POLITECNICO DI TORINO
Repository ISTITUZIONALE

Default robustness and worst-case losses in financial networks

Original
Default robustness and worst-case losses in financial networks / Calafiore, G.C., Fracastoro, G., Proskurnikov, A.V.. - In: APPLIED NETWORK SCIENCE. - ISSN 2364-8228. - 10:1(2025). [10.1007/s41109-025-00728-5]

Availability:
This version is available at: 11583/3009256 since: 2026-03-26T13:17:41Z

Publisher:
Springer Nature

Published
DOI:10.1007/s41109-025-00728-5

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

08 October 2026

Calafiore *et al. Applied Network Science* (2025) 10:48
https://doi.org/10.1007/s41109-025-00728-5

Applied Network Science

RESEARCH Open Access

Default robustness and worst-case losses in financial networks

Giuseppe C. Calafiore[1*†], Giulia Fracastoro[1†] and Anton V. Proskurnikov[1†]

†The three authors have contributed equally.

*Correspondence:
Giuseppe C. Calafiore
giuseppe.calafiore@polito.it
[1]Department of Electronics and Telecommunications (DET), Politecnico di Torino, Corso Duca Degli Abruzzi 24, Torino, Italy

Abstract

We analyze the resilience and worst-case losses in a financial network of banks linked by mutual liabilities and shared exposures to external assets. Abrupt asset price changes lead to simultaneous shocks to bank balance sheets, potentially triggering cascades of defaults. In this context, we introduce first the concept of *default resilience margin*, ϵ^*, defined as the maximum amplitude of asset prices fluctuations that the network can sustain without generating defaults. Such threshold value is computed by considering two different measures of price fluctuations, one based on the maximum individual variation of each asset, and the other based on the sum of all the asset's absolute variations. For any price perturbation having amplitude no larger than ϵ^*, the network absorbs the shocks and remains free of defaults. When the perturbation amplitude goes *beyond* the ϵ^* level, however, defaults may occur. We assume that external liabilities have higher priority than interbank claims, thereby distinguishing between banks that default on obligations to other banks and those that are fully insolvent, i.e., unable to meet even their external obligations. We propose an explicit method to determine the upper bound on shock amplitude, $\epsilon_{\rm ub}$ below which such insolvencies do not occur–that is, all banks are able to fulfill their external liabilities, though some may default on their interbank obligations. For each shock amplitude $\epsilon \in (\epsilon^*, \epsilon_{\rm ub})$, we show how to compute the worst-case systemic loss, that is, the total financial network shortfall under the worst-case scenario of price variation of given magnitude.

Keywords Financial networks, Systemic risk, Contagion resilience

Introduction

The global financial crisis of 2007–2008 sparked a surge of research into the vulnerabilities of financial systems, their susceptibility to shocks, and the phenomenon of financial contagion (Glasserman and Young 2016; Pacelli et al. 2025; Allen and Gale 2000). Significant attention has been dedicated to unraveling the intricate relationships between financial institutions and capital flows, with a particular emphasis on their impact on the resilience of the global financial system–its ability to absorb and insulate itself from external shocks. This body of research underlines the critical role of interconnectedness in determining the stability of financial systems.

Resilience and fragility in interbank networks

Springer Open

© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

The link between a financial network's structure and its resilience to shock propagation is complex and not yet fully understood. On one hand, dense interconnections can help redistribute liquidity and provide implicit insurance against moderate shocks. Through bilateral agreements, banks effectively share portfolio losses across many counterparties, making the system more robust to moderate external disturbances, as shown in models such as Allen and Gale (2000); Freixas et al. (2000); Babus (2016). On the other hand, dense networks are more vulnerable to severe shocks that surpass critical thresholds or affect key nodes. Interconnections facilitate shock transmission, increasing the risk of cascading defaults (Glasserman and Young 2016). This vulnerability has been proven in regular networks (Acemoglu et al. 2015) and supported by simulations on random graphs (Nier et al. 2007; Hurd 2023; Erol et al. 2023), which show that greater connectivity can either enhance or undermine resilience depending on the nature and scale of the shock. Both effects can be amplified by the multiplex structure of financial networks, where different layers represent distinct types of interbank dependencies and assets (Rio-Chanona et al. 2020; Jurakovaite and Gaigaliene 2024; Zino et al. 2025).

Beyond the transmission of payment shortfalls, financial contagion can also stem from asset commonality—banks' shared exposures to declines in the value of external assets (Cifuentes et al. 2005; Glasserman and Young 2016; Allen et al. 2012). Such common exposures can trigger systemic liquidity shocks that affect all holders of the impacted assets (Banerjee and Feinstein 2022), regardless of direct interbank connections. Nevertheless, network structure can amplify these risks in several ways. An initial asset shock may lead to defaults, prompting distressed banks to conduct "fire sales" of their holdings (Demirer et al. 2021). These sales depress asset prices further, potentially igniting information contagion (Glasserman and Young 2016), where falling confidence and fear of widespread defaults drive down the value of even sound assets, affecting the whole financial system. Circular mechanisms of financial contagion–particularly those driven by fire sales at distressed banks–have been examined in several studies (see, e.g., Barucca et al. (2021) and the recent overview in Caiazzo and Zazzaro (2025), which also provides pointers to related literature).

Systemic loss quantification under external shocks

The Eisenberg–Noe framework (Eisenberg and Noe 2001) is one of the most widely recognized models for analyzing financial contagion and evaluating the impact of default cascades in interbank networks. Their model addresses the clearing of mutual liabilities in the event of institutional defaults triggered by external shocks, offering an explicit method to compute both the clearing payment matrix and the total shortfall, or systemic loss (Glasserman and Young 2016). Extensions of the Eisenberg–Noe model have incorporated key features of real-world financial systems, including the prioritization of liabilities to the non-financial sector over interbank claims (Elsinger et al. 2006), default costs (Rogers and Veraart 2013), multiple debt maturities (Kusnetsov and Veraart 2019), asset liquidation costs and strategies (Feinstein 2017; Feinstein and Hałłaj 2023), and cross-holdings (Weber and Weske 2017). The computation and uniqueness of the clearing payment under external shocks have been extensively studied in the literature (Massai et al. 2022; Csóka and Herings 2024; Kanellopoulos et al. 2024). Additional lines of research have addressed more specific issues, such as the sensitivity of the clearing payments to small liquidity shocks (Liu and Staum 2010) and to slight changes in interbank liabilities (Feinstein et al. 2018). The Eisenberg-Noe framework has also served as a foundation for the

development of dynamic models of clearing (Calafiore et al. 2023; Banerjee et al. 2025) with time-evolving obligations and payments. Most of the aforementioned research, however, assumes that the asset and liability positions of all banks are fully known.

Since banks are typically exposed to common assets, fluctuations in market prices generate shocks that are, on the one hand, difficult to predict and, on the other hand, impact the balance sheets of all affected institutions. Existing research that accounts for these uncertainties primarily focuses on random asset values, as reflected in the extensive literature on systemic risk measures (see Ararat and Meimanjan (2023), Jarrow et al. (2024) and references therein), network valuation theory (Barucca et al. 2020; Pallante et al. 2025), and topology optimization to enhance resilience to random shocks (Li and Zhang 2024) or to identify network configurations most prone to contagion (Hu et al. 2024).

The problem in question and contributions

In this paper, we address a novel class of problems in systemic risk evaluation under asset commonality. Unlike existing approaches, asset price fluctuations are not assumed to follow known probabilistic distributions but are instead constrained in magnitude–for instance, each asset price may vary by at most ε. Within this framework, we explore two key questions: (1) What is the maximum level of asset price variation that a banking system can withstand without triggering any defaults? and (2) once this resilience threshold is exceeded, what is the worst-case systemic loss resulting from cascaded defaults within the financial network? Our analysis builds on an extended version of the Eisenberg-Noe model, introduced in Elsinger et al. (2006), which permits banks to experience a negative net liquidity inflow from the external sector–potentially leading to insolvency. In addition, we assume that the asset side of some banks' balance sheets includes, beyond interbank claims and cash flows from the non-financial sector, stakes in a set of external assets.

We examine two distinct categories of uncertain fluctuations in external asset prices: (i) a scenario in which the total sum of absolute variations in asset prices does not exceed a predetermined threshold of ϵ units (referred to as an ℓ_1 price perturbation), and (ii) a scenario in which the price of each individual asset fluctuates by no more than ϵ units (referred to as an ℓ_∞ price perturbation). We evaluate the network's resilience against the most severe shock of each type and assess the worst-case systemic loss caused by these shocks. The contributions of this paper are as follows:

1. We compute the network's *default resilience margin*, denoted by ϵ^*, which represents the largest value of ϵ such that the entire system remains default-free under all asset price perturbations with amplitude not exceeding ϵ. The computation of ϵ^* provides a quantitative measure of the financial network's *robustness* to asset price fluctuations.
2. When the price perturbation level ϵ exceeds the threshold ϵ^*, defaults may occur at some nodes, leading to the problem of evaluating the *worst-case systemic loss* (the total shortfall of financial institutions). The term "worst-case" refers to the maximum systemic loss resulting from the most adverse shock of magnitude not exceeding ϵ. It is important to note that the loss function is highly nonlinear in ϵ and lacks a closed-form analytical expression, necessitating numerical methods. By leveraging linear programming duality, we show that for values of ϵ exceeding ϵ^* but bounded above by a certain threshold $\epsilon_{\rm ub}$, both the worst-case loss and the corresponding worst-case shock can be efficiently computed by solving a linear programming problem.

3. We show that $\epsilon_{\rm ub} \geq \epsilon^*$ represents the largest shock magnitude the network can withstand without any bank becoming fully insolvent–that is, unable to meet its external liabilities even with all interbank payments suspended. Once ϵ exceeds this threshold, there exist shocks of magnitude ϵ that lead to insolvency at one or more banks. In contrast, for all $\epsilon \leq \epsilon_{\rm ub}$, every bank remains solvent with respect to its external obligations. We compute the value $\epsilon_{\rm ub} \geq \epsilon^*$, henceforth called the *insolvency resilience margin*.

Organization of the paper

The remainder of the paper is structured as follows. Section "Models and methods" presents a method for evaluating systemic loss caused by external shocks, based on a generalized Eisenberg–Noe model and the concept of a clearing vector. The main contributions are summarized in Section "Main results". In particular, Section "The default resilience margin" addresses the computation of the default resilience margin; Section "The worst-case loss curve" focuses on the analysis and computation of the worst-case systemic loss; and Section "The insolvency resilience margin" presents the computation of the insolvency resilience threshold. Our principal findings are illustrated via numerical examples in Section "Numerical examples". Conclusions are drawn in Section "Discussion and conclusions". For improved readability, the technical proofs and other details of our key results are contained in the Appendix, which can be found in the supplementary file.

Models and methods

Notation. Given a finite set $\mathcal{V}$, the symbol $|\mathcal{V}|$ stands for its cardinality. For two vectors $a, b \in \mathbb{R}^n$, we write $a \leq b$ if $a_i \leq b_i \; \forall i$. The relation $\geq$ for vectors is defined similarly. If $a \leq b$, then $[a, b]$ denotes the set of vectors c such that $a \leq c \leq b$. We write $a \geq 0$ if $a_i \geq 0 \, \forall i$. Operations $\min, \max$ are applied to vectors element-wise, e.g., $\max(a, b) = (\max(a_i, b_i))_{i=1}^n$. Given a vector a, we denote its positive part by $[a]^+ \doteq \max(a, 0)$. For a real number $z \in \mathbb{R}$, let $\text{sign}(z) \doteq 1$ if $z > 0$, $\text{sign}(z) \doteq -1$ if $z < 0$ and $\text{sign}(z) \doteq 0$ if $z = 0$. The symbol 1 represents a (column) vector of all ones, with dimension inferred from context.

Financial networks with asset commonality

We follow the basic model introduced in Eisenberg and Noe (2001) and Glasserman and Young (2016), with the added assumption that banks are also interconnected through asset commonality (Feinstein and Hałlaj 2023). The parameters of the financial network, described in the next paragraphs, are summarized in Table 1.

Table 1 Parameters of a financial network

$\mathcal{V} = \{1, \dots, n\}$	Set of banks
$\mathcal{M} = \{1, \dots, m\}$	Set of external assets
$\bar{p}_{ij}$	Liability from bank $i \in \mathcal{V}$ to bank $j \in \mathcal{V}$
$\bar{p}_i$	Total liability of bank $i \in \mathcal{V}$
p_{ij}	Actual clearing payment bank $i \in \mathcal{V}$ to bank $j \in \mathcal{V}$
c_i^e	Net liquidity of bank $i \in \mathcal{V}$
s_{ij}	Shares of asset $j \in \mathcal{M}$ owned by bank $i \in \mathcal{V}$
v_j	Value per shares of asset $j \in \mathcal{M}$
c_i	External net position of bank $i \in \mathcal{V}$

Interbank liabilities. A *financial network* is represented as a weighted directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \bar{P})$, where the node set $\mathcal{V} \doteq \{1, 2, \ldots, n\}$ represents financial institutions (e.g., banks, funds, insurance companies). The weighted adjacency matrix $\bar{P} = (\bar{p}_{ij})$ encodes nominal mutual liabilities, with $\bar{p}_{ij} \geq 0$ indicating that institution *i* owes $\bar{p}_{ij}$ currency units to institution $j \neq i$. The liability arises when institution *i* borrows funds from institution *j* under a contractual agreement–such as an interbank loan, repurchase agreement, or other credit arrangement. We assume that $\bar{p}_{ii} = 0$ for all $i \in \mathcal{V}$. An arc $(i, j) \in \mathcal{E}$ exists if and only if $\bar{p}_{ij} > 0$, indicating that node *i* has a liability to node *j*.

The total liability of node i to other nodes is given by its weighted out-degree, $\bar{p}_i \doteq \sum_{j \neq i} \bar{p}_{ij}$. It is convenient to define the vector of total liabilities as $\bar{p} \doteq (\bar{p}_i)_{i \in \mathcal{V}} = \bar{P}1$.

Assets and net liquidity positions. In addition to interbank obligations, banks are interconnected through portfolio overlaps (Feinstein and Hałlaj 2023), reflecting shared exposures to common external assets (e.g., government and corporate bonds, publicly traded stocks, commodities, and related financial instruments).

Specifically, we assume that banks may hold shares in a set of marketable assets, denoted by $\mathcal{M} \doteq \{1, \ldots, m\}$, with corresponding prices $v_1, \ldots, v_m \geq 0$. Bank *i* holds s_{ij} shares[1] of asset *j*, and the total market value of its external asset portfolio is given by $z_i \doteq \sum_{j=1}^{m} s_{ij} v_j = \sigma_i^\top v$, where $\sigma_i^\top \doteq [s_{i1}, \ldots, s_{im}]$ and $v^\top \doteq [v_1, \ldots, v_m]$.

Each bank has a net liquidity position $c_i^e \in \mathbb{R}$ from the external sector, representing the difference between incoming cash flows and outgoing funding obligations to entities outside the financial network. The total external net position of bank *i*–that is, the total amount it receives from outside the financial system–is thus given by $c_i \doteq c_i^e + z_i$, for all $i \in \mathcal{V}$. Denoting with $S = (s_{jk})$ the $n \times m$ matrix of asset shares and with *v* the *m*-vector of asset prices, the vector *z* containing the total asset values held by the banks is given by $z \doteq Sv$, and the net inflow vector is thus found as

$$c = c^e + Sv. \tag{1}$$

It should be noted that the net external position c_i can be negative.[2] We represent it as $c_i = c_i^+ - c_i^-$, where $c_i^+ \geq 0$ denotes the total external inflow to bank *i* (including both cash inflows and the market value of held assets), and $c_i^- \geq 0$ represents its liabilities to the external sector. With this notation, the nominal asset and liability sides of the balance sheet for bank *i* are given, respectively, by

$$\bar{\phi}_i^{\text{in}}(c) \doteq c_i^+ + \sum\nolimits_{k \neq i} \bar{p}_{ki}, \quad \bar{\phi}_i^{\text{out}}(c) \doteq c_i^- + \bar{p}_i. \tag{2}$$

Here, $\bar{\phi}_i^{\text{in}}(c)$ represents the total nominal assets of bank *i*—consisting of external inflows and incoming interbank payments—while $\bar{\phi}_i^{\text{out}}(c)$ represents its total nominal liabilities to both external and interbank creditors.

[1] Although $s_{ij} \geq 0$ in standard cases—representing a ***long*** position of bank *i* in asset *j*—the methodology developed here also accommodates ***short*** positions, allowing some entries s_{ij} to be negative. A negative value of s_{ij} indicates that bank *i* has borrowed and sold $|s_{ij}|$ shares of asset *j*, anticipating a price decline and intending to repurchase them later at a lower price.

[2] As discussed in Elsinger et al. (2006), the condition $c_i \geq 0$ can only be guaranteed when liabilities to the external sector are treated as having equal seniority with interbank liabilities—a simplifying assumption adopted by Eisenberg and Noe in Eisenberg and Noe (2001). However, if external liabilities are given priority over interbank claims, it is possible for a bank's obligations to the external sector to exceed the total value of its assets.

Actual payments, clearing vectors and system loss

Under regular operating conditions, the nominal external net value vector $c = \bar{c}$ is such that $\bar{\phi}_i^{\text{in}}(\bar{c}) \geq \bar{\phi}_i^{\text{out}}(\bar{c})$ for all *i*. This indicates that each bank is able to meet its obligations at the end of the period, possibly by liquidating some of its assets.

However, if a financial shock affects certain nodes–such as a decline in the value of external assets below expected levels—then $c_j < \bar{c}_j$ for some *j*, and it may result in an institution *i* being unable to fully meet its payment obligations, as its liabilities exceed its assets, that is, $\bar{\phi}_i^{\text{out}}(\bar{c}) > \bar{\phi}_i^{\text{in}}(\bar{c})$. In this situation of *default*, node *i* pays out according to its available resources: it first satisfies its obligations to the external sector and then reduces its payments to interbank creditors, lowering the nominal amounts $\bar{p}_{ij}$ to actual payments $p_{ij} \leq \bar{p}_{ij}$. As a result, adjacent nodes may also default, since the reduction in incoming payments decreases their asset sides.

Clearing principles in the presence of external shocks. The framework proposed by Eisenberg and Noe (2001) introduces a simple yet powerful scheme for determining this "fair" payment matrix $P \doteq (p_{ij})$, based on three fundamental axioms: (i) limited liability, (ii) absolute priority of debt claims, and (iii) the pro-rata rule. To formalize the axioms, we define the actual asset and liability sides of the banks, given a payment matrix *P* and an external net position vector *c*:

$$\phi^{\text{in}} \doteq c^{+} + P^{\top}1, \quad \phi^{\text{out}} \doteq c^{-} + p, \quad \text{where } p \doteq P1, \tag{3}$$

representing, respectively, the total inflows and outflows at each node. The vector of residual values at the nodes, upon fulfilling external liabilities, is then given by

$$d \doteq \phi^{\text{in}} - c^{-} = c + P^{\top}1. \tag{4}$$

(i) **Limited liability.**
The limited liability rule requires that the total payment of bank *i* to the rest of the network should not exceed its net residual value. Mathematically, this means that bank *i* fulfills the interbank debt claims only if $d_i > 0$, and in any case its total payment p_i cannot exceed d_i. Otherwise (if $d_i \leq 0$), bank *i* is *insolvent* and ceases all payments to other banks: $p_i = 0$. These conditions can be written as

$$0 \leq p_i \leq [d_i]^{+} \quad \forall i \in \mathcal{V}. \tag{5}$$

(ii) **Absolute priority of debt claims.**
Each node *i* either pays its obligations in full ($p_i = \bar{p}_i$) or pays all its value to the creditors ($p_i = [d_i]^{+}$). In view of the limited liability rule, this means that $p_i = \min(\bar{p}_i, [d_i]^{+})$ for every bank $i \in \mathcal{V}$.

(iii) **Pro-rata (proportionality) rule.**
Assuming that debts between the banks have equal seniority, it is natural to assume that the payments p_{ij} from bank *i* to its claimants $j \neq i$ have to be proportional[3] to

[3] The proportional division principle is widely used in bankruptcy and tax legislation, is enforced in many financial contracts, and has been shown to be the only allocation (or division) rule satisfying several key axiomatic properties (Csóka and Herings 2021). Accordingly, we adopt this principle in line with the Eisenberg-Noe framework. However, while the pro-rata rule ensures local fairness in clearing payments between neighboring nodes, relaxing this rule can substantially reduce the system shortfall (Calafiore et al. 2024).

the nominal liabilities $\bar{p}_{ij}$. To write this formally, it is convenient to introduce the stochastic[4] matrix of *relative* liabilities

$$A = (a_{ij}), \quad a_{ij} \doteq \begin{cases} \frac{\bar{p}_{ij}}{\bar{p}_i}, & \bar{p}_i > 0, \\ 1, & \bar{p}_i = 0 \text{ and } i = j, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

The pro-rata (equal priority, proportionality) rule can then be formulated as

$$p_{ij} = p_i a_{ij}, \quad \forall i, j \in \mathcal{V}, \tag{7}$$

or, in matrix format, as $P = \text{ diag } (P1)A$.

Clearing vectors and incurred system loss. The three rules (i)–(iii) above can be written as a single nonlinear equation on the vector $p = P1$: by virtue of the pro-rata rule (7), $d = c + P^\top 1 = c + A^\top p$, whence[5]

$$p = \min(\bar{p}, [c + A^\top p]^+) = \left[\min(\bar{p}, c + A^\top p)\right]^+. \tag{8}$$

Definition 1 A vector p satisfying (8) is said to be a *clearing* vector for the financial network $(\mathcal{V}, \mathcal{E}, \bar{P})$ with vector of net incoming values c.
Notice that if the actual payments $p_{ij} = a_{ij} p_i$ are determined by a clearing vector p, then the network's shortfall (the total reduction of interbank claims) is found as

$$L(p) \doteq \sum\nolimits_{i \in \mathcal{V}} (\bar{p}_i - p_i) = 1^\top (\bar{p} - p). \tag{9}$$

Definition 2 We call $L(p)$ the overall *system loss* corresponding to clearing vector p.
In the *generic* case, the clearing vector is known to be unique—this holds, for instance, when the network is strongly connected (a condition referred to as "regularity" in Acemoglu et al. (2015)) and $\sum_{i \in \mathcal{V}} c_i \neq 0$. When $c \geq 0$, uniqueness depends solely on the structural properties of the network (Calafiore et al. 2024; Csóka and Herings 2024). Necessary and sufficient conditions for uniqueness in the general signed case are provided in Massai et al. (2022). Even in cases where the clearing vector is not unique, a *maximal* clearing vector $p^*(c)$ always exists–one that minimizes the system loss function (9) and can be computed efficiently. Its formal characterization is presented in the next section, which introduces the key method used to analyze worst-case shock scenarios in the subsequent analysis.

Optimal clearing vectors and loss minimization

In this subsection, we present three propositions concerning the optimal (maximal) clearing vector, which is optimal in the sense that it minimizes the system loss. The proofs are deferred to the Appendix.

[4] By definition, matrix A is *stochastic* if $a_{ij} \geq 0$ and $\sum_j a_{ij} = 1$ for all i or, equivalently, $A1 = 1$.

[5] Equation (8) is a special case of the so-called "payment equilibrium" equation from (Acemoglu et al. 2015).

Proposition 1 The set of clearing vectors is non-empty for any vector of net incoming values $c \in \mathbb{R}$. Moreover, there exist clearing vectors $p_*(c)$ and $p^*(c)$ that are, respectively, the componentwise *minimal* and *maximal* elements of this set, in the sense that any other clearing vector p satisfies $p_*(c) \leq p \leq p^*(c)$.
Since the function $L(p)$ defined in (9) is monotone decreasing[6] in p, the maximal clearing vector $p^*(c)$, obviously, provides the minimal overall system loss among all possible clearing vectors. Therefore, we assume henceforth that financial institutions adopt the *maximal* clearing vector $p^*(c)$ and restrict our analysis to its properties.

The following proposition states that $p^*(c)$ can be computed by solving a linear programming problem, except in cases where severe drops in asset values make it impossible to avoid insolvency at certain nodes, that is, the equity of some financial institutions becomes negative even if they fully suspend their liabilities.

Proposition 2 For given $c \in \mathbb{R}^n$, nominal liability matrix $\bar{P} \in \mathbb{R}^{n,n}$, and stochastic pro-rata matrix A given in (6), consider the linear program:

$$\begin{aligned} \eta^* = \min_{p \in \mathbb{R}^n} \quad & 1^\top (\bar{p} - p) \\ \text{s.t.:} \quad & \bar{p} \geq p \geq 0 \\ & c + A^\top p \geq p. \end{aligned} \tag{10}$$

If this problem is feasible, then its optimal solution p^* is unique and coincides with the maximal clearing vector $p^* = p^*(c)$. Moreover, under p^*, all liabilities towards the external sector are paid in full, i.e., vector d in (4) is nonnegative.

Contrary, if problem (10) is infeasible, then for any possible payment vector $0 \leq p \leq \bar{p}$ (including the clearing vector defined via (8)) there will exist nodes which are insolvent to the external sector, that is, $d_i = c_i + (A^\top p)_i < 0$ and $p_i = 0$ for some i.
A proof of Proposition 2 is given in Appendix B.

Observe that the optimal value η^* of (10) is nonnegative; it equals zero when no defaults occur and is positive otherwise, representing the total system-wide loss due to defaults. We adopt the convention of setting $\eta^* = \infty$ when problem (10) is infeasible– that is, when no payment vector p satisfying $0 \leq p \leq \bar{p}$ ensures nonnegative equity for all participating banks. Moreover, the function $\eta^* = \eta^*(c)$ is *convex* over the set $\mathcal{C}_{feas}$ of vectors $c \in \mathbb{R}^n$ for which Problem (10) is feasible, as established in the following proposition.

Proposition 3 Let $\mathcal{C}_{feas}$ be the set of all vectors $c \in \mathbb{R}^n$ for which the inequalities in (10) have at least one feasible solution. The maximal clearing vector p^* is an element wise non-decreasing and concave function on $\mathcal{C}_{feas}$. Symmetrically, the optimal value η^* in (10) is non-increasing and convex over $\mathcal{C}_{feas}$.
A proof of Proposition 3 is given in Appendix C. Notice that the set $\mathcal{C}_{feas}$ is a convex closed polyhedral set, which is unbounded and contains all vectors $c \geq 0$. However, if c has negative components, it may happen that $\mathcal{C}_{\text{feas}}$ is empty. In this case, some banks– despite being fully relieved of interbank payment obligations (i.e., $p_i = 0$) – still cannot meet their external liabilities without incurring negative equity, regardless of how an admissible payment vector $p \in [0, \bar{p}]$ is chosen.

[6] This means that $\tilde{L}(p') \geq \tilde{L}(p'')$ whenever $p', p'' \in [0, \bar{p}]$, $p' \leq p''$ and $p' \neq p''$.

An illustrative example. Figure 1 shows a small example with $n = 4$ nodes, mutual liabilities $\bar{p}_{12} = 1$, $\bar{p}_{14} = 2$, $\bar{p}_{24} = 4$, $\bar{p}_{31} = 1$, $\bar{p}_{32} = 1$, $\bar{p}_{43} = 6$, exposures in external assets $s_{11} = 1$, $s_{21} = 2$ (only one external asset in this example), and nominal value of the external asset $v = v_1 = 2.2$. Assuming $c^e = 0$ in (1), we have $c = Sv$, and solving problem (10) we have that no default arises, whence $\eta^*(2.2) = 0$, and $p^* = \bar{p} = [3,\ 4,\ 2,\ 6]^\top$. However, if the value of the external asset drops to $v = 1.9$, then $\eta^*(1.9) = 0.1667$ and the clearing vector is $p^* = [2.9,\ 4.0,\ 2.0,\ 5.93]^\top \leq \bar{p}^\top$, showing that nodes 1 and 4 default. If the value of the external asset drops to $v = 1.5$ we have $\eta^*(1.5) = 0.8333$ and $p^* = [2.5,\ 4.0,\ 2.0,\ 5.6667]^\top \leq \bar{p}^\top$, still with nodes 1 and 4 in default. If ν drops further to $v = 1$, then $\eta^*(1) = 2.333$ and $p^* = [2.0,\ 3.6667,\ 2.0,\ 5.0]^\top \leq \bar{p}^\top$, where now also node 2 defaults. The main contribution of this paper is to provide an efficient computational method for determining the threshold value on the amplitude of ν (in the more general case in which ν is a vector, rather than a scalar) that guarantees a no-default operation of the financial network, and for determining the loss curve of η^* versus the perturbation level, once such threshold is surpassed.

In this example, there are no insolvencies with respect to the external sector, as all banks have zero net positions toward it. If, on the other hand, one assumes that $c_1^e = -2$, meaning that the first node owes 2 units to the external sector, then this debt clearly cannot be repaid when $v \geq 1$. If the asset price ν falls below 1, then bank 1 cannot fulfill its obligation to the external sector–even if all interbank payments are suspended, since its inflow ϕ_1^{in} obviously does not exceed $s_{11}v + \bar{p}_{31} = 1 + v$.

Main results

We henceforth assume that the net external position vector c consists of a nominal component $\bar{c}$ and a variable component, caused by the price fluctuations, that is,

$$c = \bar{c} + S\delta, \quad \bar{c} \doteq \bar{c}^e + S\bar{v}, \quad v = \bar{v} + \delta \tag{11}$$

where $\bar{c}^e$, $\bar{v}$ denote the nominal values of the net liquidity incoming from the external sector and of the asset prices, respectively, and $\delta \in \mathbb{R}^m$ is a vector of price fluctuations. We recall that, in nominal conditions, the book values of the asset and liability sides of the balance sheets for all banks (2) can be written as

$$\bar{\phi}^{\text{in}} = \bar{c}^+ + A^\top \bar{p}, \quad \bar{\phi}^{\text{out}} = \bar{c}^- + \bar{p}. \tag{12}$$

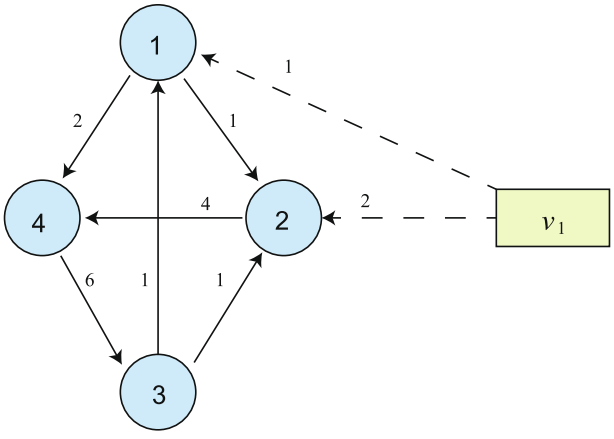

Fig. 1 An illustrative example with $n = 4$ nodes and one external asset

We consider the following assumption.

Assumption 1 (No defaults on interbank or external obligations under nominal conditions.) A, $\bar{p}$, S, $\bar{c}^e$, $\bar{v}$ are such that

$$\bar{c} + A^\top \bar{p} > \bar{p}. \tag{13}$$

That is, under nominal liabilities, nominal net external inflows, and nominal asset prices, all banks can fulfill their interbank and external obligations. ⋆

The default resilience margin

The first question we pose is how much fluctuation in the assets' prices the system can withstand before a default is triggered in some bank.

Definition 3 (Default resilience margin) Under Assumption 1, the *default resilience margin* of the banking system with respect to the norm $\|\cdot\|$ on the space of price fluctuations $\{\delta\}$ is defined as the maximum value ϵ^* of $\epsilon > 0$ such that

$$\bar{c} + S\delta + A^\top \bar{p} \geq \bar{p}, \quad \forall \delta : \|\delta\| \leq \epsilon. \tag{14}$$

In other words, the resilience margin is the maximum joint price perturbation amplitude that guarantees the system will remain default-free in the worst case.

Computing the resilience margin

The resilience margin is found by solving

$$\begin{aligned} &\epsilon^* = \max \epsilon \\ &\text{s.t.: } \bar{c} + S\delta + A^\top \bar{p} \geq \bar{p}, \, \forall \delta : \|\delta\| \leq \epsilon. \end{aligned} \tag{15}$$

Due to symmetry in $\|\delta\| \leq \epsilon$ we can replace $\delta = -\tilde{\delta}$ and rewrite the requirement as

$$S\tilde{\delta} \leq \bar{r} \doteq \bar{c} + (A^\top - I)\bar{p}, \quad \forall \tilde{\delta} : \|\tilde{\delta}\| \leq \epsilon, \tag{16}$$

which can be rewritten as

$$\max_{\tilde{\delta}:\|\tilde{\delta}\|\leq\epsilon} \sigma_i^\top \tilde{\delta} \leq \bar{r}_i \quad \forall i \in \mathcal{V} \tag{17}$$

where $\sigma_i^\top$ denotes the ith row of S, and we notice that $\bar{r}_i$ is the nominal net worth of bank i, in the absence of price perturbations. The actual solution now depends on the specific choice of the norm. Introducing the *dual* norm

$$\|\sigma\|_* \doteq \max_{\|\delta\|\leq 1} \sigma^\top \delta, \tag{18}$$

the condition (17) can be reformulated in the equivalent form

$$\epsilon \|\sigma_i\|_* \leq \bar{r}_i \quad \forall i \in \mathcal{V}. \tag{19}$$

This leads to our first result, offering a simple formula for the default resilience margin.

Theorem 1 The default resilience margin with respect to the norm $\|\cdot\|$ is found as

$$\epsilon^* \doteq \min_{i\in\mathcal{V}} \frac{\bar{r}_i}{\|\sigma_i\|_*}, \tag{20}$$

where vector $\bar{r}$ is defined in (16). For any price perturbation δ such that $\|\delta\| \leq \epsilon^*$ the financial system remains default free. There exists, however, a *worst-case* perturbation with $\|\delta\| = \varepsilon^*$, which brings the balance of some bank to zero.

The proof of Theorem 1 proceeds straightforwardly. Define ε^* as in (20). Then, inequalities (19) (and consequently, conditions (16) and (17)) become equivalent to the condition $\varepsilon \leq \varepsilon^*$. Let $\mathcal{I}^* \subseteq \mathcal{V}$ denote the set of indices on which the minimum in (20) is attained. By selecting $i_0 \in \mathcal{I}^*$, we can select a perturbation $\delta = -\tilde{\delta}_*$, where

$$\delta_* \doteq \varepsilon^* \arg\max\{\sigma_{i_0}^\top v : v \in \mathbb{R}^n, \|v\| = 1\}. \tag{21}$$

This choice effectively nullifies the balance of bank i_0, transforming the inequality (17) with $i = i_0$ into an equality.⋆

Two relevant special cases

Although the dual norm can easily be computed for any ℓ_p norm on $\mathbb{R}^n$, in this work we are primarily interested in two cases that have a clear financial significance, namely the ℓ_∞ and the ℓ_1 norm. The ℓ_∞ constraint $\|\delta\|_\infty \leq \epsilon$ implies independent, entry-wise bounds on the variation of each asset price, that is $|\delta_i| \leq \epsilon$, for $i = 1, \ldots, m$. The ℓ_1 constraint $\|\delta\|_1 \leq \epsilon$ implies instead a constraint on the sum of absolute price perturbations, that is $\sum_i |\delta_i| \leq \epsilon$.

ℓ_∞ case

For the ℓ_∞ norm $\|\cdot\| = \|\cdot\|_\infty$ the dual norm is $\|\cdot\|_* = \|\cdot\|_1$, and

$$\max_{\tilde{\delta}:\|\tilde{\delta}\|_\infty \leq \epsilon} \sigma_i^\top \tilde{\delta} = \epsilon \|\sigma_i\|_1, \tag{22}$$

which maximum is attained for $\tilde{\delta}_j^* = \epsilon \cdot \mathrm{sign}(S_{ij})$, $j = 1, \ldots, m$. The default resilience margin is thus

$$\epsilon_\infty^* = \min_{i=1,\ldots,n} \frac{\bar{r}_i}{\|\sigma_i\|_1}. \tag{23}$$

and, correspondingly, a worst-case perturbation

$$\delta^{(\infty)} = -\epsilon_\infty^* \cdot \mathrm{sign}(\sigma_o), \tag{24}$$

with $o \in \mathcal{I}^*$, where $\mathcal{I}^*$ denotes the set of indices $i \in \{1, \ldots, n\}$ for which the minimum in (23) is attained. The banks with $i \in \mathcal{I}^*$ are called *primary defaulters* since their balance sheets are brought to zero (hence to the brink of default) by a critical movement $\delta^{(\infty)}$ of the asset prices of amplitude ϵ_∞^*. Should the value of asset prices move beyond $\delta^{(\infty)}$ (positively, if $\delta_j^{(\infty)} > 0$, or negatively if $\delta_j^{(\infty)} < 0$) then all primary defaulters will actually default and, as a consequence of these defaults, they may trigger other *secondary defaults* in banks whose balance sheets become negative due to reduced income from primary defaulters, and so on. Conversely, for any price perturbation δ such that $\|\delta\|_\infty \leq \epsilon_\infty^*$ the system is guaranteed to remain default-free.

We observe that the worst-case ℓ_∞ price perturbations in (24) represent a rather pessimistic situation in which all asset prices simultaneously drop (or rise, depending on the sign of the corresponding entry in σ_o) by the maximum margin ϵ^*_∞. This aspect is mitigated by considering bounds on the joint variation of all asset prices, as captured by the ℓ_1 norm of the price perturbations, as discussed next.

ℓ_1 case

For the ℓ_1 norm case, the dual norm is the ℓ_∞ norm, and

$$\max_{\tilde{\delta}:\|\tilde{\delta}\|_1\le\epsilon} \sigma_i^\top \tilde{\delta} = \epsilon\|\sigma_i\|_\infty. \tag{25}$$

Letting $\mathcal{J}_i^* \doteq \arg\max_{j=1,\dots,m} |S_{ij}|$, the above maximum is attained for perturbations $\tilde{\delta}$ such that all entries are zero except for those with indices $j \in \mathcal{J}_i^*$, which take value $\pm\epsilon/|\mathcal{J}_i^*|$.

The default resilience margin is given by

$$\epsilon_1^* = \min_{i=1,\dots,n} \frac{\bar{r}_i}{\|\sigma_i\|_\infty}. \tag{26}$$

Denoting again with $\mathcal{I}^*$ the set of indices $i \in \{1,\dots,n\}$ for which the minimum in (26) is attained, by taking any $i \in \mathcal{I}^*$ we obtain a corresponding worst-case price perturbation scenario in the ℓ_1 case as a vector $\delta^{(1)}$ in which all entries are zero, except for those in positions $j \in \mathcal{J}_i^*$, which take value

$$\delta_j^{(1)} = -\frac{\epsilon_1^*}{|\mathcal{J}_i^*|}\text{sign}(S_{ij}), \quad \forall j \in \mathcal{J}_i^*. \tag{27}$$

Observe that in common situations the optimal index sets $\mathcal{I}^*$ and $\mathcal{J}_i^*$, $i \in \mathcal{I}^*$, will contain just one element, hence in such cases the optimal $\delta^{(1)}$ perturbation will consist of only one nonzero entry, and $\mathcal{I}^*$, $\mathcal{J}_i^*$ identify the most critical bank and the most critical asset, respectively.

In the next section we consider the situation in which the amplitude ϵ of the perturbation goes beyond the default resilience threshold ϵ^*, and hence defaults may appear. In such case, we are interested in determining the worst-case impact of the cascaded defaults on the system, i.e., in computing the worst-case loss due to defaults.

The worst-case loss curve

Consider problem (10) with $c = \bar{c} + S\delta$: its optimal value $\eta^* = \eta^*(\bar{c} + S\delta)$ is a function of the price perturbation δ, and we know that under Assumption 1 we shall have $\eta^*(\bar{c} + S\delta) = 0$ for all δ such that $\|\delta\| \le \epsilon^*$, where ϵ^* is the default resilience margin relative to the considered norm. As $\|\delta\|$ is allowed to go beyond the ϵ^* level, we shall have $\eta^*(\bar{c} + S\delta) \ge 0$ since defaults will be triggered. We are here interested in computing the worst-case value of the system-wide financial loss $\eta^*(\bar{c} + S\delta)$ that can occur when $\|\delta\| \le \epsilon$, for ϵ possibly larger than the resilience margin. This is formalized as the following max-min problem

$$\begin{aligned} \eta_{\text{wc}} = \max_{\|\delta\|\le\epsilon} \eta^*(\bar{c} + S\delta) = \max_{\|\delta\|\le\epsilon} \min_{p\in\mathbb{R}^n} \quad & 1^\top(\bar{p} - p) \\ \text{s.t.:} \quad & \bar{p} \ge p \ge 0 \\ & \bar{c} + S\delta + A^\top p \ge p. \end{aligned} \tag{28}$$

Its optimal value η_{wc} quantifies the worst-case systemic impact of asset price variations in the range $\|\delta\| \leq \epsilon$. The following key result holds.

Theorem 2 The optimal value η_{wc} of (28) can be computed by solving the optimization problem

$$\begin{array}{ll} \eta_{\text{wc}} = \max\limits_{\beta,\lambda\geq 0} & (1-\beta)^\top \bar{p} - \bar{c}^\top \lambda + \epsilon \|S^\top \lambda\|_* \\ \text{s.t.:} & \beta - 1 + (I - A)\lambda \geq 0 \end{array} \tag{29}$$

where $\|\cdot\|_*$ is the dual of norm $\|\cdot\|$.

A proof of this theorem is given in Section D of the Appendix. Observe that, in general, this formulation can be challenging to solve numerically, as it involves the *maximization* of a convex function over a polyhedron, such problems are generally NP-hard. In the next subsections, we show that the above problem actually admits an efficient formulation in the cases where the price perturbation is bounded by either the ℓ_∞ or the ℓ_1 norm.

Worst-case loss for ℓ_1-norm price variations

Consider first the case where the δ perturbation is measured via the ℓ_1 norm, $\|\delta\|_1$. In this case, the dual norm is the infinity norm, that is

$$\|S^\top \lambda\|_* = \|S^\top \lambda\|_\infty = \max_{i=1,\dots,m} |\zeta_i^\top \lambda| = \max_{i=1,\dots,m} |\zeta_i|^\top \lambda,$$

where $\zeta_i^\top$, $i = 1, \dots, m$, are the rows of $S^\top$, and since $\lambda \geq 0$ it follows that $|\zeta_i^\top \lambda| = |\zeta_i|^\top \lambda$. Problem (29) then becomes

$$\begin{array}{ll} \eta_{\text{wc}} = \max\limits_{\beta,\lambda\geq 0} & (1-\beta)^\top \bar{p} - \bar{c}^\top \lambda + \epsilon \max\limits_{i=1,\dots,m} |\zeta_i|^\top \lambda \\ \text{s.t.:} & \beta - 1 + (I - A)\lambda \geq 0. \end{array} \tag{30}$$

The latter optimization problem can be rewritten as follows:

$$\begin{array}{lll} \eta_{\text{wc}} = \max\limits_{i=1,\dots,m} \max\limits_{\beta,\lambda\geq 0} & (1-\beta)^\top \bar{p} - \bar{c}^\top \lambda + \epsilon |\zeta_i|^\top \lambda \\ \text{s.t.:} & \beta - 1 + (I - A)\lambda \geq 0. \end{array} \tag{31}$$

This means that we can compute the worst-case system loss efficiently and globally by solving m linear programs, each of which amounts to solving the inner maximization in (31) with respect to β, λ, for a fixed ζ_i. Little further elaboration will give us also a worst-case price variation vector, and the clearing vector p which is worst-case optimal. Let λ^* be an optimal λ vector for the above problem, then an ℓ_1 worst-case price perturbation can be found as $\delta_i^{(1)}$ such that

$$\delta_i^{(1)} = \begin{cases} -\epsilon \cdot \text{sign}(\zeta_i^\top \lambda^*) & \text{if } i = i^* \\ 0 & \text{otherwise,} \end{cases} \quad i = 1, \dots, m, \tag{32}$$

where i^* is any index for which $|\zeta_i^\top \lambda^*|$ is maximum.

An interesting question arises about the uniqueness of the worst-case price perturbation $\delta^{(1)}$. This may be relevant, since the worst-case perturbation identifies the subset of assets whose price perturbation is the most critical and, correspondingly, the set of

banks who will default due to such critical price fluctuation. Leveraging a known result of Mangasarian (1979), we provide in Proposition 4 an easily computable criterion for checking whether the worst-case perturbation is unique.

Worst-case loss for ℓ_∞-norm price variations

When the δ perturbation is measured via the ℓ_∞ norm, $\|\delta\|_\infty$, we have

$$\|S^\top \lambda\|_* = \|S^\top \lambda\|_1 = \sum_{i=1}^{m} |\zeta_i^\top \lambda| = \sum_{i=1}^{m} |\zeta_i|^\top \lambda = 1^\top |S|^\top \lambda \tag{33}$$

Problem (29) then becomes

$$\begin{aligned} \eta_{\text{wc}} = \max_{\beta,\lambda\geq 0} \quad & (1-\beta)^\top \bar{p} - \bar{c}^\top \lambda + \epsilon 1^\top |S|^\top \lambda \\ \text{s.t.:} \quad & \beta - 1 + (I - A)\lambda \geq 0. \end{aligned} \tag{34}$$

That is, in this case we can find exactly and efficiently the worst-case loss by solving a single LP. A worst-case price perturbation can then be found as $\delta^{(\infty)}$ such that

$$\delta_i^{(\infty)} = -\epsilon \cdot \text{sign}(\zeta_i^\top \lambda), \quad i = 1, \ldots, m. \tag{35}$$

An easily computable criterion for checking whether such worst-case perturbation is unique is provided by the following proposition.

Proposition 4 (Uniqueness of the worst-case perturbation scenario) For the ℓ_1-norm case: the worst-case perturbation scenario in eq. (32) is unique if the conditions hold:

- The max in Eq. (31) is attained at a single index *i*;
- The corresponding maximization problem in (β, λ) has a unique solution (β^*, λ^*) (a fact that can be checked by applying Proposition 5 in Appendix F);
- $\max_i |\zeta_i^\top \lambda^*|$ is attained at a single index *i*.

Similarily, for the ℓ_∞-norm case: the worst-case perturbation scenario in Eq. (35) is unique if the following two conditions hold:

- The maximization problem (34) has a unique solution (β^*, λ^*) (a fact that can be checked by applying Proposition 5 in Appendix F);
- $\zeta_i^\top \lambda^*$ is nonzero for all *i*.

The insolvency resilience margin

From (29), we observe that the optimal worst-case loss value η_{wc} is a function of the perturbation level ϵ. As established in Section "The default resilience margin", we have $\eta_{\text{wc}}(\epsilon) = 0$ for all $\epsilon \leq \epsilon^*$, where ϵ^* denotes the default resilience threshold under the chosen norm. Beyond this threshold, $\eta_{\text{wc}}(\epsilon)$ begins to increase as ϵ grows. However, once the perturbation set $\delta : |\delta| \leq \epsilon$ becomes sufficiently large, it may include some vector δ for which Problem (28) becomes infeasible, i.e., $\bar{c} + S\delta \notin \mathcal{C}_{\text{feas}}$. In this section, we compute the largest perturbation level ϵ such that Problem (28) remains feasible for all δ with $|\delta| \leq \epsilon$, called the *insolvency resilience margin*.

Definition 4 (Insolvency resilience margin) The *insolvency resilience margin* $\epsilon_{\rm ub}$ of the banking system with respect to the norm $\|\cdot\|$ is the largest value of $\epsilon \geq 0$ such that the set $\{p : \bar{p} \geq p \geq 0,\ \bar{c} + S\delta + A^\top p \geq p\}$ remains nonempty for all $\delta : \|\delta\| \leq \epsilon$.
The loss curve of $\eta_{\rm wc}$ as a function of ϵ is thus properly defined in the range $[0, \epsilon_{\rm ub}]$. Notice that η_{wc}, obviously, is non-decreasing and, furthermore, it can be easily derived from Proposition 3

that $\eta_{wc}(\epsilon)$ is convex on $[0, \epsilon_{\rm ub}]$. The insolvency resilience margin $\epsilon_{\rm ub}$ can be computed according to the following theorem.

Theorem 3 $\epsilon_{\rm ub}$ is given by the optimal value of the following LP:

$$\begin{aligned} \epsilon_{\rm ub} = \max_{p,\epsilon} \quad & \epsilon \\ \text{s.t.:} \quad & 0 \leq p \leq \bar{p} \\ & \bar{c} - \epsilon s + A^\top p \geq p, \end{aligned}$$

where s is a vector such that $s_i \doteq \|\sigma_i^\top\|_*$, $i = 1, \ldots, n$.
Theorem 3 is proven in Section E of the Appendix.

Technically, for $\epsilon > \epsilon_{\rm ub}$, there exist perturbations δ with $|\delta| \leq \epsilon$ such that the constraint $(\bar{c} + S\delta) + A^\top p \geq p$ in Problem (28) cannot be satisfied by any $p \geq 0$. In such cases, the clearing vector can no longer be characterized as the optimal solution of (10). As discussed in Section "Optimal clearing vectors and loss minimization", this situation corresponds to the inability to clear interbank liabilities without revealing the insolvency of certain banks, whose obligations to the external sector exceed the value of their assets.

Dealing with the case $\epsilon > \epsilon_{\rm ub}$ would require a formulation incorporating constraints of the form

$$p \leq \max\left(\bar{c} + S\delta + A^\top p, 0\right), \tag{36}$$

which would introduce nonlinearity and destroy the convexity of the problem. In this case, the extremal problem for determining the maximal clearing vector becomes a mixed-integer linear program (MILP), rather than a linear program; see the relevant formulation in Ararat and Meimanjan (2023). Since our method for computing the worst-case loss relies on linear programming duality, it does not extend to MILPs. Consequently, evaluating the worst-case loss for $\epsilon > \epsilon_{\rm ub}$ lies beyond the scope of this manuscript.

Numerical examples

To illustrate the proposed approach numerically, we present two examples. The first involves a small network with 8 nodes, allowing for transparent analysis and interpretation. The second considers a larger network with a core–periphery structure, a topology commonly used to represent financial networks (Hu et al. 2024; Craig and von Peter 2014; Jie and Ma 2025).

Example 1 First, we considered a schematic network with $n = 8$ nodes and $m = 4$ external assets, as shown in Fig. 2, where the numbers on the arrows represent the mutual liabilities among the banks. We assumed the vector of external inflows c^+, external outflows c^-, and the matrix of asset shares S to be as follows:

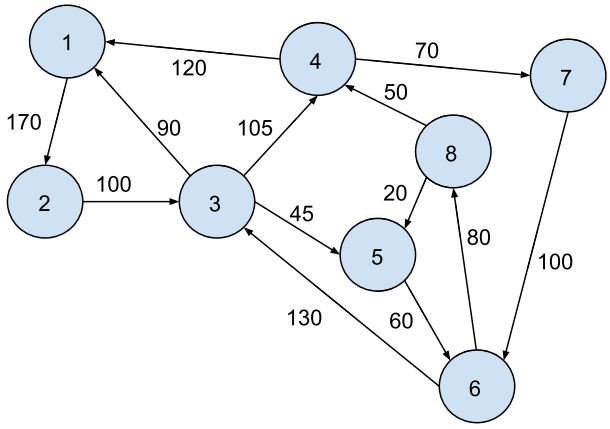

Fig. 2 A schematic network with 8 nodes

$$c^+ = \begin{bmatrix} 158 \\ 38 \\ 15 \\ 285 \\ 25 \\ 180 \\ 60 \\ 55 \end{bmatrix}, \quad c^- = \begin{bmatrix} 180 \\ 100 \\ 0 \\ 145 \\ 0 \\ 50 \\ 20 \\ 60 \end{bmatrix}, \quad S = \begin{bmatrix} 96 & 29 & 99 & 0 \\ 53 & 13 & 75 & 85 \\ 0 & 28 & 0 & 57 \\ 32 & 79 & 0 & 0 \\ 0 & 0 & 21 & 0 \\ 0 & 45 & 27 & 71 \\ 0 & 0 & 49 & 0 \\ 75 & 79 & 5 & 41 \end{bmatrix}.$$

For this network, we can compute the resilience margins ϵ^*_∞ and ϵ^*_1, as given in (23) and in (26) for the ℓ_∞ case and for the ℓ_1 case, respectively; we also computed the perturbation upper limit $\epsilon_{\rm ub}$ as shown in Theorem 3, obtaining

$$\epsilon^*_\infty = 0.0249, \quad \epsilon^*_1 = 0.0630, \quad \epsilon_{ub,\infty} = 0.1878, \quad \epsilon_{ub,1} = 0.4328.$$

We next selected 20 evenly spaced values of ϵ inside the interval $[\epsilon^*, \epsilon_{\rm ub}]$ and for each of them we computed the worst-case loss as shown in Section "Worst-case loss for ℓ_1-norm price variations" for the ℓ_1 case and in Section "Worst-case loss for ℓ_∞-norm price variations" for the ℓ_∞ case. For each value of ϵ, we also modeled a financial shock by generating a random vector of price fluctuations $\delta < 0$ such as $\|\delta\| = \epsilon$ and we computed the corresponding loss η^* by solving Problem (10).

In order to compare the worst possible losses to those produced by random shocks, we performed 1000 runs for each value of ϵ and in each run we generated a different random scenario δ. Figure 3a and Fig. 3b show the obtained results for the ℓ_1 case and the ℓ_∞ case, respectively. The red curve represents the average loss due to random shocks, and the red band represents, for each value of ϵ, the interval between the minimum and maximum losses obtained in the random numerical simulation.

It is interesting to observe that the worst-case loss curve presents some "kinks", i.e., there are ϵ levels at which the slope of the curve changes abruptly. Intuitively, this feature is due to the change in the group of defaulting banks as the level of perturbation increases. In (32) we have shown that the worst-case price perturbation δ^* in the ℓ_1 case is a vector where the only non zero component corresponds to the asset i^*, defined as the one that maximizes $|\zeta_i^\top \lambda|$. The asset i^* is thus the one whose variation causes the worst-case impact. Figure 4 shows the value of i^* as function of the perturbation radius ϵ. By comparing Figs. 3a and 4, we can observe that the kinks observed in Fig. 3a are in correspondence of the changes of values of i^*.

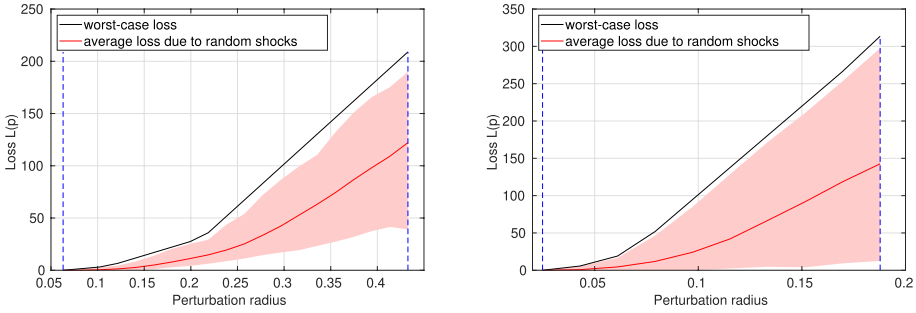

(a) ℓ_1-case (b) ℓ_∞-case

Fig. 3 Worst-case losses vs. losses under random shocks, ℓ_1 and ℓ_∞ norms. The dashed lines mark the default resilience margin ϵ^* and the insolvency resilience margin ϵ_{ub}

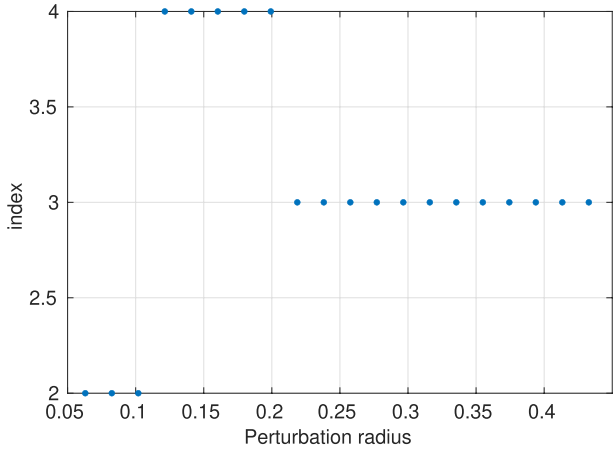

Fig. 4 Index i^* of the worst-case case perturbation δ^* in the ℓ_1 case. Using Proposition 4, we have checked that δ^* is unique for all $\epsilon > \epsilon^*$

Example 2 We now consider a core–periphery random graph with $n = 353$ nodes and $m = 5$ external assets, designed to mimic the structure of a financial network presented in Hu et al. (2024) (the details are given in Appendix G). Similarly to the previous example, we select 10 evenly spaced values of ϵ inside the interval $[\epsilon^*, \epsilon_{\mathrm{ub}}]$ and for each of them we compute the worst-case loss for the ℓ_1 case and the ℓ_∞ case. Then, for each value of ϵ, we can model a financial shocks by generating a random vector of price fluctuation $\delta < 0$ such as $\|\delta\| = \epsilon$ and compute the corresponding loss η^* by solving (10). We perform 1000 runs for each value of ϵ, generating a different random scenario δ for each run. Figures 5a and 5b show the obtained results for the ℓ_1 case and the ℓ_∞ case, respectively. Figure 6 shows the value of i^* of the worst-case case perturbation δ^* in the ℓ_1 case as a function of the perturbation radius ϵ.

Discussion and conclusions

We observe that two distinct but interrelated network effects contribute to financial contagion. First, the network of mutual liabilities creates direct contagion channels: when a bank defaults, it may fail to fully repay its obligations to a creditor bank, potentially triggering a chain reaction of defaults as each affected bank passes on reduced payments to its counterparties. Second, a separate network of common asset exposures links banks indirectly: when the market value of an external asset drops, all banks holding long

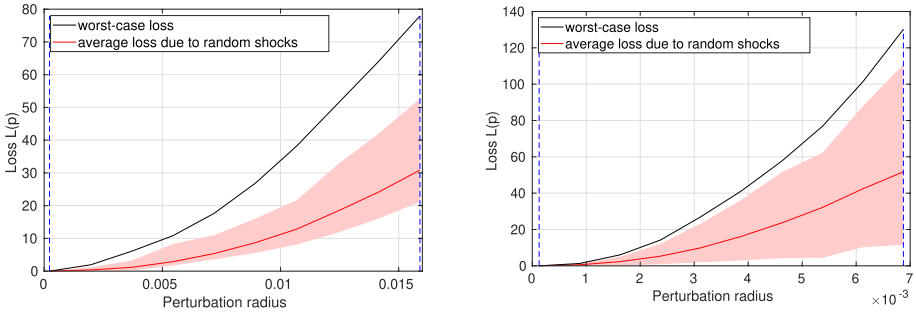

(a) ℓ_1-case (b) ℓ_∞-case

Fig. 5 Core-periphery random graph: Worst-case losses vs. losses under random shocks for the ℓ_1 and ℓ_∞ norms. The dashed lines mark the default resilience margin ϵ^* and the insolvency resilience margin ϵ_{ub}

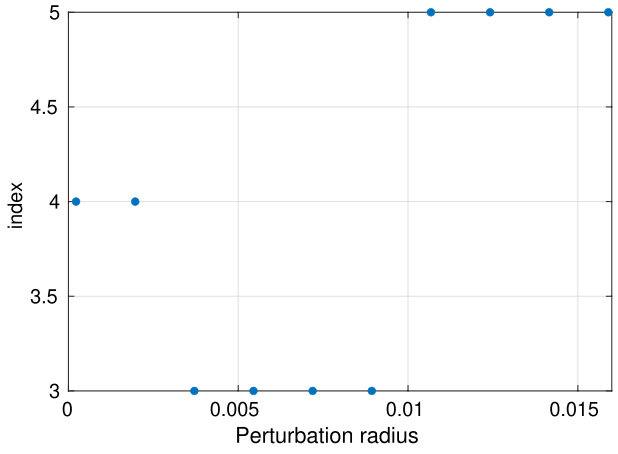

Fig. 6 Core-periphery random graph: index i^* of the worst-case perturbation δ^* in the ℓ_1 case. Using Proposition 4, we have checked that δ^* is unique for all $\epsilon > \epsilon^*$

positions in that asset (i.e., those with $s_{ij} > 0$) simultaneously suffer losses proportional to their exposure levels. As a result, even a shock to a single asset can reduce multiple entries in the net external position vector c. A similar mechanism applies in the case of price increases for banks with short exposures.

Based on an extension of the classical financial network model of Eisenberg and Noe (2001); Elsinger et al. (2006)—incorporating common exposures to a set of external assets—this paper focused on two central questions: (i) what is the largest asset price perturbation the network can withstand without triggering defaults, and (ii) what is the worst-case systemic loss incurred when perturbations exceed that threshold? To this end, and allowing for sign-unrestricted external cash flows c, we first characterized the maximal clearing vector through a linear program (Proposition 2). In Section "The default resilience margin", we introduced the concept of the *default resilience margin*–the maximum amplitude of asset price fluctuations the system can absorb without triggering any defaults. We showed how to compute this threshold under both the ℓ_∞ norm (maximum individual fluctuation) and the ℓ_1 norm (total fluctuation magnitude), providing a rigorous analytical measure of the network's *robustness* to external shocks. Section "The worst-case loss curve" explored the scenario where the perturbation exceeds the default resilience margin, leading to defaults. In this setting, we evaluated the worst-case systemic loss and demonstrated that it can be computed efficiently–either

by solving a single linear program in the ℓ_∞ case or a sequence of m linear programs in the ℓ_1 case–along with identifying the corresponding worst-case shock scenario. To visualize the system's vulnerability against shocks, we proposed the construction of the *worst-case loss curve* $\eta_{\rm wc}(\epsilon)$, which maps the magnitude of asset price perturbation ϵ to the corresponding worst-case systemic loss.

Our analysis applies to perturbations up to an upper limit $\epsilon_{\rm ub}$, which is defined and characterized in Section "The insolvency resilience margin". As discussed in that section, shocks with magnitude $\epsilon > \epsilon_{\rm ub}$ may result in the insolvency of some banks, preventing the determination of the maximal clearing vector through a linear program. Moreover, from a financial modeling perspective, scenarios involving insolvency with respect to the external sector should incorporate procedural aspects of firm liquidation and bankruptcy costs (Weber and Weske 2017). The computation of the worst-case losses in financial networks with defaults and bankruptcy costs (Banerjee and Feinstein 2022; Ararat and Meimanjan 2023) requires methodological instruments that are beyond the scope of the present paper and remains an important topic for future research.

Another challenging problem is computing the worst-case loss curve under more general clearing mechanisms—such as those designed for financial markets with multiple clearing centers or central counterparties (CCPs) Benos et al. (2024); Veraart and Aldasoro (2025); Tang and Fan (2025), as well as net liability clearing mechanisms (Ma et al. 2024), which consolidate all bilateral obligations between pairs of financial institutions into net positions.

Also, our model assumes the external assets to be liquid. Financial networks where the institutions have common exposures to *illiquid* assets are described by more advanced models than the classical Eisenberg-Noe model (see, e.g., Cifuentes et al. (2005); Feinstein (2017)) that revise the definition of a clearing vector and take liquidation costs into account. Even more complex interconnections between assets and liabilities are possible. Some recent models (Altermatt et al. 2024) examine self-fulfilling bank runs, where the value of bank assets hinges on households' ability to redeem their deposits. In these settings, liquidity misallocations–where cash is withdrawn by households not intending to consume - undermine firm revenues and, consequently, bank asset returns, reinforcing the panic. A bank run occurs when many depositors simultaneously try to withdraw funds out of fear of failure. Since banks hold only a fraction of deposits as liquid reserves, such withdrawals can trigger or accelerate insolvency. Robustness of such models against external (market-driven or liquidity-driven) asset price fluctuations remains, to the best of our knowledge, an open problem. Note that these contagion effects reflect higher-order interactions within the network (Boccaletti et al. 2023), involving multiple nodes simultaneously and not reducible to simple pairwise relationships.

Supplementary Information

The online version contains supplementary material available at https://doi.org/10.1007/s41109-025-00728-5.

Supplementary file 1 (pdf 199 KB)

Author contributions
The authors contributed equally to this work.

Funding
The work has been supported by the project 2022K8EZBW "Higher-order interactions in social dynamics with application to monetary networks", funded by European Union–Next Generation EU within the PRIN 2022 program (D.D. 104 -

02/02/2022 Ministero dell'Università e della Ricerca). This manuscript reflects only the authors' views and opinions, and the Ministry cannot be considered responsible for them. The authors have no conflicts of interest to disclose.

Data availability
Data used in the numerical examples has been created randomly. It can be reproduced by following the examples' description in the paper.

Declarations

Ethics approval and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
The authors declare no Conflict of interest.

Received: 28 February 2025 / Accepted: 28 July 2025
Published online: 30 September 2025

References

Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. Am Econ Rev 105(2):564–608
Allen F, Gale D (2000) Financial contagion. J Polit Econ 108(1):1–33
Allen F, Babus A, Carletti E (2012) Asset commonality, debt maturity and systemic risk. J Financ Econ 104(3):519–534
Altermatt L, van Buggenum H, Voellmy L (2024) Systemic bank runs without aggregate risk: how a misallocation of liquidity may trigger a solvency crisis. J Financ Econ 161:103929
Ararat C, Meimanjan N (2023) Computation of systemic risk measures: a mixed-integer programming approach. Oper Res 71(6):2130–2145
Babus A (2016) The formation of financial networks. RAND J Econ 47(2):239–272
Banerjee T, Feinstein Z (2022) Pricing of debt and equity in a financial network with comonotonic endowments. Oper Res 70(4):2085–2100
Banerjee T, Bernstein A, Feinstein Z (2025) Dynamic clearing and contagion in financial networks. Eur J Oper Res 321(2):664–675
Barucca P, Bardoscia M, Caccioli F, D'Errico M, Visentin G, Caldarelli G, Battiston S (2020) Network valuation in financial systems. Math Financ 30(4):1181–1204
Barucca P, Mahmood T, Silvestri L (2021) Common asset holdings and systemic vulnerability across multiple types of financial institution. J Financ Stab 52:100810
Benos E, Huang W, Menkveld A, Vasios M (2024) The cost of clearing fragmentation. Manag Sci 70(6):3581–3596
Boccaletti S, De Lellis P, del Genio CI, Alfaro-Bittner K, Criado R, Jalan S, Romance M (2023) The structure and dynamics of networks with higher order interactions. Phys Rep 1018:1–64
Boyd S, Vandenberghe L (2004) Convex optimization. Cambridge University Press, New York
Caiazzo E, Zazzaro A (2025) Bank diversity and financial contagion. J Financ Stab 77:101392
Calafiore GC, Fracastoro G, Proskurnikov AV (2023) Clearing payments in dynamic financial networks. Automatica 158:111299
Calafiore GC, Fracastoro G, Proskurnikov AV (2024) Optimal clearing payments in a financial contagion model. SIAM J Financ Math 15(2):473–502
Cifuentes R, Ferrucci G, Shin HS (2005) Liquidity risk and contagion. J Eur Econ Assoc 3(2–3):556–566
Craig B, von Peter G (2014) Interbank tiering and money center banks. J Financ Intermed 23(3):322–347
Csóka P, Herings PJ-J (2021) An axiomatization of the proportional rule in financial networks. Manag Sci 67(5):2799–2812
Csóka P, Herings PJ-J (2024) Uniqueness of clearing payment matrices in financial networks. Math Oper Res 49(1):232–250
Demirer M, Diebold FX, Liu L, Yilmaz K (2021) An integrated model for fire sales and default contagion. Math Financ Econ 15:59–101
Eisenberg L, Noe TH (2001) Systemic risk in financial systems. Manag Sci 47(2):236–249
Elsinger H, Lehar A, Summer M (2006) Risk assessment for banking systems. Manag Sci 52(9):1301–1314
Erol S, Parise F, Teytelboym A (2023) Contagion in graphons. J Econ Theory 211:105673
Feinstein Z (2017) Financial contagion and asset liquidation strategies. Oper Res Lett 45(2):109–114
Feinstein Z, Hałłaj G (2023) Interbank asset-liability networks with fire sale management. J Econ Dyn Control 155:104734
Feinstein Z, Pang W, Rudloff B, Schaanning E, Sturm S, Wildman M (2018) Sensitivity of the Eisenberg-Noe clearing vector to individual interbank liabilities. SIAM J Financ Math 9(4):1286–1325
Freixas X, Parigi B, Rochet J-C (2000) Systemic risk, interbank relations, and liquidity provision by the central bank. J Money Credit Bank 32(3):611–638
Glasserman P, Young HP (2016) Contagion in financial networks. J Econ Lit 54(3):779–831
Hu F, Mitchell D, Tompaidis S (2024) Robust financial networks. Oper Res 72(5):1827–1842
Hurd TR (2023) Systemic cascades on inhomogeneous random financial networks. Math Financ Econ 17:1–21
Jarrow RA, Protter P, Quintos A (2024) Computing the probability of a financial market failure: a new measure of systemic risk. Ann Oper Res 336(1):481–503
Jie EY, Ma Y (2025) An analysis of incomplete and random financial networks. J Math Finance 15:277–305
Jurakovaite O, Gaigaliene A (2024) Regionalization and continuing globalization scenarios as a result of re-globalization: a matter of ally-familiarity proximity in multiplex financial network. Appl Netw Sci 9(1):78
Kanellopoulos P, Kyropoulou M, Zhou H (2024) On priority-proportional payments in financial networks. Theoret Comput Sci 1014:114767

Kusnetsov M, Veraart LAM (2019) Interbank clearing in financial networks with multiple maturities. SIAM J Financ Math 10(1):37–67
Li B, Zhang X (2024) Systemic risk and financial networks. Q Rev Econ Finance 94:25–36
Liu M, Staum J (2010) Sensitivity analysis of the Eisenberg-Noe model of contagion. Oper Res Lett 38(5):489–491
Ma J, Zhu S, Li D (2024) Measuring financial systemic risk: net liability clearing mechanism and contagion effect. J Syst Sci Complex 37(4):1114–1146
Mangasarian OL (1979) Uniqueness of solution in linear programming. Linear Algebra Appl 25:151–162
Massai L, Como G, Fagnani F (2022) Equilibria and systemic risk in saturated networks. Math Op Res 47(3):1707–2545
Nier E, Yang J, Yorulmazer T, Alentorn A (2007) Network models and financial stability. J Econ Dyn Control 31(6):2033–2060
Pacelli V, Panetta IC, Povia MM (2025). In: Pacelli V (ed) Systemic risk and network science: a bibliometric and systematic review. Springer, Cham, pp 21–42
Pallante G, Guerini M, Napoletano M, Roventini A (2025) Robust-less-fragile: tackling systemic risk and financial contagion in a macro agent-based model. J Financ Stab 76:101352
Rio-Chanona RM, Korniyenko Y, Patnam M, Porter MA (2020) The multiplex nature of global financial contagions. Appl Netw Sci 5(1):74
Rogers LC, Veraart LA (2013) Failure and rescue in an interbank network. Manag Sci 59(4):882–898
Tang M, Fan H (2025) Systemic risk under different clearing methods for multiple derivatives. Appl Econ, 1–18
Veraart LAM, Aldasoro I (2025) Systemic risk in markets with multiple central counterparties. Math Financ 35(1):214–262
Weber S, Weske K (2017) The joint impact of bankruptcy costs, fire sales and cross-holdings on systemic risk in financial networks. Probab Uncertain Quant Risk 2:9
Zino L, Proskurnikov AV, Fracastoro G, Calafiore GC (2025) Structure and shock resistance in financial networks: A numerical study. In: Cherifi H, Donduran M, Rocha LM, Cherifi C, Varol O (eds) Complex Networks & Their Applications XIII. Springer, Cham, pp 103–114

Publisher's Note

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