Computing Tarski Fixed Points in Financial Networks

Leander Besting
Department of Computer Science, RWTH Aachen University, Germany
Martin Hoefer
Department of Computer Science, RWTH Aachen University, Germany
Lars Huth
Department of Computer Science, RWTH Aachen University, Germany

Abstract

Modern financial networks are highly connected and result in complex interdependencies of the involved institutions. In the prominent Eisenberg-Noe model [12], a fundamental aspect is *clearing* – to determine the amount of assets available to each financial institution in the presence of potential defaults and bankruptcy. A clearing state represents a fixed point that satisfies a set of natural axioms. Existence can be established (even in broad generalizations of the model) using Tarski's theorem. While a *maximal* fixed point can be computed in polynomial time, the complexity of computing other fixed points is open. In this paper, we provide an efficient algorithm to compute a *minimal* fixed point. Our algorithm applies in a broad generalization of the Eisenberg-Noe model with any monotone, piecewise-linear payment functions and default costs. We also study claims trading, a local network adjustment to improve clearing, when networks are evaluated with minimal clearing. We provide an efficient algorithm to decide existence of Pareto-improving trades and compute optimal ones if they exist.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Networks → Network algorithms

Keywords and phrases Tarski Fixed Points, Financial Networks, Minimal Clearing, Claims Trade

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.14

Funding This work was supported by DFG Research Unit ADYN (project number 411362735) and grant DFG Ho 3831/9-1 (project number 514505843).

1 Introduction

Modern financial systems exhibit highly complex debt relationships between their constituents. An important concern in these networks is systemic risk – after a shock, financial institutions become pressured to pay back (or *clear*) their debt. This leaves some of them in default. Consequently, creditors receiving little or no payments from their defaulting debtors might in turn be unable to meet their own obligations. As such, default can quickly propagate throughout the whole network. This is a realistic concern with the most well-known occurrence being the financial crisis of 2008 (and many other, less severe episodes since then).

The canonical framework to understand properties of debt clearing in financial networks is the Eisenberg-Noe model [12]. It has, for example, been used by the European Central Bank in its STAMP€ framework for financial stress-testing [32]. There are n financial institutions (termed "banks" throughout), which are represented by nodes in an edge-weighted, directed graph. There are m edges, each representing a debt claim, with an edge weight expressing the liability of the claim. Banks have (usually non-negative) external assets, which capture funds available for clearing that are not part of the claim network. The basic solution concept in the Eisenberg-Noe model is a *clearing state*, which yields an assignment of assets of banks and payments on each edge that satisfies a set of natural axioms. When a creditor bank

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).
Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thắng
Article No. 14; pp. 14:1–14:18
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

is in default, its claims will not be valued by the liability but only the amount that the creditor will pay off in accordance with their legal requirements[1]. As one of the axioms, the Eisenberg-Noe model assumes *proportional* debt clearing, i.e., for a bank in default there is a *recovery rate* given by the ratio of total assets available to the total liabilities of the bank. The recovery rate translates directly into proportional payments for all claims where the bank is the creditor. For a more formal discussion of the model, see Section 2 below.

Clearing states in the Eisenberg-Noe model are Tarski fixed points. The Knaster-Tarski theorem states that for any monotone function mapping a complete lattice to itself, the set of fixed points of the function constitutes a complete lattice. This implies, in particular, that at least one such fixed point exists and that fixed points might not be unique. The complexity of computing Tarski fixed points has been of significant interest recently, especially on the k-dimensional grid (see, e.g., [5, 6] and the references therein).

While it is known how to compute the *maximal* fixed point in the Eisenberg-Noe model in polynomial time [12, 27], we concentrate on finding other fixed points (most notably, the *minimal* one). A maximal clearing state is desirable from a centralized perspective, since it yields pointwise maximal assets and payments for each bank and claim, respectively. However, it is unlikely that in the current multi-national financial system, a central coordination agency can dictate such assets and payments in times of crises. As such, arguably, the emergence of other clearing states is more realistic. The *minimal* one has been shown to emerge as the limit of a natural decentralized clearing process [8]. More generally, other clearing states can also emerge if agents update their recovery rates in a sequential fashion [26]. However, to our knowledge, finding efficient algorithms to compute such clearing states or characterizing their computational complexity are important open problems.

Our Results and Techniques. We show that a minimal clearing state in the Eisenberg-Noe model can be computed in polynomial time. In contrast to computing the maximal one, it seems impossible to formulate the problem directly as an LP. Our approach in Section 3 is a careful adjustment of the bottom iteration for Tarski fixed points. A naive implementation of the bottom iteration (which results, e.g., from distributed clearing processes [8]) would start with the external assets at each bank and then propagate payments and assets into the system. This can take infinite time to converge to the fixed point. Instead of propagating all external assets directly, we inject external assets sequentially and simulate their (infinite) propagation by solving several carefully designed LPs. In particular, upon injection of additional assets at a node v, our algorithm distinguishes two cases: (1) all additional assets circulate and gradually reach a "sink" bank without any solvency event, and (2) some assets keep circulating until another bank becomes solvent. Case (1) yields a linear increase in the payments of every claim, and we compute the slopes by solving a polynomial-sized system of linear equations. Then we can inject the assets at v and compute the increase in assets and payments using the slopes. Case (2) arises if any positive portion of the assets injected at v eventually reaches a part C of the network that we term *flooded* – where all outgoing paths of any node eventually return to that node. More formally, consider the strongly connected components (SCCs) of the remaining network of open claims. The graph of SCCs is a DAG, and a flooded part C is exactly a sink component of the SCC graph, which contains more than one bank. In this case, we do not inject any additional assets at v but instead increase the payments within the sink component C by a circulation (consistent with proportional payments) until a bank becomes solvent. This can be computed using a suitable LP. We

---

[1] In the United States, the legal framework for this is given by Chapter 11 bankruptcy [7].

then update the SCC graph and again attempt to inject additional assets at v. Overall, we see that in each iteration, we either fully inject the external assets of a bank or create a solvent bank. Thus, at most $2n$ iterations are needed.

Our algorithmic ideas turn out to be very powerful to address many generalizations of the problem. First, we relax the condition of proportional payments to the class of arbitrary *monotone and piecewise-linear* payment functions. This class includes many important examples considered in the literature recently, including edge ranking [4, 14, 17] (or singleton liability priorities [20]), unit-ranking [4, 8, 17], or priority-proportional [13, 22]. By using appropriate granularity, we can handle arbitrary monotone payment functions via approximation with piecewise-linear ones. The extension requires applying our algorithm and the analysis in phases, where each phase ends with reaching an interval border for the payment function of any edge. In each phase, all payments behave linearly, so slopes in case (1) and circulation in case (2) can be obtained by setting up and solving appropriate systems of linear equations. The interval borders of the payment functions limit the increase in external assets of v in case (1), as well as the amount of the circulation in case (2).

Second, we show how to handle an extension of the model to *default costs* for insolvent banks [27]. Here, the assets of each defaulting bank are additionally reduced by a constant factor. This adjustment represents a linear decrease in the assets, which can be integrated into the network. Since we monotonically increase payments and assets throughout, we (only) need to be careful when a bank v becomes solvent. The challenge is that the fixed point function ceases to be continuous from below. Thus, a standard bottom iteration might fail to guarantee convergence to a minimal clearing state in the limit.

For each bank v, we can represent default cost using an auxiliary claim to an auxiliary individual "sink bank" such that this bank receives the default cost. When v becomes solvent, the default cost vanishes, we remove the auxiliary sink bank, and inject the remaining assets as additional external assets at the out-neighbors of v. This view directly indicates how our approach of sequential injection of external assets can be extended to handle default cost.

More fundamentally, in Section 3.3 we show an interesting structural equivalence. Every financial network with piecewise-linear payment functions can be transformed into an equivalent network with priority-proportional payments. The transformation increases the representation by at most a polynomial factor. This shows the generality of priority-proportional payment functions. Moreover, it allows applying an existing algorithm for the computation of *maximal* clearing states [22] to networks with arbitrary monotone, piecewise-linear payments and default cost.

Finally, in Section 4 we study *claims trades* as a network adjustment to influence the minimal clearing state. The operation was recently formalized in [16]. In a claims trade, a given bank w strives to buy a given claim $e = (u, v)$ from a creditor bank v. Formally, w should pay a return ρ from its external assets to v and become creditor of claim e. The goal is to give liquidity to v and raise the assets of v in order to mitigate contagion effects as much as possible. In contrast to previous work, we focus on the scenario when the network is evaluated with a minimal clearing state.

Ideally, one would want to compute a return ρ such that *both* v and w *strictly* benefit from the trade (w.r.t. the minimal clearing state). We show that this is impossible for all networks with strictly monotone payment functions. This result uses a novel characterization for banks, for which the clearing state is not unique (Lemma 14). Our interest then lies in creditor-positive trades, in which v strictly profits but w stays at least indifferent. For networks with piecewise-linear payment functions, we show how to decide in polynomial time if a creditor-positive claims trade exists, and how to compute one that maximizes the post-trade assets of v if it exists. Our proof shows that the set of returns that yield creditor-positive trades forms a consecutive interval. The largest such return maximizes the assets of v, and we can find it by a combination of binary search and maximization of suitable LPs.

For spatial reasons, all missing proofs are deferred to the full version of the paper.

Further Related Work. Algorithmic aspects of the Eisenberg-Noe model for financial networks have become a popular topic in recent years. A *maximal* clearing state can be computed in polynomial time for proportional payments [12], even more generally in networks with default cost [27] and with priority-proportional payments [22]. It is known that decentralized update procedures (which essentially implement a bottom-iteration) converge to the *minimal* clearing state [8], but these approaches do not necessarily run in polynomial time. Up to our knowledge, an efficient algorithm for minimal clearing has been derived only for networks with edge-ranking payments without default cost [4].

The complexity of computing clearing states has also been considered in an extension of the model to credit-default swaps [29, 30], where clearing states exist but their entries are not necessarily rational. Some notions of approximation yield PPAD-hardness results [29], even for constant approximation factors [10, 19]. Stronger notions of approximation even give rise to FIXP-hardness [18, 20].

Claims trades have recently been studied in the Eisenberg-Noe model without default cost and with *maximal* clearing [16]. Since a claims trade can be interpreted as a debt swap operation [14, 25], it is impossible that both creditor and buyer banks profit strictly. An optimal creditor-positive trade can be computed in polynomial time for proportional and edge-ranking payment functions. For general functions, it is possible to compute approximately optimal trades in polynomial time. When trading either multiple incoming or multiple outgoing claims of a single bank, finding optimal trades becomes NP-hard. For incoming claims, there is a bicriteria approximation for all monotone payment functions, for outgoing the problem is NP-hard to approximate within polynomial factors even for edge-ranking functions. These results were recently extended and improved in a model with *fractional* claims trades, for networks with proportional payments and default cost [15].

Our paper is related to a growing body of work that studies structural and algorithmic properties of game-theoretic scenarios based on the Eisenberg-Noe model. Aspects that have received attention include, for example, strategic payment allocation [4, 17, 22, 24], forgiving, cancelling or forwarding debt [21, 23, 33], donations [34], prepayments [36], lending [11], and more [3]. With the exception of [17], all these works consider only maximal clearing states.

More generally, there is work on the complexity of improvement measures for the clearing properties in Eisenberg-Noe financial networks, such as debt swapping [14, 25] or portfolio compression. In portfolio compression, a debt cycle is eliminated from the network. For proportional payments, this can have counter-intuitive effects [28, 35], and many optimization problems surrounding this operation are NP-hard [1].

2 Model and Preliminaries

Network Model. We define a financial network $\mathcal{F} = (V, E, \ell, \mathbf{a}^{\mathbf{x}})$. There is a set V of n banks, and a set E of m directed edges. Each edge $e = (u, v) \in E$ has a non-negative edge weight ℓ_e that represents the *liability* of a claim with debtor u and creditor v. Each bank v has non-negative *external assets* $a_v^{(x)} \geqslant 0$. For simplicity, we assume that the network has no self-loops or multi-edges[2]. We define the set of *outgoing* and *incoming* claims of v by $E^+(v) = \{e = (v,u) \in E \mid u \in V\}$ and $E^-(v) := \{e = (u,v) \in E \mid u \in V\}$, respectively. The *total outgoing* and *incoming liabilities* of bank v are $L^+(v) = \sum_{e \in E^+(v)} \ell_e$ and $L^-(v) = \sum_{e \in E^-(v)} \ell_e$, respectively.

Payment Functions. The basic solution concept in the Eisenberg-Noe model is a *clearing state*, which defines a consistent set of assets $a_u \geqslant 0$ for each bank $u \in V$. In the standard model, each bank is assumed to clear debt *proportionally.* Thus, with assets a_u, we have a *recovery rate* $\min(1, a_v/L^+(v))$, such that each claim $e = (u,v)$ is cleared to the same fractional extent, i.e., $p_e(a_u) = \min(1, a_v/L^+(v)) \cdot \ell_e$. Using these payments, the assets have to satisfy the natural asset axioms: The *(total) assets* of each bank are given by the external assets plus the incoming payments of other banks

$$a_v = a_v^{(x)} + \sum_{(u,v) \in E^-(v)} p_{(u,v)}(a_u) \quad \text{for each } u \in V.$$

In particular, this implies that no bank will conduct fraud by generating money or holding back assets from paying its open claims.

We consider several extensions of the model. First, we consider the model with *default cost* [27]. Each bank v has a *default rate* $\rho_v \in [0,1]$. If the bank is insolvent, the available assets a_v that can be used to clear debt are reduced to $\rho_v \cdot a_v$. The asset axioms for the vector of total assets $\mathbf{a} = (a_v)_{v \in V}$ of the banks become

$$a_v = \begin{cases} a_v^-(\mathbf{a}) & \text{if } a_v^-(\mathbf{a}) \geqslant L^+(v) \text{ (i.e., } v \text{ solvent), and} \\ \rho_v \cdot a_v^-(\mathbf{a}) & \text{otherwise.} \end{cases} \tag{1}$$

with the *incoming assets* of v (before default cost reduction) given by

$$a_v^-(\mathbf{a}) = a_v^{(x)} + \sum_{(u,v) \in E^-(v)} p_{(u,v)}(a_u).$$

We recover the standard model without default cost when $\rho_v = 1$ for all $v \in V$.

Second, we address a general class of monotone payments [4]. Each bank u has a *payment function* $p_e : \mathbb{R} \to \mathbb{R}$ for each claim $e = (u,v)$. A payment function satisfies, for each $e \in E$, $u \in U$ and $a_u, \varepsilon > 0$

$$\begin{aligned} p_e(a_u) &\in [0, \ell_e] && \text{(no edge under- or overpaid)} \\ p_e(a_u) &\leqslant p_e(a_u + \varepsilon) && \text{(monotonicity)} \\ \sum_{e \in E^+(u)} p_e(a_u) &= \min\{a_u, L^+(u)\} && \text{(no fraud)} \end{aligned} \tag{2}$$

These constraints imply that all p_e must be continuous (since $p_e(a+\rho) - p_e(a) \leqslant \rho$ for any $\rho > 0$). The axioms are trivially fulfilled for proportional payments. Many other natural examples of payment functions from this class have been considered, e.g., priority/edge-ranking [4,8] (rank edges in an order and pay them sequentially until running out of assets), constrained equal awards or losses [9] (all claims receive the same payment or the same non-payment, up their liability) or priority-proportional [22] (partition edges into sets, rank

---

[2] These aspects can be incorporated by increasing the notational overhead. We leave the straightforward extension to the interested reader.

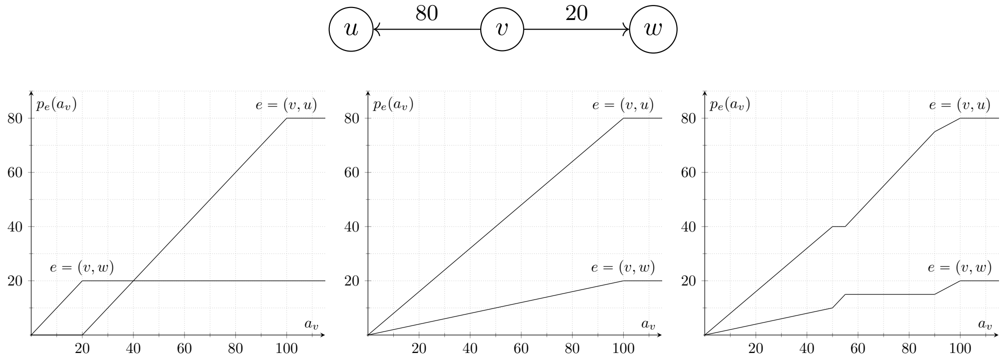

Figure 1 Consider the financial network displayed at the top. Payments of bank v on each edge for edge-ranking payment functions with (v, w) ranked first (left); proportional payment functions (middle); piecewise-linear functions with interval borders $0, 50, 55, 90, 100, \infty$ (right).

sets in an order, pay edges proportionally within each set, and sequentially over sets until running out of assets). All these examples share a natural property – they are *piecewise-linear*. In this paper, we concentrate on the class of monotone, piecewise-linear payment functions.

▶ **Definition 1.** *A* piecewise-linear *payment function $p_e : \mathbb{R} \to \mathbb{R}$ for an edge $e \in E^+(v)$ is given by $k_e \geqslant 1$ interval borders $0 = x_{e,0} < x_{e,1} < \ldots < x_{e,k_e} = L^+(v) < x_{e,k_e+1} = \infty$ and slopes $m_{e,i} \geqslant 0$ for each $i = 1, \ldots, k_e + 1$ such that*

$$p_e(a) = m_{e,i} \cdot (a - x_{e,i-1}) + p_e(x_{e,i-1}) \qquad \textit{for any } a \in [x_{e,i-1}, x_{e,i}).$$

Moreover, p_e adheres to the three axioms in (2).

Note that $p_e(a) = \ell_e$ when $a \geqslant L^+(v) = x_{e,k_e}$. Thus $m_{e,k_e+1} = 0$. We further define $m_e(a)$ as the slope $m_{e,i}$ that applies for argument a, i.e., $m_e(a) = m_{e,i}$ such that $a \in [x_{e,i}, x_{e,i+1})$. This implies that, for every $v \in V$,

$$\sum_{e \in E^+(v)} m_e(a) = \begin{cases} 1 & \text{for all } 0 \leqslant a < L^+(v) \\ 0 & \text{for all } a \geqslant L^+(v). \end{cases} \tag{3}$$

For each bank's available assets, we define the additional amount of assets until a new interval is reached by $\delta_e(a) = \min\{x_{e,i} - a \mid x_{e,i} > a\}$. We denote the total number of interval borders in $\mathcal{F}$ by $k = \sum_{e \in E} k_e \geqslant m$. Finally, for a given vector of total assets $\mathbf{a}$, we call an edge $e = (u, v)$ *active* if $m_e(a_u) > 0$. More generally, the *set of active edges* for $\mathbf{a}$ is $E_{\mathbf{a}} = \{e \in E \mid e \text{ is active for } \mathbf{a}\}$, and the *active graph* is $G_{\mathbf{a}} = (V, E_{\mathbf{a}})$. For each active edge $e \in E_{\mathbf{a}}$, the *active interval* is the index i such that $p_e(a_v) \in [x_{e,i}, x_{e,i+1})$.

Clearing States. Given a financial network $\mathcal{F}$ and piecewise-linear payment functions $\mathbf{p} = (p_e)_{e \in E}$, a *clearing state* is simply a vector of assets $\mathbf{a} = (a_v)_{v \in V}$ such that all asset axioms (1) are fulfilled. These axioms are fixed point conditions, since the value of a_v depends on other a_u (and potentially vice versa). Consider the space of all potential asset vectors $\mathcal{A} = \times_{v \in V} [0, L^-(v) + a_v^{(x)}]$. We define a function $\Phi_{\mathbf{a}^{(x)}}(\mathbf{a}) : \mathcal{A} \to \mathcal{A}$ resulting from applying the map defined by the asset axioms (1). $\Phi_{\mathbf{a}^{(x)}}$ is monotone since all p_e are monotone. More formally, if $\mathbf{a}' \geqslant \mathbf{a}$ coordinate-wise, then $\Phi_{\mathbf{a}^{(x)}}(\mathbf{a}') \geqslant \Phi_{\mathbf{a}^{(x)}}(\mathbf{a})$ coordinate-wise. The fixed points of Φ are the clearing states. Applying the Knaster-Tarski theorem, we see that the set $\mathcal{A}_f$ of fixed points and the coordinate-wise $\geqslant$-operation form a complete lattice.

We denote by $\hat{\mathbf{a}}$ the minimal fixed point and by $\check{\mathbf{a}}$ the maximal one. A natural attempt towards computation of $\hat{\mathbf{a}}$ is a *bottom iteration*: Starting from any $\mathbf{a}^0 \leqslant \hat{\mathbf{a}}$ (say, $a_v^0 = a_v^{(x)}$ for all $v \in V$) we iteratively compute $\mathbf{a}^{i+1} = \Phi_{\mathbf{a}^{(x)}}(\mathbf{a}^i) \leqslant \hat{\mathbf{a}}$. Monotonicity directly implies that $\mathbf{a}^{i+1} \geqslant \mathbf{a}^i$. In the model without default cost (all default rates $\rho_v = 1$) and arbitrary monotone payment functions, it is easy to see that $\lim_{i\to\infty} \mathbf{a}^i = \hat{\mathbf{a}}$, but there are simple instances where $\mathbf{a}^i \neq \hat{\mathbf{a}}$ for every $i \geqslant 0$. More generally, when $\rho_v < 1$ for some $v \in V$, there are simple instances where $\lim_{i\to\infty} \mathbf{a}^i \neq \hat{\mathbf{a}}$. For example, bank v can be solvent in $\hat{\mathbf{a}}$ but insolvent in $\mathbf{a}^i$, for every $i \geqslant 0$. Then the limit of the total assets of v is $\lim_{i\to\infty} a_v^i \leqslant \rho \cdot L^+(v) < L^+(v) \leqslant \hat{a}_v$.

Claims Trades. An operation to improve clearing states are claims trades. In a claims trade, we are given a claim $e = (u, v)$ and a potential buyer bank w. The bank w buys the claim by paying an amount of external assets ρ to v. We call ρ the *return*. In turn, the creditor of the claim e is changed from v to w. We assume that the return is upper bounded by $\rho \leqslant \min\{a_w^{(x)}, \ell_{(u,v)}\}$. After the trade, the external assets of w are $a_w^{(x)} - \rho$ and the ones of v are $a_v^{(x)} + \rho$. A claims trade can represent a *donation*, in which w only transfers ρ external assets without changing any edges in the network[3].

We consider claims trades when the network is evaluated with $\hat{\mathbf{a}}$. A claims trade is called *creditor-positive* if there exists a return ρ such that the assets of creditor bank v are strictly improved and the assets of buyer bank w remain at least the same. We call ρ a *creditor-positive return*. If it exists, we look for an *optimal* one, i.e., a creditor-positive return ρ^* that maximizes the post-trade assets of v. Note that any creditor-positive trade Pareto-improves the assets in the entire network.

3 Computing Clearing States

3.1 Minimal Clearing States

In this section, we explain our algorithm to compute a minimal clearing state. The standard approach to computing minimal Tarski fixed points is a bottom iteration, but implementing this directly in financial networks is not effective. Instead, Algorithm 1 maintains and increases a vector $\mathbf{b}^{(x)}$ of *reduced external assets*. It starts at $\mathbf{b}^{(x)} = \mathbf{0}$ and approaches the vector of actual external assets from below, i.e., $\mathbf{b}^{(x)} \leqslant \mathbf{a}^{(x)}$. In each iteration of the main while-loop, the algorithm computes an increase of reduced external assets at one vertex v. It also maintains a vector $\mathbf{b}$ of total assets.

For the analysis, we will maintain the invariant that at the beginning of an iteration of the main while-loop, $\mathbf{b}$ is a minimal clearing state of the network $\mathcal{F}$ with the reduced external assets $\mathbf{b}^{(x)}$.

▶ **Lemma 2.** *At the beginning of each iteration of the main while-loop,* $\mathbf{b}^{(x)} \leqslant \mathbf{a}^{(x)}$ *and* $\mathbf{b}$ *is a minimal clearing state of the financial network* $\mathcal{F}$ *with external assets* $\mathbf{b}^{(x)}$*.*

For clarity of exposition, we first prove correctness of our algorithm when there is no default cost, i.e., all default rates are $\rho_v = 1$. We then outline the adjustments for general default rates in Section 3.2 below. In particular, without default cost, $D = \varnothing$ in line 2, we make no adjustments in lines 3 and 17, and we never execute lines 13-16.

---

[3] To formulate donations as special cases of claims trades, we may simply assume that w trades an auxiliary claim (u, v) with liability $a_w^{(x)}$ and no existing payments from an auxiliary debtor bank u of v.

Algorithm 1 Computation of a minimal clearing state.

```
Input   : Financial network $\mathcal{F}$
Output : Minimal clearing state $\hat{\mathbf{a}}$
$\mathbf{b}^{(x)} \leftarrow \mathbf{0}$, $\mathbf{b} \leftarrow \mathbf{0}$                    // External and outgoing assets
$D \leftarrow \{u \in V \mid \rho_u < 1\}$
Adjust each bank $u \in D$ with auxiliary sink bank and scaled payment functions
while there is bank $v \in V$ with $b_v^{(x)} < a_v^{(x)}$ do
    // Repeated flooding of reachable sink-SCCs
    $\mathcal{C} \leftarrow$ strongly connected components (SCC) of the active graph $G_{\mathbf{b}}$
    $G_{\mathcal{C}} \leftarrow$ directed acyclic graph of SCCs of $G_{\mathbf{b}}$
    while there is non-singleton $C \in \mathcal{C}$ that is reachable from $v$ and a sink in $G_{\mathcal{C}}$ do
        Solve Flood-LP (5), let $\mathbf{d}^*$ be the optimal solution
        $\mathbf{b} \leftarrow \mathbf{b} + \mathbf{d}^*$
        Update $G_{\mathbf{b}}$, $\mathcal{C}$, and $G_{\mathcal{C}}$
    // Raise external assets of $v$
    Solve Increase-LP (6), let $(\delta^*, \mathbf{d}^*)$ be an optimal solution
    $b_v^{(x)} \leftarrow b_v^{(x)} + \delta^*$ and $\mathbf{b} \leftarrow \mathbf{b} + \mathbf{d}^*$
    // Adjust the network for banks with default cost
    forall banks $u \in D$ that became solvent do
        $b_w^{(x)} \leftarrow b_w^{(x)} + p_e(b_u)$, for each $e = (u, w) \in E$
        $a_w^{(x)} \leftarrow a_w^{(x)} + \ell_e$, for each $e = (u, w) \in E$
        Remove all edges $e = (u, w)$ from $\mathcal{F}$
Remove auxiliary sink banks and scale assets for every insolvent $u \in D$
return $\mathbf{b}$
```

For the proof of Lemma 2, the properties clearly hold in the beginning since $\mathbf{b}^{(x)} = \mathbf{b} = \mathbf{0}$. Suppose they are true at the beginning of iteration i. We argue that the properties hold in the end of iteration i, i.e., at the beginning of iteration $i + 1$. Our first insight shows that the minimal clearing state is non-decreasing in the external assets of each bank.

▶ **Lemma 3.** *Consider a financial network $\mathcal{F}$. Suppose all banks have reduced external assets $\mathbf{0} \leqslant \mathbf{b}^{(x)} \leqslant \mathbf{a}^{(x)}$. If $a_u^{(x)} > b_u^{(x)}$, then in the minimal clearing state $\hat{\mathbf{b}}$ resulting from $\mathbf{b}^{(x)}$ we have $\hat{a}_u > \hat{b}_u$.*

Proof. For contradiction, assume that $\hat{a}_u \leqslant \hat{b}_u$. Consider the pointwise minimum $\mathbf{b}' = \hat{\mathbf{a}} \wedge \hat{\mathbf{b}}$, i.e., $b'_v = \min\{\hat{a}_v, \hat{b}_v\}$ for all $v \in V$. Since $\mathbf{b}' \leqslant \hat{\mathbf{b}}$, the bottom iteration shows that $\Phi_{\mathbf{b}^{(x)}}(\mathbf{b}') \geqslant \mathbf{b}'$. Now $\hat{\mathbf{a}} \geqslant \mathbf{b}'$, $a_u^{(x)} > b_u^{(x)}$ and Φ is strictly monotone with respect to external assets and weakly monotone with respect to total assets. This implies that

$$\Phi_{\mathbf{a}^{(x)}}(\hat{\mathbf{a}})_u > \Phi_{\mathbf{b}^{(x)}}(\hat{\mathbf{a}})_u \geqslant \Phi_{\mathbf{b}^{(x)}}(\mathbf{b}')_u \geqslant \mathbf{b}'_u = \hat{a}_u.$$

Thus, $\hat{\mathbf{a}}$ is not a clearing state – a contradiction. ◀

Flooding SCCs. Now consider the bank v with $b_v^{(x)} < a_v^{(x)}$ chosen for the increase in iteration i of the main while-loop. Let us concentrate on regions of the active graph $G_{\mathbf{b}}$ that are reachable from v and strongly connected.

▶ **Definition 4.** *We define a* phase *as a subset of asset vectors* $\mathcal{P} \subseteq \mathcal{A}$ *such that for all* $\mathbf{a}, \mathbf{a}' \in \mathcal{P}$ *we have the same active edges* $E_{\mathbf{a}} = E_{\mathbf{a}'}$*, and for each* $e \in E_{\mathbf{a}}$*, the same active interval in* $\mathbf{a}$ *and* $\mathbf{a}'$*.*

▶ **Definition 5.** *We say* v causes a flood in $G_{\mathbf{b}}$ *if there is a strongly connected component* $C \subseteq V$ *of* $G_{\mathbf{b}}$ *that is (1) reachable from* v*, (2) non-singleton, and (3) a sink-component, i.e., does not have outgoing edges to banks outside* C*.*

▶ **Lemma 6.** *Suppose we increase the external assets of bank* v *from* $b_v^{(x)}$ *to* $b_v^{(x)} + \varepsilon$ *for some positive amount* $0 < \varepsilon \leqslant a_v^{(x)} - b_v^{(x)}$*. Let* $\mathbf{b}'$ *be the new minimal clearing state after the increase. If* v *causes a flood in* $G_{\mathbf{b}}$*, then* $\mathbf{b}'$ *cannot be in the same phase as* $\mathbf{b}$*, for any* $\varepsilon > 0$*.*

Proof. By Lemma 3, we know that in the minimal clearing states $b'_v > b_v$. As a consequence, each active outgoing edge $e = (v, w)$ must have strictly higher payments in $\mathbf{b}'$, i.e., $p_e(b'_v) > p_e(b_v)$ because the slope $m_e(b_v) > 0$. This shows

$$
\begin{aligned}
b'_w &= b_w^{(x)} + p_e(b'_v) + \sum_{(u,w) \in E} p_{(u,w)}(b'_u) \geqslant b_w^{(x)} + p_e(b'_v) + \sum_{(u,w) \in E} p_{(u,w)}(b_u) \\
&> b_w^{(x)} + p_e(b_v) + \sum_{(u,w) \in E} p_{(u,w)}(b_u) = b_w
\end{aligned}
$$

Consider the sets of banks and active edges reachable from v in $G_{\mathbf{b}}$. Applying the insight inductively shows that every reachable bank has strictly higher total assets, and every reachable active edge strictly higher payments in $\mathbf{b}'$ than in $\mathbf{b}$, respectively.

Now, suppose for contradiction that $\mathbf{b}' > \mathbf{b}$ and both minimal clearing states are in the same phase. Consider a non-singleton sink-SCC C reachable from v. For every bank $w \in C$, we consider $d_w = b'_w - b_w > 0$. C is not a singleton, so all banks w in C are insolvent. $\mathbf{b}'$ and $\mathbf{b}$ are in the same phase, so (3) implies that all of d_w gets paid to out-neighbors in C. In turn, these payments become additional incoming assets at other nodes of C. Summing over all additional incoming assets of $w \in C$ from neighbors of C, we see that

$$
\sum_{w \in C} d_w \leqslant \sum_{w \in C} \sum_{\substack{v' \in C \\ (v',w) \in E^-_{\mathbf{b}}(w)}} m_{(v',w)}(b'_{v'}) \cdot d_{v'} \tag{4}
$$

Now let w be a node of C that is closest to v, and let (u, w) be an edge on a shortest v-w-path in $G_{\mathbf{b}}$. Since u is reachable from v, we know that $b'_u > b_u$. If $v \in C$, then we can assume $u = v$, and since C is non-singleton, an out-neighbor $w \in C$ with $(v, w) \in E$ must exist. Let $d_{(u,w)} = p_{(u,w)}(b'_u) - p_{(u,w)}(b_u)$, then $d_{(u,w)} > 0$. Since $\mathbf{b}'$ is a clearing state, the asset axiom holds for w. As such, the additional assets d_w are lower bounded by the sum of additional incoming assets over (u, w) and from in-neighbors of C, i.e.,

$$
d_w \geqslant d_{(u,w)} + \sum_{\substack{v' \in C \\ (v',w) \in E^-_{\mathbf{b}}(w)}} m_{(v',w)}(b'_{v'}) \cdot d_{v'} > \sum_{\substack{v' \in C \\ (v',w) \in E^-_{\mathbf{b}}(w)}} m_{(v',w)}(b'_{v'}) \cdot d_{v'}
$$

With (4) this proves that there must be some bank $w' \in C$ such that

$$
d_{w'} < \sum_{\substack{v' \in C \\ (v',w') \in E^-_{\mathbf{b}}(w')}} m_{(v',w')}(b'_{v'}) \cdot d_{v'},
$$

i.e., the additional assets of w' are *strictly less* than its additional incoming assets from C. Thus, $\mathbf{b}'$ is not a clearing state – a contradiction. ◀

$\mathbf{b}'$ cannot be in the same phase as $\mathbf{b}$, since raising the external assets at v would leave only inconsistent assignments of assets for every reachable, non-singleton, sink-SCC C. Then again, C is reachable and external assets of v must be increased to obtain $\hat{\mathbf{a}}$, so the assets in C must also grow. Therefore, we have to raise the clearing state $\mathbf{b}$ within C to escape the current phase and eventually enable an increase of external assets at v. In Algorithm 1, we raise the assets in C by a minimal circulation that maintains the properties of a clearing state (i.e., additional incoming assets are additional outgoing assets) and suffices to advance to the next phase. This is achieved by solving the following LP. We term this a *flooding* of C with assets. Recall that $\delta_e(a)$ is the smallest amount of additional assets required to advance the payment function p_e to the next interval.

$$\begin{aligned} \text{Max.} \quad & \sum_{w \in C} d_w \\ \text{s.t.} \quad & d_w = \sum_{\substack{v' \in C \\ (v',w) \in E^-_{\mathbf{b}}(w)}} m_{(v',w)}(b_{v'}) \cdot d_{v'} && \forall w \in C \\ & d_w \leqslant \delta_e(b_w) && \forall e \in E^+_{\mathbf{b}}(w) \\ & d_w \geqslant 0 && \forall w \in C \end{aligned} \tag{5}$$

Solving LP (5). Let us define the weighted $|C| \times |C|$ adjacency matrix $\mathbf{M}$ for component C with entries $m_{v',w} = m_{(v',w)}(b_{v'})$. C being a SCC and property (3) imply that $\mathbf{M}$ is row-stochastic and irreducible. The first set of constraints can be written as $\mathbf{d} = \mathbf{dM}$. Thus, $\mathbf{d}$ is an eigenvector of $\mathbf{M}$ with eigenvalue 1. The Perron-Frobenius theorem implies $\mathbf{d}$ is non-negative and unique. Thus, we can construct $\mathbf{M}$, compute $\mathbf{d}$, and scale it to the largest multiple such that all constraints $d_w \leqslant \delta_e(b_w)$ remain satisfied. Note that power iteration methods for approximating $\mathbf{d}$ might not be applicable since $\mathbf{M}$ is not necessarily aperiodic.

Increasing External Assets. Flooding of components monotonically increases the assets and changes the phase, but leaves the external assets of v untouched. Thus, after a finite number of repetitions, we must reach a clearing state $\mathbf{b} < \hat{\mathbf{a}}$ such that all sink-SCCs reachable from v are singletons, i.e., solvent banks. In this case, the next lemma shows that there exists a sufficiently small $\delta > 0$ such that for any increase of external assets of v to less than $b_v^{(x)} + \delta$, the resulting minimal clearing state $\mathbf{b}'$ remains in the phase of $\mathbf{b}$.

▶ **Lemma 7.** *Suppose we increase the external assets of bank v from $b_v^{(x)}$ to $b_v^{(x)} + \varepsilon$ for some positive amount $0 < \varepsilon \leqslant a_v^{(x)} - b_v^{(x)}$. Let $\mathbf{b}'$ be the new minimal clearing state after the increase. If v does not cause a flood in $G_{\mathbf{b}}$, there is a $\delta > 0$ such that $\mathbf{b}'$ is in the same phase as $\mathbf{b}$, for every $\varepsilon < \delta$.*

Proof. Similar to our observation above, we consider the additional assets $d_w = b'_w - b_w \geqslant 0$. Since the external assets of v are strictly increased, we see that $d_v > 0$. As $\mathbf{b}'$ is a minimal clearing state, we can assume all banks u that are not reachable from v maintain their current assets and have $d_u = 0$, as they remain unaffected from the increase in external assets at v. Assuming that $\mathbf{b}'$ and $\mathbf{b}$ are in the same phase, the largest increase in the external assets of v is given by LP (6) below. The first set of equalities asserts that the increase in assets of v is given by the additional external assets and the additional incoming assets over edges from $E^-_{\mathbf{b}}(v)$. Similarly, for banks that are reachable from v, the increase in assets is given by the increase in incoming assets. Clearly, $a_v^{(x)} - b_v^{(x)}$ is a trivial upper bound on the maximum increase. To stay in the same phase, we ensure that the (open) active interval on each edge remains the same, which yields a (strict) upper bound of $\delta_e(b_w)$ for each $e = (u, w) \in E_{\mathbf{b}}$. Using weak inequalities, the optimum solution represents the supremum δ as stated in the lemma. It represents the smallest increase to advance the minimal clearing state to the next phase when increasing the external assets of v.

$$\begin{aligned} \text{Max.}\quad & \delta \\ \text{s.t.}\quad & d_v = \delta + \sum_{(u,v)\in E^-_{\mathbf{b}}(v)} m_{(u,v)}(b_u) \cdot d_u \\ & d_w = \sum_{(u,w)\in E^-_{\mathbf{b}}(w)} m_{(u,v)}(b_u) \cdot d_w && \text{for all } w \neq v \text{ reachable from } v \\ & d_w = 0 && \text{for all } w \neq v \text{ unreachable from } v \\ & d_w \leqslant \delta_e(b_w) && \forall e \in E^+_{\mathbf{b}}(w) \\ & \delta \leqslant a_v^{(x)} - b_v^{(x)} \\ & d_w \geqslant 0 && \forall w \in V \end{aligned} \tag{6}$$

To show correctness, we argue that the LP indeed allows a unique optimal solution $(\delta^*, \mathbf{d}^*)$. We can restrict attention to the set $U \subseteq V$ of banks that are reachable from v in $G_{\mathbf{b}}$. Then the first two constraints in (6) above compose a system of linear equations describing the asset increase that can be expressed by

$$\mathbf{d} = \delta \cdot \mathbf{e}_v + \mathbf{d}\,\mathbf{M}. \tag{7}$$

Here $\mathbf{d}$ is an $|U|$-dimensional row vector with entries d_u, $\mathbf{e}_v$ is an $|U|$-dimensional unit row vector with entry 1 for v and 0 otherwise, and $\mathbf{M}$ is an $|U|$-dimensional square matrix with entries $m_{u,w} = m_{(u,w)}(b_u)$ for all $(u, w) \in E_{\mathbf{b}}, u \in U$, and 0 otherwise. The vector $\mathbf{d}$ can be given by

$$\mathbf{d} = (\mathbf{I} - \mathbf{M})^{-1}\, \delta \cdot \mathbf{e}_v \quad .$$

Let us observe that the inverse exists. By (3), $\mathbf{I} - \mathbf{M}$ is weakly diagonally dominant. Since v does not cause a flood, every bank $u \in U$ has a path to a solvent bank in $G_{\mathbf{b}}$. Rows corresponding to solvent banks are strictly diagonally dominant. These properties give rise to a chained variant of diagonal dominance [2] and imply that $\mathbf{I} - \mathbf{M}$ is invertible [31].

Thus, the increase in the assets of the minimal clearing state scales linearly in the increase of $b_v^{(x)}$. Since all active intervals of $\mathbf{b}$ are open, for every sufficiently small value of $\delta > 0$, the resulting asset vector $\mathbf{b} + \mathbf{d}$ indeed remains in the same phase. In particular, let

$$\mathbf{s} = (\mathbf{I} - \mathbf{M})^{-1} \cdot \mathbf{e}_v$$

be the vector of slopes. For the supremum for all increases that keep $\mathbf{b} + \mathbf{d}$ in the same phase, we require $d_u = s_u \delta \leqslant \delta_{(u,w)}(b_u)$ for all $(u, w) \in E_{\mathbf{b}}$, which are the fourth set of constraints in (6). Clearly, we also require $\delta \leqslant a_v^{(x)} - b_v^{(x)}$, the maximal increase in external assets of v. The lemma follows using

$$\delta = \min\{a_v^{(x)} - b_v^{(x)}, \min_{(u,w)\in E_{\mathbf{b}}} \delta_{(u,w)}(b_u)/s_u\} > 0. \tag{8}$$

◀

Solving LP (6). The proof of the previous lemma shows that to solve the LP, we can compute the set U of reachable banks and set up the matrix $(\mathbf{I} - \mathbf{M})$. We then solve $(\mathbf{I} - \mathbf{M})\mathbf{s} = \mathbf{e}_v$ (e.g., by Gaussian elimination) to obtain the slopes $\mathbf{s}$, and determine δ by computing the minima in (8).

Correctness and Polynomial Time. We are now ready to prove Lemma 2 as the main invariant of the algorithm.

Proof of Lemma 2. Consider round i of the while-loop and the vertex v chosen for the increase in external assets $b_v^{(x)}$. Suppose v causes a flood. Lemma 6 shows that by repeatedly flooding the corresponding sink-SCCs, we maintain a clearing state that remains below any minimal clearing state resulting from any strict increase in $b_v^{(x)}$.

There can only be a finite number of flooding operations. Afterwards, the proof of Lemma 7 reveals that the increase in the minimal clearing state is linear in the increase in external assets of v. We execute the smallest increase δ that either yields $b_v^{(x)} + \delta = a_v^{(x)}$ or changes the phase. In any case, in the beginning of iteration $i + 1$, the vector $\mathbf{b}$ is again the minimal clearing state for $\mathcal{F}$ with the larger vector of external assets $\mathbf{b}^{(x)}$. ◀

▶ **Theorem 8.** *Algorithm 1 computes the minimal clearing state of a given financial network without default cost in time* $O((n + k) \cdot (n^3 + m))$.

Proof. Lemma 2 shows that Algorithm 1 computes the correct minimal clearing state. Regarding the running time, computing the active graph can be done in time $O(n + m)$. Finding a reachable, non-singleton, sink-SCC (or verifying that none exists) can be done in time $O(n + m)$ using standard depth-first-search methods. Solving LPs (5) and (6) requires polynomial time. The running times for this are dominated by the computation of the eigenvector $\mathbf{d}^*$ and solving linear equations for vector $\mathbf{s}$, respectively. Each of these requires at most time $O(n^3)$. Upon solution of any of these two LPs, we advance to a new phase or meet the desired external assets of a bank. Thus, the number of times we need to solve an LP is upper bounded by $O(n + k)$. ◀

3.2 Default Cost

We now turn to the extension of Algorithm 1 to banks with default cost. We explain and justify the adjustments made in lines 2-3, 17 and lines 13-16 of the Algorithm 1.

▶ **Lemma 9.** *For each financial network* $\mathcal{F}$ *with minimal clearing state* $\hat{\mathbf{a}}$ *and default cost, there is a network* $\mathcal{F}'$ *without default cost with a minimal clearing state* $\hat{\mathbf{a}}'$ *equivalent to* $\hat{\mathbf{a}}$.

Proof. We implement default cost reductions of all banks in D by adjusting the network $\mathcal{F}$.

A bank with $\rho_u = 0$ makes no payments. Hence, we can w.l.o.g. assume it has no outgoing edges in $\mathcal{F}$. Omitting the default cost reduction for the assets of u is inconsequential for the remaining clearing state.

Now consider bank u with $0 < \rho_u < 1$ that is insolvent in $\hat{\mathbf{a}}$. We adjust the network $\mathcal{F}$ to $\mathcal{F}'$ as follows. Create an auxiliary sink bank u_s and an edge $e = (u, u_s)$ with liability $\ell_e = (1 - \rho_u)L^+(u)$. Define a new payment function $p'_e(a) = (1 - \rho_u) \cdot a$ for $a < L^+(u)$ and $p_e(a) = \ell_e$ for $a \geqslant L^+(u)$. For each other edge $e' = (u, v)$ in the network, we set the new payments to $p'_e(a) = p_e(\rho_u \cdot a)$. All other banks and payments remain as in $\mathcal{F}$. We assume that u has no default cost in $\mathcal{F}'$.

In $\mathcal{F}'$, instead of reducing the assets of u, the default cost reduction is directly incorporated into the payment functions, and the remaining assets are pushed to an auxiliary sink. It is straightforward to verify that $\hat{\mathbf{a}}' = (\hat{a}_u/\rho_u, \hat{\mathbf{a}}_{-u})$ is a clearing state in $\mathcal{F}'$ that is equivalent to $\hat{\mathbf{a}}$. In particular, $p'_e(a/\rho_u) = p_e(a)$ for each (original) edge $e = (u, v) \in E$. Any smaller clearing state $\mathbf{b}'$ in $\mathcal{F}'$ can be mapped back to a smaller clearing state $\mathbf{b} = (\rho_u b_u, \mathbf{b}_{-u})$ for $\mathcal{F}$. This proves the statement for insolvent banks.

For a bank $u \in D$ that is solvent in $\hat{\mathbf{a}}$, we have $p_e(\hat{a}_u) = \ell_e$ for all $e = (u, v) \in E$. We can remove all outgoing edges $(u, v) \in E$ and raise the external assets of each out-neighbor v by $\ell_{(u,v)}$. This creates an equivalent network $\mathcal{F}'$ for which $\hat{\mathbf{a}}' = \hat{\mathbf{a}}$ remains the minimal clearing state. This proves the statement for solvent banks. ◀

▶ **Corollary 10.** *Using the adjustment in Lemma 9, Algorithm 1 computes the minimal clearing state of a given financial network with default cost in time* $O((k + n + m) \cdot (n^3 + m))$.

Proof. In the algorithm, all banks are insolvent initially. Thus, we adjust the network in lines 2-3 according to Lemma 9 and apply the algorithm in network $\mathcal{F}'$, where all banks of D have adjusted payment functions and auxiliary sinks. This shows that, in particular, no (insolvent) bank from D can be part of any non-singleton sink-SCC. Thus, the flooding step is relevant only for components of banks without default cost. If a bank $u \in D$ remains insolvent until the end of the algorithm, we revert the adjustment in line 17 and scale down the assets to $\rho_u \cdot b_u$ to represent the default cost reduction in the original network $\mathcal{F}$.

The representation becomes problematic when a bank $u \in D$ becomes solvent w.r.t. the liabilities in $\mathcal{F}$. Then the external plus incoming assets of u equal the liabilities in $\mathcal{F}$ (i.e., without auxiliary edge (u, u_s)). We adjust the network in two steps: (1) Remove the auxiliary sink bank and increase the total payments towards the (original) outgoing edges by $\rho_u \cdot L^+(u)$. All outgoing edges $e \in E^+(u)$ become fully paid. By Lemma 9 we (2) remove them and raise the external assets of each out-neighbor by ℓ_e.

We implement these two steps in lines 13-16 as follows. Once a bank $u \in D$ becomes solvent (w.r.t. the original unadjusted network $\mathcal{F}$), by Lemma 2 it is solvent in $\hat{\mathbf{a}}$. Thus, we adjust $\mathbf{a}^{(x)}$ by removing the edges and adding external assets of $\ell_{(u,v)}$ to $a_v^{(x)}$. We execute a similar adjustment for $\mathbf{b}$ and $\mathbf{b}^{(x)}$: Add the *current payments* $p_{(u,v)}(b_u)$ to $b_v^{(x)}$ for each $(u, v) \in E$. This gives an equivalent representation of $\hat{\mathbf{a}}$ and $\mathbf{b}$ after removal of the out-edges of u. The remaining increase of $\rho_u L^+(u)$ in payments then gets executed via an increase in external assets of the (former) out-neighbors of u. This aligns directly with the invariant of Lemma 2 and proves that the algorithm remains correct for networks with default cost.

The upper bound on the running time in Theorem 8 suffers by at most m additional iterations due to an increase of some external assets by the liability of an incoming edge. ◀

3.3 Characterization and Maximal Clearing States

In this section, we briefly discuss how to compute the *maximal* clearing state $\check{\mathbf{a}}$. More generally, we show an equivalence of financial networks with general piecewise-linear payment functions to networks with *priority-proportional* functions.

▶ **Definition 11.** *For a bank* $v \in V$, *a collection of* priority-proportional *payment functions* $(p_e)_{e \in E^+(v)}$ *with* $p_e : \mathbb{R} \to \mathbb{R}$ *is given by a partition of* $E^+(v)$ *into* $k_v \geqslant 1$ *sets* $(E_1, \ldots, E_{k_v})$ *such that* $E_i \cap E_j = \varnothing$ *for* $i \neq j$, $\bigcup_i E_i = E^+(v)$, *and for every* $i = 1, \ldots, k_v$ *and* $e \in E_i$,

$$p_e(a) = \begin{cases} 0 & \textit{if } a < x_{v,i-1} \\ \frac{\ell_e}{\sum_{e' \in E_i} \ell_{e'}} \cdot (a - x_{v,i-1}) & \textit{if } a \in [x_{v,i-1}, x_{v,i}) \\ \ell_e & \textit{if } a \geqslant x_{v,i}, \end{cases}$$

where $x_{v,i} = \sum_{j \leqslant i} \sum_{e' \in E_j} \ell_{e'}$ *and* $x_{v,k_v+1} = \infty$.

Using priority-proportional payment functions, v clusters its outgoing edges into sets of decreasing priority. It starts to pay all edges of class E_1 proportionally until they are fully paid. Then it uses any remaining assets to pay class E_2 proportionally, and so on. Our main insight is that in terms of clearing, the class of networks with priority-proportional functions is equivalent to the class of all networks with piecewise-linear ones.

▶ **Lemma 12.** *For any financial network $\mathcal{F}$ with piecewise-linear payment functions, there is a network $\mathcal{F}'$ with priority-proportional payment functions such that $\mathbf{a}$ is a clearing state in $\mathcal{F}$ if and only if $\mathbf{a}$ is a clearing state in $\mathcal{F}'$.*

The lemma yields a simple polynomial-time transformation to obtain the network $\mathcal{F}'$ with $n + O(km)$ banks and $O(km)$ edges. For $\mathcal{F}'$ the algorithm from [22] computes $\check{\mathbf{a}}$ efficiently.

▶ **Corollary 13.** *There is a polynomial-time algorithm to compute the maximal clearing state in every financial network with piecewise-linear payment functions.*

4 Claims Trades for Minimal Clearing States

In this section, we consider the properties of *claims trading*, a decentralized network adjustment, when applied with *minimal* clearing. In claims trading, a bank w can buy an edge (u, v) by transferring some of its external assets to v. Throughout this section, we focus on networks without default cost, i.e., $\rho_v = 1$ for all $v \in V$.

We start by showing that in a claims trade, it is impossible that *both* creditor v and buyer w *strictly* improve their assets. We prove this result for financial networks in which all payment functions are *strictly* monotone, i.e., $p_e(a_u) < p_e(a_u + \varepsilon)$ for every $\varepsilon \in (0, \ell_e - a_u]$. Our proof uses the following lemma, which states novel properties of uniqueness for clearing states in such networks.

▶ **Lemma 14.** *Consider a bank v in a financial network $\mathcal{F}$ without default cost and with strictly monotone payment functions. If the clearing states are not unique w.r.t. v (i.e., $\hat{a}_v \neq \check{a}_v$), then $\hat{a}_v = 0$.*

Proof. Let $d_u = \check{a}_u - \hat{a}_u$ be the difference in payments between the greatest and least fixed point. Moreover, let $d_e = p_e(\check{a}_u) - p_e(\hat{a}_u)$ for every edge $e = (u, w) \in E$. Consider any bank v with $d_v > 0$. Then every node u reachable from v in the active graph $G_{\hat{\mathbf{a}}}$ must have $d_u > 0$. Moreover, d represents a circulation, so $\sum_{e \in E^+(u)} d_e = d_u = \sum_{e \in E^-(u)} d_e$ for every bank $u \in V$. This shows that at the end of Algorithm 1, v must be part of a non-sink SCC. Since all payment functions are *strictly* monotone, the active graph $G_{\mathbf{b}}$ is monotonically getting sparser during the execution of Algorithm 1 due to solvency of banks. Moreover, since there is no default cost, the set of sinks in the network is monotonically growing. Thus, if v is part of a non-sink SCC in the end, it must be part of a non-sink SCC throughout the entire execution of the algorithm. However, once assets of v get raised, all non-sink SCCs reachable from v must be flooded. As a consequence, v can only be part of such an SCC at the end of the algorithm when $\hat{a}_v = 0$. ◀

▶ **Theorem 15.** *Consider a financial network $\mathcal{F}$ without default cost and with strictly monotone payment functions. There exists no claims trade that strictly improves the assets of* both *a buyer w and a creditor v w.r.t. the minimal clearing state.*

Proof. Consider a claims trade with claim (u, v) and buyer w. In order to pay any return, w must have external assets $a_w^{(x)} > 0$. By Lemma 14, this implies $\hat{a}_w = \check{a}_w$.

First, suppose that $\hat{a}_v < \check{a}_v$. Then, Lemma 14 implies $\hat{a}_v = 0$. As discussed in the proof of the lemma, v is part of a non-sink SCC in $G_{\hat{\mathbf{a}}}$. Moreover, w must be unreachable from v in $G_{\hat{\mathbf{a}}}$. Since all payment functions are strictly monotonic, the active graph becomes only sparser for larger payments. Hence, w remains unreachable from v even when v has higher assets. Consequently, it is impossible for w to recover any portion of the return ρ paid to v by larger incoming payments. Thus, there cannot be a creditor-positive trade when $\hat{a}_v < \check{a}_v$.

Second, suppose that $\hat{a}_v = \check{a}_v$. Now, if there was such a trade, w.r.t. $\hat{\mathbf{a}}$, this trade would also strictly improve the assets of both parties w.r.t. $\check{\mathbf{a}}$. This is impossible [16]. ◀

4.1 Computing Optimal Creditor-Positive Trades

We focus on computing *creditor-positive* trades in this subsection. The main result is that existence of such trades can be decided in polynomial time for (even non-strictly) monotone, piecewise-linear payment functions. Moreover, an optimal creditor-positive return can be computed in polynomial time.

Even if there exists a creditor-positive trade, it is not directly obvious that an *optimal* return must exist. For example, consider the related problem of *cash injection*. The goal is to allocate M external assets in the network to maximize the total assets of all banks. For networks with proportional payments, an optimal solution can be computed in polynomial time when the network is evaluated by the *maximal* clearing state [23]. For *minimal* clearing states, it is easy to see that there are simple networks where no optimal solution exists[4].

We first analyze the structure of the set of creditor-positive returns. Let $\hat{\mathbf{a}}$ be the pre-trade minimal clearing state and $\rho_{\min} = p_e(\hat{a}_u)$ be the pre-trade payment on claim $e = (u, v)$.

▶ **Lemma 16.** *The set of creditor-positive returns forms a (possibly empty) interval* $(\rho_{\min}, \rho^*]$.

Based on this structural insight, we proceed to show the main result of this section.

▶ **Theorem 17.** *Consider a financial network $\mathcal{F}$ without default cost and with piecewise-linear payment functions. For a given claim $e = (u, v)$ and a buyer w, it can be decided in polynomial-time if a creditor-positive trade exists. If the trade exists, the optimal creditor-positive return can be computed in polynomial time.*

Proof. Suppose we use a return of $\rho_{\min}$. As observed above, the resulting minimal clearing state is $\hat{\mathbf{b}} = \hat{\mathbf{a}}$. Consider the active graph $G_{\hat{\mathbf{b}}}$. Let us increase the return ρ by a sufficiently small value δ and denote the resulting minimal clearing state by $\hat{\mathbf{b}}^{(\delta)}$.

First, suppose we do not pass an interval border on any edge, i.e., $G'_{\hat{\mathbf{b}}^{(\delta)}} = G_{\hat{\mathbf{b}}}$. Then the effect on the minimal clearing state must be linear, i.e., the change in the assets is given by

$$\mathbf{d} = \hat{\mathbf{b}}^{(\delta)} - \hat{\mathbf{b}} = \delta \cdot \mathbf{e}_v - \delta \cdot \mathbf{e}_w + \mathbf{d}\,\mathbf{M},$$

where $\mathbf{e}_v$ and $\mathbf{e}_w$ are unit vectors with an entry of 1 for v and w, respectively, and $\mathbf{M}$ is the matrix of all slopes of edges in $G_{\hat{\mathbf{b}}}$ (c.f. (7) above). $\mathbf{d}$ is linear in δ with slopes

$$\mathbf{s} = (\mathbf{I} - \mathbf{M})^{-1} \cdot (\mathbf{e}_v - \mathbf{e}_w).$$

---

[4] Suppose the financial network has two components; a cycle of 2 banks and a path of $n-2$ banks. All edges have liability 1, all banks have no external assets. Suppose we want to inject $M = 1$. Initially, $\hat{\mathbf{a}} = \mathbf{0}$. Assigning $1 - \varepsilon$ to the head of the path and ε to the banks in the cycle yields total assets of $(n-2)(1-\varepsilon) + 2 + \varepsilon$. This expression is maximized for $\varepsilon = 0$, but it applies only when $\varepsilon \in (0, 1]$. For $\varepsilon = 0$, the total assets in $\hat{\mathbf{a}}$ drop to $n - 2$. As such, there is no optimal solution.

As observed in the proof of Lemma 16, we have $s_w \leqslant 0 \leqslant s_v$. If $s_w = 0$, we can increase the return and it will stay creditor-positive. Given that the assets of w remain a_w, this Pareto-improves all assets in the network (and, thus, $\mathbf{s} \geqslant \mathbf{0}$ coordinate-wise).

An increase can be implemented very similarly as in Algorithm 1. We adjust all payments linearly until we reach an interval border for some payment function.

Second, suppose $G_{\hat{\mathbf{b}}}$ passes an interval border upon increase of ρ. More precisely, there is a positive slope for v and slope 0 for w, and raising the return requires a change in the active graph. If any larger creditor-positive return exists, it would further raise the assets of v (and keep the assets of w at a_w). This means that we first have to apply the flooding operation on all non-sink SCCs reachable from v as in Algorithm 1 above. Then, we again determine the slopes of further increase as before and check if the slope of w remains 0 or becomes negative.

We check the existence of a creditor-positive return as follows: Compute the slopes $\mathbf{s}$ for $G_{\hat{\mathbf{b}}}$. The initial slopes must be $s_v > 0 = s_w$, otherwise no creditor-positive return exists. If they are, and $G_{\hat{\mathbf{b}}}$ is located at an interval border, flood the appropriate SCCs and check the slopes again. If they continue to be $s_v > 0 = s_w$, a creditor-positive return exists; otherwise not.

If there is a creditor-positive return, we can search iteratively by increasing the return, changing the active graph, and flooding components as long as the resulting slopes are $s_v > 0 = s_w$. This approach repeatedly requires to solve systems of linear equations. It is very similar to Algorithm 1, and we obtain the same asymptotic upper bound on the running time. Alternatively, we can binary search over the interval $[\rho_{\min}, \min\{a_v^{(x)}, \ell_e\}]$. Once we find a return that is creditor-positive, we refine it by computing the slopes and increasing the return until the active graph hits the next interval border. This approach is faster if ρ^* is large and there are many interval borders for small creditor-positive returns. It can be slower if there are few interval borders and $\rho^* \ll \min\{a_v^{(x)}, \ell_e\}$. ◀

References

1 Hamed Amini and Zachary Feinstein. Optimal network compression. *Europ. J. Oper. Res.*, 306(3):1439–1455, 2023. `doi:10.1016/J.EJOR.2022.07.026`.

2 Peter Azimzadeh, Parsiad; Forsyth. Weakly chained matrices, policy iteration, and impulse control. *SIAM J. Numer. Anal.*, 54(3):1341–1364, 2016. `doi:10.1137/15M1043431`.

3 Nils Bertschinger, Martin Hoefer, Simon Krogmann, Pascal Lenzner, Steffen Schuldenzucker, and Lisa Wilhelmi. Equilibria and convergence in fire sale games. In *Proc. 22nd Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*, pages 215–223, 2023. `doi:10.5555/3545946.3598640`.

4 Nils Bertschinger, Martin Hoefer, and Daniel Schmand. Flow allocation games. *Math. Oper. Res.*, 50(1):68–89, 2025. `doi:10.1287/MOOR.2022.0355`.

5 Simina Brânzei, Reed C. Phillips, and Nicholas J. Recker. Tarski lower bounds from multi-dimensional herringbones. In *Proc. 2025 Workshop Approx. Algorithms for Comb. Opt. Problems (APPROX)*, pages 52:1–52:12, 2025. `doi:10.4230/LIPIcs.APPROX/RANDOM.2025.52`.

6 Xi Chen, Yuhao Li, and Mihalis Yannakakis. Reducing Tarski to Unique Tarski (in the black-box model). In *Proc. 38th Comput. Complex. Conf. (CCC)*, pages 21:1–21:23, 2023. `doi:10.4230/LIPIcs.CCC.2023.21`.

7 United States Courts. Bankruptcy basics. https://www.uscourts.gov/court-programs/bankruptcy/bankruptcy-basics. [Online; accessed March 27, 2025].

8 Péter Csóka and Jean-Jacques Herings. Decentralized clearing in financial networks. *Manag. Sci.*, 64(10):4681–4699, 2018. `doi:10.1287/MNSC.2017.2847`.

9 Péter Csóka and Jean-Jacques Herings. Uniqueness of clearing payment matrices in financial networks. *Math. Oper. Res.*, 49(1):232–250, 2024. `doi:10.1287/MOOR.2023.1354`.

10 Simon Dohn, Kristoffer Arnsfelt Hansen, and Asger Klinkby. Improved hardness results for the clearing problem in financial networks with credit default swaps. In *Proc. 18th Symp. Algorithmic Game Theory (SAGT)*, pages 81–98, 2025. `doi:10.1007/978-3-032-03639-1_5`.

11 Béni Egressy, Andreas Plesner, and Roger Wattenhofer. What is the price for lending in financial networks? In *Proc. 25th Int. Conf. Princ. Pract. Multi-Agent Syst. (PRIMA)*, pages 120–135, 2024. `doi:10.1007/978-3-031-77367-9_11`.

12 Larry Eisenberg and Thomas Noe. Systemic risk in financial systems. *Manag. Sci.*, 47(2):236–249, 2001. `doi:10.1287/MNSC.47.2.236.9835`.

13 Helmut Elsinger. Financial networks, cross holdings, and limited liability. Working Papers 156, Oesterreichische Nationalbank (Austrian Central Bank), May 2009.

14 Henri Froese, Martin Hoefer, and Lisa Wilhelmi. Dynamic debt swapping in financial networks. In *Proc. 4th Symp. Algorith. Found. Dynamic Networks (SAND)*, pages 2:1–2:16, 2025. `doi:10.4230/LIPIcs.SAND.2025.2`.

15 Martin Hoefer, Lars Huth, and Lisa Wilhelmi. Fractional claims trades and donations in financial networks. *CoRR*, abs/2502.06515, 2025. `doi:10.48550/arXiv.2502.06515`.

16 Martin Hoefer, Carmine Ventre, and Lisa Wilhelmi. Algorithms for claims trading. In *Proc. 41st Symp. Theoret. Aspects Comput. Sci. (STACS)*, pages 42:1–42:17, 2024. `doi:10.4230/LIPIcs.STACS.2024.42`.

17 Martin Hoefer and Lisa Wilhelmi. Seniorities and minimal clearing in financial network games. In *Proc. 15th Symp. Algorithmic Game Theory (SAGT)*, pages 187–204, 2022. `doi:10.1007/978-3-031-15714-1_11`.

18 Stavros Ioannidis, Bart de Keijzer, and Carmine Ventre. Strong approximations and irrationality in financial networks with derivatives. In *Proc. 49th Int. Colloq. Autom. Lang. Programming (ICALP)*, pages 76:1–76:18, 2022. `doi:10.4230/LIPIcs.ICALP.2022.76`.

19 Stavros Ioannidis, Bart de Keijzer, and Carmine Ventre. Clearing financial networks with derivatives: From intractability to algorithms. *CoRR*, abs/2312.05139, 2023. `doi:10.48550/arXiv.2312.05139`.

20 Stavros Ioannidis, Bart de Keijzer, and Carmine Ventre. Financial networks with singleton liability priorities. *Theor. Comput. Sci.*, 963:113965, 2023. `doi:10.1016/J.TCS.2023.113965`.

21 Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. Debt transfers in financial networks: Complexity and equilibria. In *Proc. 22nd Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*, pages 260–268, 2023. `doi:10.5555/3545946.3598645`.

22 Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. On priority-proportional payments in financial networks. *Theor. Comput. Sci.*, 1014:114767, 2024. `doi:10.1016/J.TCS.2024.114767`.

23 Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. Optimal bailouts and strategic debt forgiveness in financial networks. *Artif. Intell.*, 349:104424, 2025. `doi:10.1016/J.ARTINT.2025.104424`.

24 Pál András Papp and Roger Wattenhofer. Network-aware strategies in financial systems. In *Proc. 47th Int. Colloq. Autom. Lang. Programming (ICALP)*, pages 91:1–91:17, 2020. `doi:10.4230/LIPIcs.ICALP.2020.91`.

25 Pál András Papp and Roger Wattenhofer. Debt swapping for risk mitigation in financial networks. In *Proc. 22nd Conf. Econ. Comput. (EC)*, pages 765–784, 2021. `doi:10.1145/3465456.3467638`.

26 Pál András Papp and Roger Wattenhofer. Sequential defaulting in financial networks. In *Proc. 12th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, pages 52:1–52:20, 2021. `doi:10.4230/LIPIcs.ITCS.2021.52`.

27 Leonard Rogers and Luitgard Veraart. Failure and rescue in an interbank network. *Manag. Sci.*, 59(4):882–898, 2013. `doi:10.1287/MNSC.1120.1569`.

28 Steffen Schuldenzucker and Sven Seuken. Portfolio compression in financial networks: Incentives and systemic risk. In *Proc. 21st Conf. Econ. Comput. (EC)*, page 79, 2020. `doi:10.1145/3391403.3399538`.

29 Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Finding clearing payments in financial networks with credit default swaps is PPAD-complete. In *Proc. 8th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, pages 32:1–32:20, 2017. `doi:10.4230/LIPIcs.ITCS.2017.32`.

30 Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Default ambiguity: Credit default swaps create new systemic risks in financial networks. *Manag. Sci.*, 66(5):1981–1998, 2020. `doi:10.1287/MNSC.2019.3304`.

31 P. N. Shivakumar and Kim Ho Chew. A sufficient condition for nonvanishing of determinants. *Proc. Amer. Math. Soc.*, 43(1):63–66, 1974.

32 Jérôme Henry Stéphane Dees and Reiner Martin. Stamp€: Stress-test analytics for macroprudential purposes in the euro area, 2017. URL: `https://www.ecb.europa.eu/press/conferences/shared/pdf/20170511_2nd_mp_policy/DeesHenryMartin-Stampe-Stress-Test_Analytics_for_Macroprudential_Purposes_in_the_euro_area.en.pdf`.

33 Jinyun Tong, Bart de Keijzer, and Carmine Ventre. Selfishly cancelling debts can reduce systemic risk. In *Proc. 27th European Conf. Aritf. Intell. (ECAI)*, pages 3397–3404, 2024. `doi:10.3233/FAIA240890`.

34 Jinyun Tong, Bart De Keijzer, and Carmine Ventre. Reducing systemic risk in financial networks through donations. In *Proc. 27th European Conf. Aritf. Intell. (ECAI)*, pages 3405–3412, 2024. `doi:10.3233/FAIA240891`.

35 Luitgard Veraart. When does portfolio compression reduce systemic risk? *Math. Finance*, 32(3):727–778, 2022.

36 Hao Zhou, Yongzhao Wang, Konstantinos Varsos, Nicholas Bishop, Rahul Savani, Anisoara Calinescu, and Michael Wooldridge. Selfishly prepaying in financial credit networks. *J. Artif. Intell. Res.*, 81:877–906, 2024. `doi:10.1613/JAIR.1.16527`.