Thus, higher trade credit exposure is associated with lower system-wide shock tolerance under the specified accounting and contagion rules.

The sector-specific thresholds exhibit a modest separation at lower exposure levels. As shown in Figure 6d, at $\mu_{ff} = 0.30$, the firm-sector threshold is 0.202000, compared with 0.211394 for the banking sector, giving $S_{\text{sector}} = 0.009394$. The separation decreases to

0.007665 at $\mu_{ff} = 0.3333$ and 0.002130 at $\mu_{ff} = 0.3667$. From the baseline value of $\mu_{ff} = 0.40$ onward, the estimated thresholds of the two sectors coincide. Thus, the modest sectoral separation observed at low exposure levels disappears as μ_{ff} reaches its baseline value.

Overall, the two parameter changes are associated with opposite movements in system-wide shock tolerance under the specified model configurations. Increasing d_f, which produces a denser interfirm network, is associated with a higher system-wide collapse threshold, whereas increasing trade credit exposure is associated with a lower threshold. In both experiments, the bank- and firm-sector thresholds remain more closely aligned than in the balance-sheet configurations examined in Section 3.4.

3.6. Robustness and Numerical Sensitivity

We examine the robustness of the main results to two methodological choices: the bankruptcy-fraction cutoff used to define sectoral collapse and the order in which bankrupt entities are liquidated during contagion. We also report the feasibility of the generated configurations in the baseline experiment. Additional sensitivity analysis with respect to bank and firm liquidation recovery rates is reported in Section S3 of the Supplementary Materials.

3.6.1. Sensitivity to the Collapse Criterion

The preceding analysis defines sectoral collapse using a bankruptcy-fraction cutoff of $q = 0.99$. We assess sensitivity to this choice by recalculating collapse classifications using $q = 0.90$ and $q = 0.95$, without rerunning the underlying simulations. In the baseline three-layer experiment, none of the 50,500 realizations changes its system-wide-collapse classification. Consequently, the estimated system-wide collapse-probability curve, collapse threshold, and $W_{10–90}$ are identical under the three cutoffs.

The system-wide-collapse classifications also remain unchanged across all simulated shock intensities in the five degree-controlled finite-size experiments. In the matched comparisons of Section 3.3, the paired system-wide-collapse classifications are likewise unchanged under both recovery-rate matching and recovery-amount matching. Thus, over the tested range of q, the system-wide-collapse classifications and the corresponding threshold-based results are insensitive to the bankruptcy-fraction cutoff.

3.6.2. Sensitivity to Liquidation Order

The contagion algorithm liquidates bankrupt entities sequentially, with each liquidation updating the financial positions of other entities. Individual outcomes may therefore depend on liquidation order. Under the baseline convention, banks are processed before firms in each contagion round after the initial shocks, and entities within each sector are processed in ascending order of their indices. We examine four alternative liquidation priorities: random order, larger entities first, smaller entities first, and firms before banks. Entity size is measured by initial total assets. Each alternative uses the same initial financial configuration and shocked firms as its baseline counterpart, while the initial shock sequence is unchanged. Only the liquidation order in subsequent contagion rounds is varied.

The experiment covers 16 shock intensities between 0.132 and 0.154, with 500 paired realizations at each intensity. A reproduction test confirms that the baseline implementation yields identical results. Table 8 reports the estimated system-wide collapse thresholds together with the proportion of paired realizations whose system-wide-collapse classification differs from that under the baseline order.

Table 8. Sensitivity of the system-wide collapse threshold and classifications to liquidation order. Threshold differences are measured relative to the baseline order. The overall flip rate is the proportion of the 8000 paired realizations whose system-wide-collapse classifications differ between the alternative and baseline liquidation orders. The maximum flip rate is the largest such proportion at an individual shock intensity.

Liquidation Priority	$\varepsilon_{50}^{\mathrm{sys}}$	$\Delta\varepsilon_{50}$	Overall Flip Rate	Maximum Flip Rate
Baseline	0.142629	–	–	–
Random	0.142742	+0.000113	2.51%	4.4%
Larger first	0.142705	+0.000076	2.75%	4.8%
Smaller first	0.142694	+0.000065	2.61%	4.8%
Firms first	0.142629	0	0.0125%	0.2%

The alternative liquidation orders produce only small changes in the estimated system-wide collapse threshold. The largest difference is 0.000113 under random priority, with a 95% paired-bootstrap confidence interval of $[-0.000010, 0.000430]$. The corresponding confidence intervals for the larger-first and smaller-first conventions also include zero, while the firms-first convention yields the same estimated threshold as the baseline.

At the realization level, however, liquidation order alters the system-wide-collapse classification for some paired realizations. Random, larger-first, and smaller-first priorities produce classification flips in approximately 2.5–2.8% of these pairs, with flip rates reaching 4.8% at individual shock intensities. By contrast, the firms-first convention changes only one of the 8000 paired classifications. Thus, alternative liquidation orders can alter the system-wide-collapse classification of individual realizations while producing only small changes in the estimated system-wide collapse threshold.

3.6.3. Feasibility of Baseline Configurations

The baseline experiment uses 101 shock intensities from 0.120 to 0.220, with 500 accepted realizations at each grid point. A new network and its associated balance sheets are generated for each realization. The generator checks the nonnegativity of residual bank deposits, firm fixed assets, and firm equity, and configurations that fail these checks are rejected. Of the 50,520 candidate systems generated for the baseline experiment, 20 are rejected, leaving 50,500 accepted realizations and a rejection rate of approximately 0.040%. These checks confirm that the accepted baseline configurations satisfy the feasibility conditions implemented in the model.

4. Conclusions

This study examines system-wide collapse in a three-layer financial network integrating interbank lending, bank–firm loans, and interfirm trade credit within a unified balance-sheet framework. In the baseline experiment, the bankruptcy fractions of both sectors exhibit pronounced bimodality, with the between-state component accounting for nearly all of the variance. The steep rise in mean bankruptcy fractions is therefore driven primarily by an increasing probability of system-wide collapse, rather than by a gradual upward shift in bankruptcy fractions across most realizations. Across five approximately degree-controlled system sizes, susceptibility peaks grow nearly linearly with sector size, while the shock-intensity interval over which the system-wide collapse probability rises from 10% to 90% generally narrows. Together, these results provide consistent finite-size evidence of phase-transition-like behavior in system-wide collapse.

The results also show that no single measure fully characterizes the system's response to shocks. The system-wide collapse probability measures collapse likelihood at a given shock intensity, while the system-wide and sector-specific collapse thresholds characterize shock tolerance and differences between bank and firm responses. Layer-specific loss ratios

provide complementary information on financial damage. In particular, losses may remain substantial without system-wide collapse and may continue to increase even after the system-wide collapse probability has reached one. Collapse likelihood, shock tolerance, and financial damage therefore describe distinct aspects of systemic risk in the model.

The comparative experiments further show how endogenous interfirm credit propagation and selected model parameters are associated with changes in shock tolerance. In the matched counterfactual experiments, the two-layer configurations without endogenous interfirm credit propagation have higher system-wide collapse thresholds than the three-layer model under both recovery-matching conventions, although the magnitude of the difference depends on how recovery resources are matched. The balance-sheet experiments show that higher bank capitalization and higher firm cash shares are associated with greater system-wide shock tolerance under the corresponding balance-sheet adjustments, while the sector-specific responses can differ substantially. In the interfirm-layer experiments, increasing d_f, which produces a denser interfirm network, is associated with a higher system-wide collapse threshold, whereas increasing the share of accounts payable in firms' total assets is associated with a lower threshold. These results indicate that interfirm propagation, balance-sheet composition, link formation, and credit exposure are all associated with changes in the system-wide collapse threshold under the specified model configurations.

The main results are also robust to the methodological choices examined here. Reclassifying sectoral collapse using bankruptcy-fraction cutoffs of 0.90, 0.95, and 0.99 leaves the system-wide-collapse classifications unchanged in the baseline, finite-size, and matched-model experiments. Alternative liquidation orders produce only small changes in the estimated system-wide collapse threshold, although they alter the system-wide-collapse classification of some individual realizations. In addition, the accepted baseline configurations satisfy the feasibility conditions implemented in the balance-sheet generator.

These findings are conditional on the model's specified network-generation, balance-sheet, shock, and contagion mechanisms. The generated networks are stylized configurations rather than empirically calibrated representations of actual financial systems, and the present analysis considers firm-originated shocks within a specified contagion process. It therefore does not establish whether the same patterns persist under alternative shock origins, targeting strategies, correlated shocks, or additional channels such as common asset exposures, fire sales, deposit runs, refinancing, and adaptive behavior. Further empirical calibration and extensions incorporating richer financial interactions and dynamic responses are needed to assess the broader applicability of the results.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/e28101148/s1: Section S1: Pseudocode for the Default-Contagion and Liquidation Procedure (Algorithms S1 and S2); Section S2: Degree Distributions of the Baseline Networks (Figure S1); Section S3: Sensitivity to Liquidation Recovery Rates (Table S1 and Figure S2). The supplementary files also include the simulation code and numerical data used to generate the results reported in the article and Supplementary Materials. The code was developed and validated using MATLAB R2024b. Documentation, data dictionaries, and licensing information are included in the accompanying code and data packages.

Author Contributions: Conceptualization, N.X. and Y.D. (Yazheng Dang); methodology, N.X. and Y.D. (Yamin Dou); software, N.X. and Y.D. (Yamin Dou); validation, N.X. and Y.D. (Yazheng Dang); formal analysis, N.X. and Y.D. (Yamin Dou); investigation, Y.D. (Yamin Dou); visualization, Y.D. (Yamin Dou); writing—original draft preparation, N.X. and Y.D. (Yamin Dou); writing—review and editing, N.X. and Y.D. (Yazheng Dang); supervision, Y.D. (Yazheng Dang); N.X. and Y.D. (Yamin Dou) contributed equally to this work. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No empirical datasets were used in this study. The numerical data generated and analyzed in the simulations, together with the MATLAB code used to produce the reported results, are provided in the Supplementary Materials. The code was developed and validated using MATLAB R2024b. Additional documentation describing the files and data structure is included with the supplementary code and data packages.

Acknowledgments: The authors would like to thank Changgui Gu for helpful discussions and valuable suggestions.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Giglio, S.; Kelly, B.; Pruitt, S. Systemic risk and the macroeconomy: An empirical evaluation. *J. Financ. Econ.* **2016**, *119*, 457–471. [CrossRef]
2. Zhang, D.; Hu, M.; Ji, Q. Financial markets under the global pandemic of COVID-19. *Finance Res. Lett.* **2020**, *36*, 101528. [CrossRef] [PubMed]
3. Liu, X.; Megginson, W.; Tran, N.; Wei, S. Who loses most when big banks suddenly fail? Evidence from Silicon Valley Bank collapse. *Finance Res. Lett.* **2024**, *59*, 104806. [CrossRef]
4. Galati, L.; Capalbo, F. Silicon Valley Bank bankruptcy and stablecoins stability. *Int. Rev. Financ. Anal.* **2024**, *91*, 103001. [CrossRef]
5. Martins, A.M. The collapse of Silicon Valley Bank and Credit Suisse and their impact on other U.S. banks. *Appl. Econ. Lett.* **2025**, *32*, 1126–1130. [CrossRef]
6. Allen, F.; Gale, D. Financial contagion. *J. Polit. Econ.* **2000**, *108*, 1–33. [CrossRef]
7. Freixas, X.; Parigi, B.M.; Rochet, J.-C. Systemic risk, interbank relations, and liquidity provision by the central bank. *J. Money Credit Bank.* **2000**, *32*, 611–638. [CrossRef]
8. Iori, G.; Jafarey, S.; Padilla, F.G. Systemic risk on the interbank market. *J. Econ. Behav. Organ.* **2006**, *61*, 525–542. [CrossRef]
9. Gai, P.; Kapadia, S. Contagion in financial networks. *Proc. R. Soc. A* **2010**, *466*, 2401–2423. [CrossRef]
10. Lenzu, S.; Tedeschi, G. Systemic risk on different interbank network topologies. *Physica A* **2012**, *391*, 4331–4341. [CrossRef]
11. Georg, C.-P. The effect of the interbank network structure on contagion and common shocks. *J. Bank. Finance* **2013**, *37*, 2216–2228. [CrossRef]
12. Boss, M.; Elsinger, H.; Summer, M.; Thurner, S. Network topology of the interbank market. *Quant. Finance* **2004**, *4*, 677–684. [CrossRef]
13. Soramäki, K.; Bech, M.L.; Arnold, J.; Glass, R.J.; Beyeler, W.E. The topology of interbank payment flows. *Physica A* **2007**, *379*, 317–333. [CrossRef]
14. Kanno, M. The network structure and systemic risk in the Japanese interbank market. *Jpn. World Econ.* **2015**, *36*, 102–112. [CrossRef]
15. Sun, A.J.; Chan-Lau, J.A. Financial networks and interconnectedness in an advanced emerging market economy. *Quant. Finance* **2017**, *17*, 1833–1858. [CrossRef]
16. Upper, C.; Worms, A. Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *Eur. Econ. Rev.* **2004**, *48*, 827–849. [CrossRef]
17. Degryse, H.; Nguyen, G. Interbank exposures: An empirical examination of contagion risk in the Belgian banking system. *Int. J. Cent. Bank.* **2007**, *3*, 123–171.
18. Langfield, S.; Liu, Z.; Ota, T. Mapping the UK interbank system. *J. Bank. Finance* **2014**, *45*, 288–303. [CrossRef]
19. Montagna, M.; Kok, C. *Multi-Layered Interbank Model for Assessing Systemic Risk*; ECB Working Paper, No. 1944; European Central Bank (ECB): Frankfurt am Main, Germany, 2016. [CrossRef] [PubMed]
20. Poledna, S.; Molina-Borboa, J.L.; Martínez-Jaramillo, S.; van der Leij, M.; Thurner, S. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *J. Financ. Stab.* **2015**, *20*, 70–81. [CrossRef]
21. Bargigli, L.; di Iasio, G.; Infante, L.; Lillo, F.; Pierobon, F. The multiplex structure of interbank networks. *Quant. Finance* **2015**, *15*, 673–691. [CrossRef]
22. Bardoscia, M.; Barucca, P.; Battiston, S.; Caccioli, F.; Cimini, G.; Garlaschelli, D.; Saracco, F.; Squartini, T.; Caldarelli, G. The physics of financial networks. *Nat. Rev. Phys.* **2021**, *3*, 490–507. [CrossRef]

23. de Castro Miranda, R.C.; Tabak, B.M. *Contagion Risk within Firm–Bank Bivariate Networks*; Working Paper Series, No. 322; Central Bank of Brazil: Brasília, Brazil, 2013.
24. Silva, T.C.; da Silva Alexandre, M.; Tabak, B.M. Bank lending and systemic risk: A financial-real sector network approach with feedback. *J. Financ. Stab.* **2018**, *38*, 98–118. [CrossRef]
25. Lux, T. A model of the topology of the bank–firm credit network and its role as channel of contagion. *J. Econ. Dyn. Control* **2016**, *66*, 36–53. [CrossRef]
26. Gao, Q. Systemic risk analysis of multi-layer financial network system based on multiple interconnections between banks, firms, and assets. *Entropy* **2022**, *24*, 1252. [CrossRef] [PubMed]
27. Bernard, A.B.; Dhyne, E.; Magerman, G.; Manova, K.; Moxnes, A. The origins of firm heterogeneity: A production network approach. *J. Polit. Econ.* **2022**, *130*, 1765–1804. [CrossRef]
28. Boissay, F. *Credit Chains and the Propagation of Financial Distress*; ECB Working Paper, No. 573; European Central Bank (ECB): Frankfurt am Main, Germany, 2006. [CrossRef]
29. Wang, L.; Li, S.; Wu, C. Risk contagion in inter-firm credit guarantee network. *Physica A* **2019**, *526*, 120842. [CrossRef]
30. Battiston, S.; Delli Gatti, D.; Gallegati, M.; Greenwald, B.; Stiglitz, J.E. Credit chains and bankruptcy propagation in production networks. *J. Econ. Dyn. Control* **2007**, *31*, 2061–2084. [CrossRef]
31. Weisbuch, G.; Battiston, S. From production networks to geographical economics. *J. Econ. Behav. Organ.* **2007**, *64*, 448–469. [CrossRef]
32. Reisch, T.; Heiler, G.; Diem, C.; Klimek, P.; Thurner, S. Monitoring supply networks from mobile phone data for estimating the systemic risk of an economy. *Sci. Rep.* **2022**, *12*, 13347. [CrossRef] [PubMed]
33. Mungo, L.; Lafond, F.; Astudillo-Estévez, P.; Farmer, J.D. Reconstructing production networks using machine learning. *J. Econ. Dyn. Control* **2023**, *148*, 104607. [CrossRef]
34. Inoue, H.; Todo, Y. Firm-level propagation of shocks through supply-chain networks. *Nat. Sustain.* **2019**, *2*, 841–847. [CrossRef]
35. Barro, D.; Basso, A. Credit contagion in a network of firms with spatial interaction. *Eur. J. Oper. Res.* **2010**, *205*, 459–468. [CrossRef]
36. Delli Gatti, D.; Gallegati, M.; Greenwald, B.; Russo, A.; Stiglitz, J.E. The financial accelerator in an evolving credit network. *J. Econ. Dyn. Control* **2010**, *34*, 1627–1650. [CrossRef]
37. Sui, X.; Li, L.; Chen, X. Risk contagion caused by interactions between credit and guarantee networks. *Physica A* **2020**, *539*, 122867. [CrossRef]
38. Landaberry, V.; Caccioli, F.; Rodriguez-Martinez, A.; Baron, A.; Martinez-Jaramillo, S.; Lluberas, R. The contribution of the intra-firm exposures network to systemic risk. *Lat. Am. J. Cent. Bank.* **2021**, *2*, 100032. [CrossRef]
39. Li, S.; Sui, X. Contagion risk in endogenous financial networks. *Chaos Solitons Fractals* **2016**, *91*, 591–597. [CrossRef]
40. Jin, Q.; Sun, L.; Chen, Y.; Hu, Z.-L. Financial risk contagion based on dynamic multi-layer network between banks and firms. *Physica A* **2024**, *638*, 129624. [CrossRef]
41. Lei, J.; Shen, P.; Li, Z. Multilayer bank–firm networks and financial risk contagion. *Physica A* **2025**, *675*, 130838. [CrossRef]
42. Axtell, R.L. Zipf distribution of U.S. firm sizes. *Science* **2001**, *293*, 1818–1820. [CrossRef] [PubMed]
43. Bremus, F.; Buch, C.M.; Russ, K.N.; Schnitzer, M. Big banks and macroeconomic outcomes: Theory and cross-country evidence of granularity. *J. Money Credit Bank.* **2018**, *50*, 1785–1825. [CrossRef]
44. De Masi, G.; Gallegati, M. Bank–firms topology in Italy. *Empir. Econ.* **2012**, *43*, 851–866. [CrossRef]
45. Montagna, M.; Lux, T. *Hubs and Resilience: Towards More Realistic Models of the Interbank Markets*; Kiel Working Paper, No. 1826; Kiel Institute for the World Economy: Kiel, Germany, 2013.
46. Saito, Y.U.; Watanabe, T.; Iwamura, M. Do larger firms have more interfirm relationships? *Physica A* **2007**, *383*, 158–163. [CrossRef]
47. Fujiwara, Y.; Aoyama, H. Large-scale structure of a nation-wide production network. *Eur. Phys. J. B* **2010**, *77*, 565–580. [CrossRef]
48. Atalay, E.; Hortaçsu, A.; Roberts, J.; Syverson, C. Network structure of production. *Proc. Natl. Acad. Sci. USA* **2011**, *108*, 5199–5202. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.