Algorithms for Claims Trading

Martin Hoefer
Goethe University Frankfurt, Germany

Carmine Ventre
King's College London, UK

Lisa Wilhelmi
Goethe University Frankfurt, Germany

Abstract

The recent banking crisis has again emphasized the importance of understanding and mitigating systemic risk in financial networks. In this paper, we study a market-driven approach to rescue a bank in distress based on the idea of *claims trading*, a notion defined in Chapter 11 of the U.S. Bankruptcy Code. We formalize the idea in the context of the seminal model of financial networks by Eisenberg and Noe [5]. For two given banks v and w, we consider the operation that w takes over some claims of v and in return gives liquidity to v (or creditors of v) to ultimately rescue v (or mitigate contagion effects). We study the structural properties and computational complexity of decision and optimization problems for several variants of claims trading.

When trading incoming edges of v (i.e., claims for which v is the creditor), we show that there is no trade in which *both banks v and w strictly* improve their assets. We therefore consider *creditor-positive* trades, in which v profits strictly and w remains indifferent. For a given set C of incoming edges of v, we provide an efficient algorithm to compute payments by w that result in a creditor-positive trade and maximal assets of v. When the set C must also be chosen, the problem becomes weakly NP-hard. Our main result here is a bicriteria FPTAS to compute an approximate trade, which allows for slightly increased payments by w. The approximate trade results in nearly the optimal amount of assets of v in any exact trade. Our results extend to the case in which banks use general monotone payment functions to settle their debt and the emerging clearing state can be computed efficiently.

In contrast, for trading outgoing edges of v (i.e., claims for which v is the debtor), the goal is to maximize the increase in assets for the creditors of v. Notably, for these results the characteristics of the payment functions of the banks are essential. For payments ranking creditors one by one, we show NP-hardness of approximation within a factor polynomial in the network size, in both problem variants when the set of claims C is part of the input or not. Instead, for payments proportional to the value of each debt, our results indicate more favorable conditions.

2012 ACM Subject Classification Theory of computation → Theory and algorithms for application domains; Networks → Network algorithms

Keywords and phrases Financial Networks, Claims Trade, Systemic Risk

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.42

Related Version *Full Version*: https://arxiv.org/pdf/2402.13627.pdf

Funding Martin Hoefer and Lisa Wilhelmi gratefully acknowledge support by DFG Research Unit ADYN under grant DFG 411362735.

Acknowledgements This paper was initiated at the Dagstuhl seminar 22452 "Computational Social Dynamics". The authors thank the organizers and participants of the seminar.

1 Introduction

The global banking crisis of March 2023 caused turmoil in a market fearful of the repeat of the Great Financial Crisis of 2007. These recent events serve as a stark reminder of the paramount importance of the study of systemic risk in financial networks. In this growing

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).
Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov; Article No. 42; pp. 42:1–42:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

body of work, the focus is mainly on the complexity of computing *clearing states*, known to measure the exposure of the different banks in the network to insolvencies within, see, e.g., [5, 10, 23], and strategic aspects of the banks' behavior, cf. [1, 8, 11, 17]. However, to calm the market and prevent contagion, regulators and central banks are more interested in finding ways to rescue banks in distress, reassure investors that the system is stable and avoid further bank runs. In fact, Silicon Valley Bank, Signature Bank and Credit Suisse – the three banks at the heart of the crisis last March – were all acquired by other banks in the network, and, by modifying the network, this has seemingly mitigated systemic risk.

A line of research in financial networks on interventions in the network is recently discussed in [7, 18], the main idea being that banks can swap debt contracts. In particular, the authors of [7] study the extent to which a sequence of debt swaps can reduce the risk in the network, in the sense that bank assets Pareto-improve. Notably, swaps can occur anywhere in the network, even if the focus is strict improvement of the assets of a given bank.

In this work, we build on this idea and initiate research on the computation of a network-based "rescue package" deal for a given bank with the objective of making it solvent. This is exactly the problem that regulators faced in March 2023 for the aforementioned banks. However, *acquisitions* do not seem to be the right operations in these instances since they have two main drawbacks from a societal perspective (as also witnessed by the reactions to recent deals). Firstly, the acquiring bank rarely has enough time or freedom to evaluate the purchase and make a sensible business decision. Secondly, and consequently, it often requires a security for bailout from the central bank, in the form of significant protection against potential losses from risks associated with the transaction. For example, in the acquisition of Credit Suisse, UBS had little choice in the matter, as reported by Bloomberg news [2], and received a guarantee worth CHF 9 billion, as confirmed by the Swiss Federal Council [3].

We instead study a market-driven approach to rescue banks in distress based on the idea of *claims trading.* Claims trading is defined in Chapter 11 of the U.S. Bankruptcy Code. We formalize the idea and analyze the consequences of such trades in the context of financial networks. When a company is in financial distress, its creditors can assert their rights to repayment by submitting a claim. At this point, a creditor can either wait for the positions to unwind and get (a part of) the claim once the bankruptcy is settled, or she can sell her credit claim to a willing buyer for some immediate liquidity. The former approach is equivalent to the mainstream work on systemic risk since the insolvency of a bank can directly cascade through the network via lower payments to its creditors. We want to explore ways to find interested buyers that purchase the claims of an insolvent bank v and give liquidity to the network that ultimately rescues v. Ideally, the buyers should avoid any loss so that the cash invested in buying the claim will return via increased payments within the network; this way incentives of buyers are aligned, and systemic risk is reduced at no extra cost to the network.

We design efficient algorithms to compute claims trades or settle the inherent complexity status of the problems. The importance of algorithms computing claim trades that resolve complicated systemic issues in finance cannot be underestimated. In practice, deals are concocted when markets are closed, and algorithms that efficiently compute solutions in these pressurised situations become essential.

Related Work. Much of the work on systemic risk in financial networks, including ours, builds upon [5]. In this seminal work, the authors propose a model and prove existence and properties of clearing states. Moreover, they also provide a polynomial-time algorithm for their computation. The model in [5] has been extended along many dimensions by follow-up work; for example, the authors of [20] add default costs whereas financial derivatives are considered in [22]. Computation of clearing states for the latter model is studied in [9,10,23] for different notions of approximation and payment schemes adopted. The solution space of clearing states for financial networks with derivatives is studied in [19].

The study of strategic behavior in financial networks was initiated in [1], where banks are assumed to strategize in the way they allocate money to their creditors. A similar approach is used in [8,11]. A different model, featuring derivatives, and banks strategically donating money or cancelling debts is studied in [17]. The idea of cancelling debts is further explored in [12]. The authors of [12,16] consider computational complexity of computing optimal or approximate bailout policies from the central bank external to the network. In contrast, in our work all transfers of assets are intrinsic to the network, and the bank providing the assets must not be harmed. In [13], the authors study computational complexity of strategic changes to the underlying network via debt transfers.

Debt swapping is introduced in [18] – the authors focus more on the existence and properties of swaps with and without shocks to the system. As discussed above, the authors of [7] share goals that are somewhat similar to ours but use a different operation to update the network. A related line of work considers portfolio compression, an accounting operation by means of which all the cycles in the network are deleted. The effects on systemic risk of portfolio compression are studied in [21,24]. However, it is important to note that portfolio compression can lead to a worse outcome for banks that are not contained in the cycle [21] and consequently it is not clear why banks should accept to modify their balance sheets in this way, as argued empirically in [15].

To the best of our knowledge, ours is the first work to study claims trading in the analysis of financial networks. Claims trading in bankruptcy has been studied by law scholars, who for example argue that its effects in that context are variegated and nuanced in general [14] but do not concern the governance of the bankruptcy process [6].

Contribution. We focus on the elementary setting with one given bank v to save (e.g., Credit Suisse) and one bank w that may rescue it (e.g., UBS). We consider the following problem: Are there claims of v that can be sold to w so that v becomes solvent (i.e., after the claims trade, v can fully pay all its liabilities)? This problem gives rise to a suite of algorithmic questions, depending on the remit of the algorithmic decision, such as: How many claims are we allowed to trade? Which claims of v should we trade? What are the payments that must be transferred from w to v to make the trade worthwhile for w? Our treatment is steered by the following structural insight: We prove that it is impossible for *both v and w to strictly* profit from the claims trade. Accordingly, we restrict our attention to creditor-positive trades that strictly improve v without harming w.

For our first set of algorithmic results in Section 3, we fix one claim with creditor v to be traded with w. Does this represent a feasible (i.e., creditor-positive) trade? This can be decided by simply computing clearing states that determine the payment towards each debt in the network. The problem becomes interesting if we also determine the *haircut rate* $\alpha \in [0,1]$ of the trade – in order to provide liquidity, w may be willing to pay an α-fraction of the claim's liability to v. Depending on the payment functions used by banks to distribute money to their debts, we design different polynomial-time algorithms that determine feasibility of the trade and also α^* (if any), the value of α maximizing the assets of v (or a close approximation of α^* if the payment functions are too granular vis-a-vis the input size). Let us highlight that these results also apply to the case in which v is the *debtor* of the claim to trade; in fact, we prove that every creditor-positive trade Pareto-improves the clearing state – each bank in the network is (weakly) better off after the trade. By maximizing the assets of the creditor of the traded claim we, thus, also maximize assets of the debtor.

We consider trading multiple claims with creditor v in Section 4. For a fixed set of claims our results from Section 3 extend rather directly. The picture becomes less benign when we also have to choose the subset of claims to be traded. Indeed, in Section 4.2 we show that it is weakly NP-hard to decide if there is a subset of claims along with suitable haircut rates to obtain a creditor-positive claims trade that makes v solvent. In our most technical contribution, we show that there exists a bicriteria FPTAS for deciding this problem. If an exact trade exists that yields total assets of A^* for v, we find an approximate trade with assets at least $A^* - \delta$ for exponentially small δ, which allows haircut rates of at most $1 + \varepsilon$. The FPTAS applies to all financial networks with general monotone payment functions for which a clearing state can be computed efficiently. On a technical level, we fix a desired value A for the total assets of v. Using a subroutine we determine if there is an approximate trade that yields this asset level for v. En route, we discover an intricate monotonicity property – if there is an exact trade that yields assets A^* for v, then for every $A \leq A^*$ there is an approximate trade with assets A for v. Notably, monotonicity can break above A^*. Still, we can apply binary search to find an approximate trade with assets at least $A^* - \delta$.

Finally, for trading multiple claims with debtor u in Section 5 – rather than trying to save u – our goal is improve conditions for the creditors of u to minimize the contagion effects by u's bankruptcy. Interestingly, the results here depend significantly on the choice of the payment functions. For payments based on a ranking of the creditors, we show that the problem becomes NP-hard to approximate within a factor polynomial in the network size. In contrast, for payments proportional to the value of each debt, we can solve the problem for a given set of claims, but it becomes strongly NP-hard when having to choose the set of claims.

All missing proofs are deferred to the full version of this paper.

2 Model and Preliminaries

A financial network $\mathcal{F} = (G, \boldsymbol{\ell}, \mathbf{a}^x, \mathbf{f})$ is expressed as a directed multigraph[1] $G = (V, E, \mathrm{de}, \mathrm{cr})$ without self loops. We denote $n = |V|$. Every node $v \in V$ in the graph represents a financial institution or *bank*. Every edge $e \in E$ represents a *debt contract* or *claim* involving two banks. For each edge $e \in E$, $\mathrm{de}(e)$ specifies the debtor (i.e., the source) and $\mathrm{cr}(e)$ the creditor (i.e., target). Edge $e \in E$ has a weight $\ell_e \in \mathbb{N}_{>0}$. In other words, in the context of debt contract e, bank $\mathrm{de}(e)$ owes $\mathrm{cr}(e)$ an amount of ℓ_e. We denote the set of outgoing and incoming edges of a bank v by $E^+(v) = \{e \in E \mid v = \mathrm{de}(e)\}$ and $E^-(v) = \{e \in E \mid v = \mathrm{cr}(e)\}$. Since we allow multi-edges, several debt contracts with possibly different liabilities could exist between the same pair of banks. The *total liabilities* L_v of v are the sum of weights of all outgoing edges of v, i.e., $\sum_{e \in E^+(v)} \ell_e = L_v$. Furthermore, every bank v holds *external assets* $a_v^x \in \mathbb{N}$. They can be interpreted as an amount of money the bank receives from outside the network.

Let $b_v \in [a_v^x, a_v^x + \sum_{e \in E^-(v)} \ell_e]$ be the total funds of bank v. Bank v distributes her total funds according to a given *payment function* $\mathbf{f}_v = (f_e)_{e \in E^+(v)}$, where $f_e : \mathbb{R} \to [0, \ell_e]$. For every outgoing edge, the function $f_e(b_v)$ defines the amount of money v pays towards e. We follow previous literature and assume the following conditions for every payment function:

(1) Every function $f_e(b_v)$ is non-decreasing and bounded by $0 \leq f_e(b_v) \leq \ell_e$.
(2) Every bank pays all funds until all liabilities are settled: $\sum_{e \in E^+(v)} f_v(b_v) = \min\{b_v, L_v\}$.
(3) The sum of payments of a bank is limited by the total funds: $\sum_{e \in E^+(v)} f_v(b_v) \leq b_v$.

Here (2) implies (3), and we mention (3) explicitly for clarity. For a *monotone* function $\mathbf{f}_v$, v weakly increases the payment on every outgoing edge when receiving additional funds.

[1] Claims trades in simple graphs can result in graphs with multi-edges. This can sometimes be avoided by analyzing the trades in equivalent simple graphs with suitable auxiliary banks. Since all our arguments can also be applied in the context of multigraphs, we discuss the more general model.

Clearing States. Let $\mathbf{p} = (p_e)_{e \in E}$ be the arising payments in the network when every bank v distributes the funds according to her payment functions $\mathbf{f}_v$. The *incoming payments* of v are given by $\sum_{e \in E^-(v)} p_e$. The *total assets* a_v are defined as the external assets plus the incoming payments, i.e., $a_v = a_v^x + \sum_{e \in E^-(v)} p_e$. Observe that the above conditions (1), (2) and (3) are fixed-point constraints. A vector of total assets $\mathbf{a} = (a_v)_{v \in V}$ is called *feasible* if it satisfies all fixed-point constraints. More formally, for every feasible $\mathbf{a}$ it holds that $a_v = a_v^x + \sum_{e \in E^-(v)} f_e(a_{\mathrm{de}(e)})$. The payments $\mathbf{p}$ corresponding to a feasible vector $\mathbf{a}$ are called a *clearing state*. For fixed payment functions, multiple clearing states may exist. We assume throughout that every payment function $\mathbf{f}_v$ is *monotone*, i.e., $f_e(x) \geq f_e(y)$ for all $x \geq y \geq 0$ and every $e \in E^+(v)$. This implies that all clearing states form a complete lattice [1, 4]. Thus, the point-wise minimal and maximal clearing states are unique. We follow previous literature and assume that the *maximal* clearing state arises in the network.

Payment Functions. In the seminal work of Eisenberg and Noe [5] and the majority of subsequent works, all banks are assumed to allocate their assets using *proportional* payment functions. The *recovery rate* $r_v = \min\{a_v/L_v, 1\}$ is the fraction of total liabilities v can pay off, and the payments on edge $e \in E^+(v)$ are defined proportionally by $f_e(a_v) = r_v \cdot \ell_e$. Hence, if $r_v = 1$, then v will fully settle all liabilities. Otherwise, v is in default, $r_v < 1$, and the liabilities are settled partially in proportion to their weight. These payments are often used when all debt contracts fall due at the same date. If, on the other hand, different debt contracts are assigned different priorities or maturity dates, payments are more suitably expressed by *edge-ranking* payment functions. Then, the debt contracts in $E^+(v)$ are ordered by a permutation $\boldsymbol{\pi}_v$. First, v makes payments towards the highest ranked edge $\pi_v(1)$ until the edge is saturated or v has no remaining assets. Once $\pi_v(1)$ is fully paid off, v pays off the second highest ranked edge $\pi_v(2)$ until the edge is saturated or v has no remaining assets. The process continues and ends when either all liabilities are settled or v exhausted all assets.

Both proportional and edge-ranking payments are monotone. For proportional payments, the clearing state can be computed in polynomial time [5]; for edge-ranking payments in strongly polynomial time [1].

In this paper, we obtain some results explicitly for networks with proportional and edge-ranking payment functions. Most of our results, however, generalize to arbitrary monotone payment functions when there is an efficient *clearing oracle*, i.e., there exists an algorithm that receives a network $\mathcal{F}$ as input and outputs the clearing state $\mathbf{p}$ in polynomial time.

Claims Trades. When a bank u is in default and unable to settle all debt, this introduces risk into the network. In particular, the creditors of u do not receive their full liabilities. This could lead to further defaults of the creditor banks. In order to reduce the risk of spreading default, the creditors of u can sell claims they make towards u. More in detail, consider banks u, v and w with edge e with $\mathrm{de}(e) = u$ and $\mathrm{cr}(e) = v$, $\ell_e \geq 0$. Suppose u is in default. If v and w perform a claims trade, w becomes the new creditor of bank u with the same liability. Consequently, any payment from u towards the claim will be received by w. In return for the traded claim, v receives a *return* ρ from w, i.e., an immediate payment of $\rho = \alpha \cdot \ell_e$, for some $\alpha \in [0, 1]$. We call α the *haircut rate*. To separate the return from the payments in the clearing state, we model the return by a transfer of external assets from w to v. Note that w can invest at most her external assets as a return, so every trade must satisfy $\alpha \ell_e \leq a_w^x$. After a trade the external assets of v and w might no longer be integer values.

We proceed to define three variants of claims trades. We are given a financial network with distinct banks $u, v, w \in V$, an edge $e \in E$ with $(\mathrm{de}(e), \mathrm{cr}(e)) = (u, v)$ and haircut rate $\alpha \in [0, 1]$. For a *(single) claims trade* of e to w we perform the following adjustments to

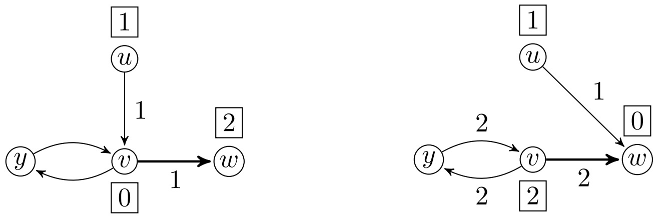

Figure 1 The network from Example 1 before the trade is depicted left, and right after the trade. All liabilities equal 2. Edges are labeled with positive payments (if any) in the clearing state.

the network: (1) change the creditor of e to $\mathrm{cr}'(e) = w$, (2) change external assets of v to $a^x_v + \alpha \cdot \ell_e$ and (3) change external assets of w to $a^x_w - \alpha \cdot \ell_e \geq 0$. We denote the resulting *post-trade network* by $\mathcal{F}' = (G', \boldsymbol{\ell}, \mathbf{a}'^x, \mathbf{f})$, and the resulting clearing state in $\mathcal{F}'$ by $\mathbf{p}'$. For a given trade of e to w, we call v the *creditor* and w the *buyer*. Observe that the total assets of v after the trade are given by $a'_v = a^x_v + \alpha \cdot \ell_e + \sum_{e' \in E'^-(v)} p'_{e'}$. Similarly, the total assets of w after the trade are $a'_w = a^x_w - \alpha \cdot \ell_e + \sum_{e' \in E'^-(w)} p'_{e'}$.

The claims trade operation can be directly extended to a trade of multiple edges. As outlined in the introduction, we are interested in the effects when a single bank (in default) trades claims with another bank w (such as a central bank). We study the differences when trading incoming or outgoing edges. Observe that both generalize single claims trades.

For a *multi-trade of incoming edges*, there are distinct banks v, w in a network $\mathcal{F}$, a set C of k distinct incoming edges $e_1, \ldots, e_k \in E^-(v)$, and haircut rates $\alpha_1, \ldots, \alpha_k$, such that $\mathrm{de}(e_i) \neq w$, for all i. After the trade, a new network $\mathcal{F}'$ emerges: We change $\mathrm{cr}'(e_i) = w$, for all $i = 1, \ldots, k$, adjust external assets for v to $a^x_v + \sum_{i=1}^k \alpha_i \ell_{e_i}$, and for w to $a^x_w - \sum_{i=1}^k \alpha_i \ell_{e_i} \geq 0$.

For a *multi-trade of outgoing edges*, there are distinct banks u, w in a network $\mathcal{F}$, a set C of k distinct outgoing edges $e_1, \ldots, e_k \in E^+(v)$ with $\mathrm{cr}(e_i) = v_i$, and haircut rates $\alpha_1, \ldots, \alpha_k$, such that $\mathrm{cr}(e_i) \neq w$, for all i. After the trade, a new network $\mathcal{F}'$ emerges: We change $\mathrm{cr}'(e_i) = w$, for all $i = 1 \ldots, k$, adjust external assets for each v_i to $a^x_{v_i} + \alpha_i \ell_{e_i}$, and for w to $a^x_w - \sum_{i=1}^k \alpha_i \ell_{e_i} \geq 0$.

We proceed with a small example of trading a single claim.

► **Example 1.** Consider the example network depicted in Figure 1 (left) on a simple directed graph. The liability of every edge is 2. The only banks with non-zero external assets are u and w, where $a^x_u = 1$ and $a^x_w = 2$. $a^x_v = 0$ is also explicitly displayed for convenience. Banks u, w and y each have at most one outgoing edge. They pay all their assets (if any) to the unique outgoing edge until it is saturated. This implies payments of 1 on edge (u, v). v is the only bank with a non-trivial payment function – suppose it uses an edge-ranking function with priority $\pi_v(1) = (v, w)$ and $\pi_v(2) = (v, y)$. Then, v pays the incoming assets of 1 to w, and there are no payments on the cycle of v and y. To see this, assume $p_{(y,v)} = x > 0$. By the edge-ranking function, from these additional assets v allocates a portion of $\min(x, 1)$ to (v, w) and the rest to (v, y). Hence, the total assets of y are $\max(x - 1, 0)$ while the outgoing payments are x, which contradicts the feasibility constraint (3). Overall, in the clearing state, the total assets are $a_u = a_v = 1, a_w = 3$ and $a_y = 0$.

Suppose we perform the trade of edge $e = (u, v)$ to w with $\alpha = 1$, see Fig. 1 (right) for the resulting network. w buys edge e and pays the full liability ℓ_e to v. The external assets of v increase to 2 and allow v to settle all debt. The total assets become 1 for u, 2 for y, 3 for w and 4 for v. The total assets of u and w are unaffected by the trade, the total assets of v and y strictly increase. Overall, the clearing state is point-wise non-decreasing.

A similar observation can be made when v uses proportional payments. Before the trade, v pays 1 to y and w. After the trade, both edges can be paid fully. ⌟

Properties of Claims Trades. In Example 1, v strictly benefits from the trade while w is indifferent. Interestingly, it is impossible for both creditor and buyer to strictly profit from a single trade. This property holds true for the more general class of multi-trades of incoming edges, and it applies in any network $\mathcal{F}$ with monotone payment functions.

The proof builds on a connection with *debt swaps* studied in [7,18]. A debt swap exchanges the creditors of two edges with the same liabilities. We show that a claims trade and the resulting payments can be represented by a debt swap in an auxiliary network.

▶ **Definition 2** (Debt Swap). *Consider a financial network $\mathcal{F}$ with four distinct nodes $u_1, u_2, v_1, v_2 \in V$ and edges $e_1, e_2 \in E$, where $u_1 = de(e_1), v_1 = cr(e_1)$ and $u_2 = de(e_2), v_2 = cr(e_2)$. Suppose the liabilities are $\ell_{e_1} = \ell_{e_2}$. A debt swap σ of e_1 and e_2 creates a new network $\mathcal{F}^\sigma$ with $G^\sigma = (V, E, de, cr^\sigma)$ where $cr^\sigma(e_1) = v_2$, $cr^\sigma(e_2) = v_1$ and $cr^\sigma(e) = cr(e)$ for all $e \in E \setminus \{e_1, e_2\}$.*

▶ **Proposition 3.** *For every financial network with monotone payment functions, there exists no multi-trade of incoming edges such that both creditor v and buyer w strictly improve their total assets.*

Proof. For a given network $\mathcal{F}$, consider a multi-trade of incoming edges and construct a new network $\hat{\mathcal{F}}$ by adding an auxiliary bank $\hat{v}$ to $\mathcal{F}$ without external assets. $\hat{v}$ serves as an "accumulator" for the payments along the edges e_i. We change the targets of the edges in C to $\mathrm{cr}(e_i) = \hat{v}$. We add an edge $\hat{e}$ with $\mathrm{de}(\hat{e}) = \hat{v}$ and $\mathrm{cr}(\hat{e}) = v$ and liability $\ell_{\hat{e}} = \sum_{i=1}^k \ell_{e_i}$. Every payment that gets paid to v via e_i in $\mathcal{F}$ now first goes to $\hat{v}$, and then gets forwarded from $\hat{v}$ to v, since $\hat{e}$ has sufficiently high liability. Consider the clearing state $\hat{\mathbf{p}}$ in the resulting network $\hat{\mathcal{F}}$. Obviously, for the new edge $\hat{p}_{\hat{e}} = \sum_{i=1}^k \hat{p}_{e_i}$. As such, every (non-auxiliary) bank from $\mathcal{F}$ receives the same external assets and eventually the same incoming and outgoing payments in $\hat{\mathcal{F}}$. Consequently, both $\mathcal{F}$ and $\hat{\mathcal{F}}$ give rise to the same clearing state, i.e., $p_e = \hat{p}_e$, for all $e \in E$, and the same assets for every (non-auxiliary) bank.

The new network $\hat{\mathcal{F}}$ allows to conveniently route all payments along edges in C to w by trading the single accumulator edge $\hat{e}$ to w. Thus, the multi-trade of incoming edges C in $\mathcal{F}$ is equivalent to trading the single claim $\hat{e}$ to w in $\hat{\mathcal{F}}$, for a suitably chosen haircut rate $\hat{\alpha}$ such that $\hat{\alpha} \cdot \sum_{i=1}^k \ell_{e_i} = \sum_{i=1}^k \alpha_i \cdot \ell_{e_i}$.

Now let us further adjust $\hat{\mathcal{F}}$ to $\tilde{\mathcal{F}}$ by introducing an auxiliary bank $\tilde{w}$. Intuitively, we "outsource" parts of external assets from w to $\tilde{w}$. Formally, external assets of w are reduced to $a_w^x - \sum_{i=1}^k \alpha_i \cdot \ell_{e_i} \geq 0$, external assets of $\tilde{w}$ are $\sum_{i=1}^k \alpha_i \cdot \ell_{e_i}$. We add an edge $\tilde{e}$ with $\mathrm{de}(\tilde{e}) = \tilde{w}$ and $\mathrm{cr}(\tilde{e}) = w$ as well as liabilities $\ell_{\tilde{e}} = \ell_{\hat{e}} = \sum_{i=1}^k \ell_{e_i}$. The clearing state $\tilde{\mathbf{p}}$ in the resulting network $\tilde{\mathcal{F}}$ is $\tilde{p}_{\tilde{e}} = \sum_{i=1}^k \alpha_i \ell_{e_i}$, since $\tilde{e}$ is the only outgoing edge of $\tilde{w}$ and $\ell_{\tilde{e}} \geq a_{\tilde{w}}^x$. Hence, w and (consequently) every non-auxiliary bank from $\mathcal{F}$ receives the same total assets in $\tilde{\mathbf{p}}$. Indeed, $\mathcal{F}$, $\hat{\mathcal{F}}$ and $\tilde{\mathcal{F}}$ yield equivalent clearing states with $p_e = \hat{p}_e = \tilde{p}_e$, for all $e \in E$.

In $\tilde{\mathcal{F}}$ we can implement the return payments from w to v by re-routing the "outsource" edge e' to v instead of w. Thus, the claim trade of $\hat{e}$ in $\hat{\mathcal{F}}$ can be expressed by a swap of creditors of $\hat{e}$ and $\tilde{e}$ in $\tilde{\mathcal{F}}$. Now since $\ell_{\tilde{e}} = \ell_{\hat{e}}$, this swap of creditors represents a debt swap. Thus, the multi-trade in $\mathcal{F}$ is equivalent to single trade in $\hat{\mathcal{F}}$ and a debt swap in $\tilde{\mathcal{F}}$. No debt swap can strictly improve both creditor banks [7, Corollary 6]. Thus, no multi-trade of incoming edges can strictly improve both creditor and buyer. ◀

The above proof implies a structural equivalence. Using the network $\hat{\mathcal{F}}$, we reduced a multi-trade of incoming edges to a single claim trade.

▶ **Corollary 4.** *For every multi-trade of incoming edges in a network $\mathcal{F}$, there is an adjusted network $\hat{\mathcal{F}}$ such that the result of the multi-trade in $\mathcal{F}$ is the result of a single trade in $\hat{\mathcal{F}}$.*

Our motivation is to analyze claims trades to improve the situation of a creditor in default by trading claims with a buyer. Since it is impossible to strictly improve the conditions of both banks, we focus on strictly improving the creditor and weakly improving the buyer. Note that the trade performed in Example 1 satisfies this property.

▶ **Definition 5** (Creditor-positive trade). *A multi-edge trade of incoming edges of bank v to bank w is called creditor-positive if $a'_v > a_v$ and $a'_w \geq a_w$.*

For the proof of Proposition 3, we express the multi-trade by a debt swap in an auxiliary network. For a creditor-positive trade, the associated debt swap satisfies the same property, i.e., it is a so-called *semi-positive* debt swap. In any network $\mathcal{F}$ with monotone payment functions, a semi-positive debt swap Pareto-improves the clearing state and, hence, the total assets of *every* bank [7]. This directly implies the next corollary.

▶ **Corollary 6.** *In every financial network with monotone payments, every creditor-positive trade Pareto-improves the clearing state.*

A creditor-positive trade reduces the impact of a defaulting debtor on the creditor. No bank in the entire network suffers. Hence, these trades contribute to the stabilization of the entire financial network. We focus on creditor-positive trades for the remainder of the paper.

3 Trading a Single Claim

In this section, we study a given single creditor-positive trade and optimize the effects on the assets in the network. For exposition, we mostly focus on financial networks with proportional or edge-ranking payments.

The choice of α affects the external assets of v and, thus, payments throughout the network. If a given trade is creditor-positive for some $\alpha \in [0, 1]$, we say that α is creditor-positive. Can we efficiently decide the existence of a creditor-positive α? What is the *optimal* α to maximize the improvement $a'_v - a_v$ of v? Clearly, a trade with optimal α maximizes the total assets a'_v. Since $a'_w = a_w$ in every creditor-positive trade, maximizing a'_v also maximizes the payments of v, the incoming payments of v's creditors, and, inductively, the payments and assets of every edge and bank in the network. A creditor-positive α that maximizes a'_v also simultaneously maximizes (1) the Pareto-improvement of payments for each edge in the network, and (2) the return ρ by w. This holds for all networks with monotone payment functions.

To answer the above questions, we modify $\mathcal{F}$ into a *return network* $\mathcal{F}^{\text{ret}}$ defined as follows. We switch edge e to $\text{cr}(e) = w$ and add a *return edge* e_r with $\text{de}(e_r) = w$ and $\text{cr}(e_r) = v$. The payment on this edge models the return from w to v, so the liability is $\ell_{e_r} = \min\{\ell_e, a_w^x\}$. Since we consider creditor-positive trades, we modify the payment function of w as follows. For all $e' \in E^+(w) \setminus \{e_r\}$, we set $f_{e'}^{\text{ret}}(x) = f_{e'}(x)$ for all $x \leq a_w$ and $f_{e'}^{\text{ret}}(x) = f_{e'}(a_w)$ for all $x \geq a_w$. For e_r we set $f_{e_r}^{\text{ret}}(x) = 0$ for all $x \leq a_w$ and $f_{e_r}^{\text{ret}}(x) = x - a_w$ for all $x \geq a_w$. Similarly, we modify the liabilities to $\ell_{e'} = f_{e'}(a_w)$. Intuitively, in $\mathcal{F}^{\text{ret}}$ w maintains its payments up to a total outgoing assets of a_w. It forwards any assets exceeding a_w as return via e_r to v.

▶ **Lemma 7.** *Consider the clearing state $\mathbf{p}^{ret}$ in $\mathcal{F}^{ret}$.*

(a) *Suppose there is an optimal creditor-positive α with return $\rho = \alpha\ell_e$. Then $\mathcal{F}^{ret}$ has $a_w^x > p_e$. In $\mathbf{p}^{ret}$ we obtain assets of $a_w^{ret} \in (a_w, a_w + \ell_{e_r}]$ and $a_v^{ret} > a_v$, and $p_{e_r}^{ret} = \rho$.*

(b) *If $a_w^x > p_e$ and $\mathbf{p}^{ret}$ yields assets of $a_w^{ret} \in (a_w, a_w + \ell_{e_r}]$ and $a_v^{ret} > a_v$, then payment p_{e_r} represents a return of an optimal creditor-positive trade.*

Proof. We first show (a). Suppose there is an optimal creditor-positive α. It results in a return $\rho = \alpha\ell_e \le \min\{\ell_e, a_w^x\} = \ell_{e_r}$, assets of a_w for w, and $a'_v > a_v$ for v. When we assign payments $\hat{p}_e = p'_e$ for all $e' \in E$ and set the payment on e_r to $\hat{p}_{e_r} = \rho$, we obtain a vector of payments $\hat{\mathbf{p}}$ in $\mathcal{F}^{\text{ret}}$ that satisfies all fixed-point conditions.

We first show that this implies $a_w^x > p_e$, the payment on e in $\mathbf{p}$ before the trade. Consider the assets of w. We have $\hat{a}_v = a'_v > a_v$. Recall $\mathbf{p}' \ge \mathbf{p}$ by Corollary 6, so

$$\begin{aligned}\hat{a}_v &= a_v^x + \rho + \sum_{e' \in E^-(v)\setminus\{e\}} \hat{p}_{e'} = a_v^x + \rho + \sum_{e' \in E^-(v)\setminus\{e\}} p'_{e'} \\ &\ge a_v^x + \rho + \sum_{e' \in E^-(v)\setminus\{e\}} p_{e'} = a_v + \rho - p_e\end{aligned}$$

Hence $a_w^x \ge \rho > p_e$, as desired.

For the other conditions, consider the clearing state $\mathbf{p}^{\text{ret}}$ in $\mathcal{F}^{\text{ret}}$. Due to maximality of the clearing state, $\mathbf{p}^{\text{ret}} \ge \hat{\mathbf{p}}$. Thus, $a_w^{\text{ret}} > a_w$, $a^{\text{ret}} > a_v$ and $p_{e_r}^{\text{ret}} \ge \rho$. We show that, indeed, $\hat{\mathbf{p}} = \mathbf{p}^{\text{ret}}$, and that the condition $a_w + \ell_{e_r} \ge \hat{a}_w$ is satisfied.

Case 1: The clearing state satisfies $a_w + \ell_{e_r} \ge a_w^{\text{ret}}$. Then we prove below that $\mathbf{p}^{\text{ret}}$ is equivalent to a creditor-positive trade with payments that Pareto-dominate $\hat{\mathbf{p}}$ and, consequently, higher assets for v with $\hat{a}_v \ge a_v^{\text{ret}}$. As such, $\mathbf{p}^{\text{ret}}$ represents a better creditor-positive trade, a contradiction to $\hat{\mathbf{p}}$ stemming from an optimal one.

Case 2: The clearing state satisfies $a_w + \ell_{e_r} < a_w^{\text{ret}}$. Then $a_v^{\text{ret}} > a_v$, and w is solvent in $\mathcal{F}^{\text{ret}}$. Indeed, w could transfer even more assets to the edges of $E^+(w) \setminus \{e_r\}$. This implies that with return ℓ_{e_r}, there is a clearing state in $\mathcal{F}'$ that can *strictly improve both* v and w. This is a contradiction to Corollary 3.

To prove (b), suppose $\mathbf{p}^{\text{ret}}$ fulfills the conditions. Then, clearly, the payment $p_{e_r}^{\text{ret}}$ represents a feasible return. The payments $p_{e'}^{\text{ret}}$ on the other edges $e' \in E \setminus \{e_r\}$ fulfill the fixed-point conditions in $\mathcal{F}'$. Now for contradiction assume that $p'_{e'} > p_{e'}^{\text{ret}}$ for some e'. Then $e' \ne e_r$, since we assume $p_{e_r}^{\text{ret}}$ is the return used to construct $\mathcal{F}'$. Hence, any strict increase in $\mathbf{p}'$ could be manifested in $\mathbf{p}^{\text{ret}}$ as well, which contradicts the maximality of $\mathbf{p}^{\text{ret}}$ in $\mathcal{F}^{\text{ret}}$. ◀

▶ **Corollary 8.** *Consider a given single claims trade of e to w.*

(a) *A return of $a_w^x \ge \rho > p_e$ is necessary to make the trade creditor-positive. For $\rho = p_e$, we obtain $\mathbf{p}' = \mathbf{p}$.*

(b) *Consider all creditor-positive α. A value α with return $\rho = \alpha\ell_e$ maximizes the assets of v if and only if it maximizes the payment on every single edge in $\mathcal{F}'$, the assets of every single bank, as well as the value of ρ and α.*

▶ **Proposition 9.** *For a given financial network with edge-ranking payments and a single claims trade, there is an efficient algorithm to compute an optimal creditor-positive $\alpha^* \in [0, 1]$ or decide that none exists.*

Proof. We construct network $\mathcal{F}^{\text{ret}}$ as described above. Observe that the adjusted payment function $\mathbf{f}_w^{\text{ret}}$ is again an edge-ranking function – it first fills edges according to $\mathbf{f}_w$ until assets a_w are paid. Thus, at most one edge $e' \in E^+(w)$ is paid partially. For this edge, the liabilities are reduced to $f_{e'}(a_w)$. For all other edges, the liabilities either remain untouched or are decreased to 0. Then the additional assets are directed to e_r. Thus, $\mathbf{f}_w^{\text{ret}}$ can be represented by the same ranking as $\mathbf{f}_w$ up to (and including) edge e', and then using e_r as the next (and last) edge in the order. Hence, we can compute $\mathcal{F}^{\text{ret}}$ in strongly polynomial time. By checking the conditions of Lemma 7, we can verify in polynomial time whether or not a creditor-positive trade exists and obtain the optimal return as the payment on e_r. ◀

▶ **Proposition 10.** *For a given financial network with proportional payments and a single claims trade, there is an efficient algorithm to compute an optimal creditor-positive* $\alpha^* \in [0, 1]$ *or decide that none exists.*

Finally, our main result in this section shows that for general monotone payments with efficient clearing oracle, we can obtain an approximately optimal solution via binary search.

▶ **Theorem 11.** *Consider a given financial network with monotone payment functions and efficient clearing oracle. For a given single claims trade, there exists an additive FPTAS for approximating the optimal improvement of* v *from any creditor-positive* α*.*

Our algorithm uses binary search. Towards this end, we first show, for a given target value $A \geq a_v$, how to verify the existence of a trade that achieves at least a value A for the total assets of v. For intuition, we use a *split network* $\mathcal{F}^{\text{sp}}$ obtained from $\mathcal{F}'$ after the trade as follows: We replace v and w by source and sink banks v_{in}, v_{out}, and w_{in}, w_{out}. v_{in} has the incoming edges of v, w_{in} the ones of w (including e). The outgoing edges of v (w) are attached to v_{out} (w_{out}). We set the external assets of v_{out} and w_{out} to A and a_w, and these banks use the payment functions of v and w, respectively. As such, the clearing state $\mathbf{p}^{\text{sp}}$ in $\mathcal{F}^{\text{sp}}$ can be computed using the clearing oracle.

Consider the incoming payments of $\mathbf{p}^{\text{sp}}$ at v_{in} and w_{in}. These payments shall exactly recover the expenses at v_{out} and w_{out} – modulo external assets and the return payment from w to v. We define the *budget difference* by

$$d_w^{\text{sp}} = \left(a_w^x + \sum_{e' \in E^-(w_{in})} p_{e'}^{\text{sp}} \right) - a_w \qquad \text{and} \qquad d_v^{\text{sp}} = A - \left(a_v^x + \sum_{e' \in E^-(v_{in})} p_{e'}^{\text{sp}} \right) .$$

d_w^{sp} is the surplus money earned by w_{in} that shall be invested into the return, d_v^{sp} is the excess money spent by v_{out} that must be recovered through the return.

▶ **Lemma 12.** *For a given single claims trade and a given target value* $A > a_v$*, there is a creditor-positive trade with a value at least* A *for* v *if and only if* $d_v^{sp} = d_w^{sp} > 0$*.*

Proof. We first show that if $d_v^{\text{sp}} = d_w^{\text{sp}} > 0$, then there exists a creditor-positive trade with asset at least A for v. Suppose we consider $\mathbf{p}^{\text{sp}}$ in the network $\mathcal{F}'$ using return $\rho = d_v^{\text{sp}} = d_w^{\text{sp}}$. This exactly equilibrates the budgets of v and w – v receives d_v^{sp}, the money needed to obtain total assets of A. Also, w spends exactly d_w^{sp}, the money needed to obtain total assets of a_w. Hence, $\mathbf{p}^{\text{sp}}$ satisfies all fixed-point conditions in $\mathcal{F}'$. As such, $\mathbf{p}' \geq \mathbf{p}^{\text{sp}}$ coordinate-wise due to maximality of the clearing state. This implies that using return ρ, the clearing state $\mathbf{p}'$ yields $a'_v \geq A > a_v$ and $a'_w \geq a_w$. A creditor-positive trade with return ρ exists.

Now for the other direction, consider an optimal creditor-positive trade, which yields the highest asset level A^* and consider any $A \in (a_v, A^*]$. We show that in this case $d_v^{\text{sp}} = d_w^{\text{sp}} > 0$ holds in the clearing state $\mathbf{p}^{\text{sp}}$ of $\mathcal{F}^{\text{sp}}$ with external assets A for v_{out}.

Consider the optimal trade, its return $\rho^* > 0$ and the emerging payments $\mathbf{p}^*$ in $\mathcal{F}'$ after this trade. Now in the corresponding split network $\mathcal{F}^{*,\text{sp}}$ with external assets of A^* for v_{out}, the payments $\mathbf{p}^*$ yield $d_v^* = d_w^* = \rho^*$, by definition of $\mathbf{p}^*$. The previous paragraph and maximality of A^* then imply that $\mathbf{p}^*$ must also be the clearing state $\mathbf{p}^{*,\text{sp}} = \mathbf{p}^*$ of $\mathcal{F}^{*,\text{sp}}$.

Now suppose in $\mathcal{F}^{*,\text{sp}}$ we reduce the external assets of v_{out} by $\varepsilon = A^* - A > 0$. Then $\mathcal{F}^{\text{sp}}$ evolves. Since we reduce the assets of a single source v_{out} by ε, we obtain $\mathbf{p}^{*,\text{sp}} \geq \mathbf{p}^{\text{sp}}$. Moreover, by non-expansivity [7, Lemma 33], the total incoming assets of all sinks must reduce by *at most* ε. For the sinks v_{in} and w_{in} we set

$$\varepsilon_v = a_{v_{in}}^{*,\text{sp}} - a_{v_{in}}^{\text{sp}} = \sum_{e' \in E^-(v_{in})} p_{e'}^* - p_{e'}^{\text{sp}} \qquad\qquad \varepsilon_w = a_{w_{in}}^{*,\text{sp}} - a_{w_{in}}^{\text{sp}} = \sum_{e' \in E^-(w_{in})} p_{e'}^* - p_{e'}^{\text{sp}}$$

and, thus, $d_v^{\text{sp}} = d_v^* - (\varepsilon - \varepsilon_v)$ and $d_w^{\text{sp}} = d_w^* - \varepsilon_w$. Non-expansivity implies $\varepsilon_v + \varepsilon_w \le \varepsilon$.

First, we observe that $\varepsilon_v + \varepsilon_w < \varepsilon$ is impossible. Then $\varepsilon_w < \varepsilon - \varepsilon_v$, so $d_w^{\text{sp}} > d_v^{\text{sp}}$, i.e., w has more excess money in $\mathbf{p}^{\text{sp}}$ than required by v. Consider a return of $\rho = d_v^{\text{sp}}$ and $\mathbf{p}^{\text{sp}}$ as payment vector in the resulting network $\mathcal{F}'$. Then all banks are feasible w.r.t. the fixed-point conditions, except for w which has strictly more income than outgoing assets. Hence, the clearing state satisfies $\mathbf{p}' \ge \mathbf{p}^{\text{sp}}$, $a'_v \ge A > a_v$, and $a'_w > a_w$, a contradiction to Proposition 3.

Second, suppose that $\varepsilon_v + \varepsilon_w = \varepsilon$, then $d_v^{\text{sp}} = d_w^{\text{sp}}$. Then the clearing state $\mathbf{p}^{\text{sp}}$ exactly fulfills the fixed-point conditions for all banks in $\mathcal{F}'$ and yields assets $A > a_v$ for v and a_w for w with $\rho = d_v^{\text{sp}}$. Note that $\rho > 0$, since otherwise we contradict the maximality of the initial clearing state $\mathbf{p}$. Therefore, the existence of a creditor-positive trade with assets $A^* > A$ for v implies that $d_v^{\text{sp}} = d_w^{\text{sp}} > 0$ for $\mathbf{p}^{\text{sp}}$ in $\mathcal{F}^{\text{sp}}$ emerging from A. ◀

We are now ready to prove Theorem 11.

Proof of Theorem 11. Our algorithm works by testing different target values A for the total assets of v. For a given target value A, we then use Lemma 12 to verify existence of a return ρ achieving at least assets A for v. The maximum achievable assets for v are $M_v = \sum_{e' \in E^-(v) \setminus e} \ell_{e'} + a_v^x + \min\{a_w^x, \ell_e\}$. We determine the maximal achievable A using binary search on the interval $(a_v, M_v]$.

More formally, we choose $\delta > 0$ and apply binary search over the set $\{a_v + \delta, a_v + 2\delta, \ldots, M_v\}$. Verifying the condition in Lemma 12 can be done in polynomial time via a call to the clearing oracle in $\mathcal{F}^{\text{sp}}$. If the algorithm discovers that the condition does not hold for all tested values, then no creditor-positive trade with asset level at least $a_v + \delta$ for v exists. Otherwise, let $\hat{A}$ be largest discovered value for which the test is positive. Then, any value of at least $\hat{A} + \varepsilon$ cannot be achieved for any return ρ. Hence, the optimal achievable total assets of v in any creditor-positive trade are bounded by $A^* \in [\hat{A}, \hat{A} + \delta]$, and the additive approximation follows $\hat{A} - a_v \ge (A^* - a_v) - \delta$.

For the running time, we require at most $\lceil \log_2(1 + (M_v - a_v)/\delta) \rceil$ oracle calls, which is polynomial in the input size and $1/\delta$. ◀

Since the number of possible (single) claims trades in a network is limited by $|E| \cdot |V|$, we can use the algorithm above to compute every creditor-positive trade with an (approximately) optimized haircut rate for a given network in polynomial time.

4 Multi-Trades of Incoming Edges

4.1 Fixed Set of Claims

In this section, we are interested in multi-trades of incoming edges of a creditor bank v to a buyer bank w. This arises naturally, for example, when a high fraction of v's debtors are in default or v is "too big to fail". Then bankruptcy of v would cause significant damage to the entire network.

We are given a financial network $\mathcal{F}$ with two distinct banks v and w, and a set C of k incoming edges of v. Suppose the haircut rates α_i can be chosen individually for each $e_i \in C$ as part of the trade. We call a vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_k)$ of haircut rates *creditor-positive* if the resulting multi-trade is creditor-positive. Our goal is to select creditor-positive $\alpha_i \in [0,1]$, for every $i \in [k]$, in order to maximize the improvement of v, i.e., $a'_v - a_v = \sum_{i=1}^k \alpha_i \cdot \ell_{e_i} + \sum_{e' \in E'^-(v)} p'_{e'} - \sum_{e' \in E^-(v)} p_{e'}$. Observe that we can restrict our attention to vectors with uniform $\alpha_i = \alpha'$ for all $i \in [k]$ and some $\alpha' \in [0,1]$ – given any $\boldsymbol{\alpha}$, choose $\boldsymbol{\alpha}'$ with $\alpha'_i = \alpha'$ such that $\alpha' \cdot \sum_{i=1}^k \ell_{e_i} = \sum_{i=1}^k \alpha'_i \cdot \ell_{e_i}$. This results in $\alpha' \in [0,1]$, the same return, and the same assets of v as for $\boldsymbol{\alpha}$.

Our result is a reduction to single trades.

▶ **Proposition 13.** *Consider a financial network with monotone payment functions and efficient clearing oracle. For a given multi-trade of incoming edges, there is an additive FPTAS for approximating the optimal improvement of v from any creditor-positive $\boldsymbol{\alpha}$.*

Proof. Consider a financial network $\mathcal{F}$ with banks v and w and edges C, where $|C| = k$. By Corollary 4, the multi-trade in $\mathcal{F}$ can be modeled by a single claims trade with edge $\hat{e}$ in adjusted network $\hat{\mathcal{F}}$. Invoke the FPTAS to compute a haircut rate α for the single claim in $\hat{\mathcal{F}}$. This results in assets of $\alpha \cdot \sum_{i=1}^{k} \ell_{e_i} + \sum_{e \in E'^-(v)} p'_e$ for v in $\hat{\mathcal{F}}$. Clearly, the same value is obtained with the multi-trade when all haircut rates are set to α, i.e., $\alpha_i = \alpha \ \forall i \in [k]$. Clearly, this choice of haircut rates also yields an (approximately) optimal solution for the multi-trade. ◀

Combining the insight with Propositions 10 and 9, we obtain the following corollary.

▶ **Corollary 14.** *Consider a financial network with proportional or edge-ranking payments. For a given multi-trade of incoming edges, there are efficient algorithms to compute an optimal creditor-positive $\boldsymbol{\alpha}^*$ or decide that none exists.*

4.2 Choosing Subsets of Claims

For a fixed pair of creditor v and buyer w, the incoming edges of v yield an exponential number of different edge sets C that might be used for a multi-trade. Thus, a creditor-positive multi-trade *cannot* be derived trivially by checking feasibility for all sets C. For improving the assets of v by a multi-trade with buyer w, the selection of claims to be traded is critical. How can we compute a (near-)optimal set of incoming edges $C \subseteq (E^-(v) \setminus E^+(w))$ for a creditor-positive multi-trade with w such that we maximize the improvement of v?

The challenge is to find a set of claims C with creditor v and appropriate individual haircut rates α_i, for $e_i \in C$. The resulting multi-trade shall be creditor-positive and yield the maximal improvement for v (over all creditor-positive multi-trades of incoming edges of v).

We show that this problem is NP-hard, for every set of monotone payment functions. Formally, we show it is NP-hard to decide whether creditor v can be *saved* by a creditor-positive multi-trade of incoming edges, i.e., whether total assets of L_v can be achieved. We call this problem IncomingSave-VR (for variable haircut rates).

In the class of networks we construct for the reduction, every bank has at most one outgoing edge. Hence, all payments will be independent of the payment function that is used. Moreover, once a set of claims C is chosen, finding optimal haircut rates for the multi-trade of C to w is trivial in this class of networks. Hardness arises from the choice of C.

▶ **Theorem 15.** *IncomingSave-VR is weakly NP-hard.*

4.2.1 Approximate Claims Trades

Contrasting NP-hardness, we show that the problem to compute a multi-trade improving v by a given amount can be solved efficiently when slightly relaxing the liability condition.

▶ **Definition 16** (ε-Approximate Multi-Trade)**.** *A multi-trade C with creditor v, buyer w, return $\rho > 0$ and haircut rates $\alpha_i \in [0, 1+\varepsilon]$, for all $e_i \in C$, is called ε-approximate if $\rho \leq (1+\varepsilon) \cdot \sum_{e_i \in C} \ell_{e_i}$.*

Consider a *creditor-positive ε-approximate* trade. Such a trade (1) strictly increases the assets of v and exactly maintains the ones of w, (2) is affordable by w, i.e., $\rho = \sum_{e_i \in C} \alpha_i \ell_{e_i} \leq a^x_w$, and (3) satisfies exact fixed-point conditions in the emerging clearing state. It is approximate only in the liability condition of the trade.

We construct a bicriteria FPTAS to compute a creditor-positive multi-trade of incoming edges. Suppose $\varepsilon, \delta > 0$ such that $1/\varepsilon$ is polynomial and $1/\delta$ is exponential in the representation size of $\mathcal{F}$. Our FPTAS guarantees that the computed trade is ε-approximate and yields assets of at least $A^* - \delta$ for v, where A^* are the assets of v resulting from an *optimal exact* creditor-positive trade. The FPTAS uses a connection to the Knapsack problem.

We proceed in several steps: First, we consider computing an (exact) trade that achieves a target asset value A for the creditor. For this problem, we derive Knapsack-style constraints capturing a set of three necessary and sufficient conditions of a valid creditor-positive trade. We then adapt the dynamic program for Knapsack to construct an FPTAS to compute an ε-approximate multi-trade with assets value at least A for v in polynomial time. Finally, we show how to use binary search to find a trade with asset level at least $A^* - \delta$.

Necessary and Sufficient Conditions. As a first step, we consider exact trades that achieve assets of at least A for v. Suppose there is such a trade with a set C of traded edges, and let $k = |C|$. Consider $\mathcal{F}'$ after trade C has occurred with a suitably chosen return. Since C is fixed, by Corollary 4 we can express the outcome of the trade using a single claims trade. Now apply the split network $\mathcal{F}^{\text{sp}}$ and Lemma 12. Hence, using

$$d^{\text{sp}}_v = A - \left(a^x_v + \sum_{e' \in (E^-(v) \setminus C)} p'_{e'} \right) \quad \text{and} \quad d^{\text{sp}}_w = \left(a^x_w + \sum_{e' \in (E^-(w) \cup C)} p'_{e'} \right) - a_w \ ,$$

a creditor-positive trade with set C and asset value at least A exists if and only if $d^{\text{sp}}_v = d^{\text{sp}}_w > 0$. This implies that

$$(A - a^x_v) + (a_w - a^x_w) = \sum_{e \in E^-(v) \cup E^-(w)} p'_e \tag{1}$$

must hold. This condition is *independent* of the set C of traded edges. As such (1) is a necessary condition that any creditor-positive trade with asset level at least A for v can exist.

With return $\rho = d^{\text{sp}}_v$ for the given set C, we satisfy the fixed-point conditions in $\mathcal{F}'$. By Lemma 12 the optimal trade using the given set C only yields a larger return, i.e., $\rho \geq d^{\text{sp}}_v > 0$. Moreover, the liabilities of the traded edges must be high enough to allow the return ρ, i.e., $\sum_{e_i \in C} \ell_{e_i} \geq \rho$. Using $P' = \sum_{e \in E^-(v)} p'_e$ this necessary condition is expressed by

$$\sum_{e \in C} \ell_e \geq d^{\text{sp}}_v = A - a^x_v - P' + \sum_{e \in C} p'_e \quad \iff \quad \sum_{e \in C} (\ell_e - p'_e) \geq A - a^x_v - P' \ . \tag{2}$$

w must be able to pay the required return. Since the return is solely funded by external assets, we obtain the necessary condition

$$a^x_w \geq \rho = d^{\text{sp}}_v \quad \iff \quad \sum_{e \in C} p'_e \leq a^x_w - A + a^x_v + P' \ . \tag{3}$$

While each condition (1)-(3) is necessary, it is easy to see that in combination they are sufficient. Indeed, if they hold, then there is a creditor-positive trade of set C with return $\rho \geq d^{\text{sp}}_v = d^{\text{sp}}_w > 0$ that respects the exact liabilities of traded edges, is affordable by w, and achieves asset level at least A for v. We summarize the argument in the following lemma:

▶ **Lemma 17.** *For a given set C of incoming edges of v, the following are equivalent:*
1. *There is a multi-trade of C to w that achieves an asset level at least A for v.*
2. *Equation* (1) *holds, and the set C satisfies* (2) *and* (3).

Knapsack-Style FPTAS. While condition (1) can be checked directly after computing the clearing state $\mathbf{p}'$, determining the existence of a set C that satisfies conditions (2) and (3) can be cast as a Knapsack decision problem: For each edge $e \in E^-(v)$ the payments p'_e are the non-negative *weight* of e, and the residual $\ell_e - p'_e$ is the non-negative *value* of e. Decide the existence of a subset of edges with total value lower bounded by (2) and total weight upper bounded by (3).

We next adapt the standard FPTAS for Knapsack to compute an *approximate* multi-trade. We *round up* the residual $\ell_e - p'_e$ of every edge to the next multiple of a parameter s. This can be interpreted as increasing the liabilities ℓ_e by a small amount. We then determine if (2) and (3) allow a feasible solution by using the standard dynamic program for Knapsack in polynomial time. We term this procedure the Level-FPTAS.

▶ **Lemma 18** (Level-FPTAS). *For a given number A, suppose there exists a creditor-positive multi-trade of incoming edges that yields assets at least A for v. Then, for every $\varepsilon > 0$, there is an algorithm to compute an ε-approximate creditor-positive trade with assets at least A for v in time polynomial in the size of $\mathcal{F}$ and $1/\varepsilon$.*

Proof. First, check feasibility of condition (1) since otherwise the desired trade does not exist. Then, consider all incoming edges $e \in E^-(v)$ of v and define $m = |E^-(v)|$. We denote by $r_e = \ell_e - p'_e$ the residual of e. Let r_{max} be the maximal residual capacity with respect to $\mathbf{p}'$ of any edge that satisfies the weight constraint, i.e., $r_{max} = \max\{r_e \mid e \in E^-(v), p'_e \leq a^x_w - A + a^x_v + P'\}$. Round the residual capacities *up* using a scaling factor $s = \frac{\varepsilon \cdot r_{max}}{m}$. Then determine the optimal solution C^* for the rounded values $\tilde{r}_e = s \cdot \lceil (r_e)/s \rceil$ via the standard dynamic program for Knapsack. The running time is bounded by $O(m^3/\varepsilon)$.

Using a return of $\rho = A - \left(a^x_v + \sum_{e\in E^-(v)\setminus C^*} p'_e\right)$ the trade of C^* yields a clearing state in $\mathcal{F}'$ with assets at least A for v. By definition, C^* satisfies (3), and since the instance satisfies (1), $\rho \leq a^x_w$ is also guaranteed.

Regarding the liabilities, note that

$$\sum_{e\in C^*} \tilde{r}_e \;\leq\; \sum_{e\in C^*} r_e + s \;\leq\; \varepsilon r_{max} + \sum_{e\in C^*} r_e \;\leq\; (1+\varepsilon) \sum_{e\in C^*} r_e \;\leq\; \sum_{e\in C^*} \ell_e(1+\varepsilon) - p'_e \;.$$

There exists an exact trade C with assets at least A for v, so the trade with C satisfies (2) and (3). As such,

$$\sum_{e\in C^*} \tilde{r}_e \quad\geq\quad \sum_{e\in C} \tilde{r}_e \quad\geq\quad \sum_{e\in C} r_e \quad\geq\quad A - a^x_v - P' \;,$$

so the optimal solution C^* satisfies (2) using the rounded residuals. Hence,

$$\begin{aligned}\rho \quad &= \quad A - \left(a^x_v + \sum_{e\in E^-(v)\setminus C^*} p'_e\right) \quad = \quad A - a^x_v - P' + \sum_{e\in C^*} p'_e \quad \leq \quad \sum_{e\in C^*} \tilde{r}_e + p'_e \\ &\leq \quad (1+\varepsilon) \sum_{e\in C^*} \ell_e \;,\end{aligned}$$

so the return generated by C^* violates the liability condition by at most a factor of $1+\varepsilon$. ◀

The lemma gives rise to an efficient algorithm computing an ε-approximate multi-trade with assets at least A, whenever an exact trade with assets at least A exists. Similar to Theorem 11, we use this test to construct a bicriteria FPTAS.

Bicriteria-FPTAS. As our main result for multi-trades of incoming edges, we obtain a bicriteria FPTAS. Suppose assets A^* for v are achievable by an exact creditor-positive multi-trade of incoming edges. We will compute an ε-approximate one resulting in assets at least $A^* - \delta$ for v, for any $\varepsilon, \delta > 0$. We say such a trade is δ-*optimal.*

▶ **Theorem 19.** *Consider a financial network with monotone payment functions and efficient clearing oracle, creditor v and buyer w. If there exists a creditor-positive multi-trade of incoming edges, then an ε-approximate δ-optimal trade can be computed in time polynomial in the size of $\mathcal{F}$, $1/\varepsilon$ and $\log 1/\delta$, for every $\varepsilon, \delta > 0$,*

Proof. We use the binary search idea put forward in Theorem 11. We choose $\delta > 0$ and apply binary search over the set $\{a_v + \delta, a_v + 2\delta, \ldots, M_v\}$ of potential asset values for v. Recall that M_v is an upper bound for maximal achievable assets of v. For multi-trades, M_v is upper bounded by $a_v^x + a_w^x + \sum_{e \in E^-(v)} \ell_e$. The goal is to find an asset value that is as large as possible.

Running the Level-FPTAS with any value from the interval $A' \in (a_v, A^*]$, we are guaranteed to receive an approximate multi-trade with asset level at least A' for v in polynomial time. As such, the binary search will never terminate with a value of $A' \leq A^* - \delta$. The search terminates in at most $\lceil \log_2(1 + (M_v - a_v)/\delta) \rceil$ steps. ◀

When called with an asset level $A' > A^*$, the Level-FPTAS might or might not return a corresponding multi-trade – rounding up the residuals can introduce non-monotone behavior. As such, using binary search our algorithm does not necessarily return an optimal asset value of v for any creditor-positive ε-approximate multi-trade. However, since the Level-FPTAS never fails to return a multi-trade for any asset level $A' \leq A^*$, we are guaranteed that assets of more than $A^* - \delta$ for v are achieved.

Ranking Payments. For edge-ranking payments, the set of meaningful values to be tested for A^* in the binary search can be restricted to a grid of at most exponential precision in the input size. This allows to compute an ε-approximate multi-trade with assets at least A^*, i.e., such a trade is δ-optimal with $\delta = 0$.

▶ **Corollary 20.** *Consider a financial network with edge-ranking functions, creditor v and buyer w. If there exists a creditor-positive multi-trade of incoming edges, then an ε-approximate 0-optimal trade can be computed in time polynomial in the size of $\mathcal{F}$ and $1/\varepsilon$.*

Proof. Consider an optimal exact multi-trade C with return ρ that achieves optimal assets of A^* for v. Recall that all liabilities and external assets are integers, and so is M_v. If A^* is integral, then we can run the binary search with $\delta = 1$ and obtain an approximate trade with assets of (more than $A^* - 1$ and, thus) at least A^* for v.

To show that A^* is integral, consider an optimal creditor-positive trade C achieving assets A^*. We resort to the equivalent representation as a single claims trade (Corollary 4). For this single trade, consider the return network $\mathcal{F}^{\text{ret}}$. An optimal return ρ for a trade of edge set C evolves as the payment $p_{e_r}^{\text{ret}}$ on e_r in the clearing state $\mathbf{p}^{\text{ret}}$. Recall that the payment function of w in the return network is also an edge-ranking function (c.f. Proposition 9). For edge-ranking payments, if all liabilities and external assets are integers, then the clearing state has integral payments [1]. The assets of every bank in $\mathbf{p}^{\text{ret}}$ (and A^*) are integral. ◀

5 Multi-Trades of Outgoing Edges

In this section, we study multi-trades of *outgoing* edges of a bankrupt bank u. We strive to improve the assets of u's creditors directly (and not indirectly via u through trades of incoming edges). It might not be feasible to save a particular bank u, e.g., when its debt is too high in relation to the claims. In such cases, we attempt to minimize the *contagion* of bankruptcy from u to her creditors by conducting multi-trades of outgoing edges of u. We execute multi-trades that *maximize* the total profit of all creditors, not just those involved in the trade. No creditor nor buyer w should be harmed by the trade.

▶ **Definition 21** (Pareto-positive trade). *Let $v_1, v_2, \dots, v_l$ be u's creditors. A multi-trade of outgoing edges of u to w is called Pareto-positive, if $a'_{v_i} > a_{v_i}$ for at least one creditor v_i, $a'_{v_i} \geq a_{v_i}$ for all creditors and $a'_w \geq a_w$.*

Suppose we are given a financial network with banks u, w and a set C of k outgoing edges of u. Denote the creditors of edges C by $V_C = \{v_i \mid e_i \in C, \mathrm{cr}(e_i) = v_i\}$. A collection of haircut rates $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_k)$ is called *Pareto-positive* if C together with $\boldsymbol{\alpha}$ forms a Pareto-positive multi-trade. The objective is to derive the optimal values of $\boldsymbol{\alpha}$ which maximize the sum of profit of creditors $v_1, v_2, \dots, v_l$, i.e., $\max \sum_{i=1}^{l} a'_{v_i} - a_{v_i}$.

Consider the problem where set C is not given as part of the input but is chosen as part of the solution. The goal is to select a subset of u's outgoing edges $C \subseteq E^+(v)$ together with a vector of haircut rates $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_{|C|})$ such that the multi-trade is Pareto-positive and maximizes the improvement of u's creditors, i.e., $\sum_{i=1}^{l} a'_{v_i} - a_{v_i}$.

In the previous section, we obtained a bicriteria FPTAS for this problem when we trade incoming edges of an insolvent bank. Interestingly, the results hold for all monotone payment functions for which there is an efficient clearing oracle (e.g., edge-ranking or proportional payments). Our results here show a strong contrast – depending on the payment functions trading outgoing edges can be much harder. For edge-ranking functions (and variable haircut rates), we denote the problem by OutgoingER-VR.

▶ **Theorem 22.** *OutgoingER-VR is strongly NP-hard. For any constant $\varepsilon > 0$ there exists no efficient $n^{1/2-\varepsilon}$-approximation algorithm for OutgoingER-VR unless P = NP.*

Now suppose the set of traded edges C is given as part of the input. Interestingly, the hardness for edge-ranking payments continues to apply when C is fixed a priori.

▶ **Corollary 23.** *Consider a financial network with banks u, w, a set of outgoing edges C of u and edge-ranking payment rules. It is strongly NP-hard to determine Pareto-positive haircut rates that maximize the sum of profits of u's creditors. For any constant $\varepsilon > 0$, there exists no efficient $n^{1/2-\varepsilon}$-approximation algorithm unless NP = P.*

Finally, we briefly observe that these problems for outgoing edges depend crucially on the set of payment functions. For proportional payments (and variable $\boldsymbol{\alpha}$), the problem for a given set C can be solved efficiently (even if the set C of edges involves different debtors). When the set C of outgoing edges is chosen as part of the solution, we refer to the problem as OutgoingPROP-VR and obtain NP-hardness. The approximability status of these problems for different payment functions is an interesting direction for future work.

▶ **Proposition 24.** *For a given financial network with proportional payments and a set C of k edges, there exists an efficient algorithm that computes an optimal Pareto-positive $\boldsymbol{\alpha}^* \in [0, 1]^k$ or decides that none exists.*

▶ **Theorem 25.** *OutgoingPROP-VR is strongly NP-hard.*

References

1 Nils Bertschinger, Martin Hoefer, and Daniel Schmand. Strategic payments in financial networks. In *Proc. 11th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, pages 46:1–46:16, 2020.
2 Bloomberg. URL: `https://www.bloomberg.com/news/features/2023-03-20/credit-suisse-ubs-takeover-how-a-166-year-old-bank-collapsed`.
3 Swiss Federal Council. URL: `https://www.admin.ch/gov/en/start/documentation/media-releases.msg-id-93793.html`.
4 Péter Csóka and Jean-Jacques Herings. Decentralized clearing in financial networks. *Manag. Sci.*, 64(10):4681–4699, 2018.
5 Larry Eisenberg and Thomas Noe. Systemic risk in financial systems. *Manag. Sci.*, 47(2):236–249, 2001.
6 Jared Ellias. Bankruptcy claims trading. *J. Empirical Legal Stud.*, 15(4):772–799, 2018.
7 Henri Froese, Martin Hoefer, and Lisa Wilhelmi. The complexity of debt swapping. *arXiv preprint arXiv:2302.11250*, 2023.
8 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.
9 Stavros Ioannidis, Bart de Keijzer, and Carmine Ventre. Financial networks with singleton liability priorities. In *Proc. 15th Symp. Algorithmic Game Theory (SAGT)*, pages 205–222, 2022.
10 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.
11 Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. Financial network games. In *Proc. 2nd Int. Conf. AI in Finance (ICAIF)*, pages 26:1–26:9, 2021.
12 Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. Forgiving debt in financial network games. In *Proc. 31st Int. Joint Conf. Artif. Intell. (IJCAI)*, pages 335–341, 2022.
13 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.
14 Adam Levitin. Bankruptcy markets: Making sense of claims trading. *Brooklyn J. Corporate, Financial & Commercial Law*, 4(1), 2009.
15 Katherine Mayo and Michael P. Wellman. A strategic analysis of portfolio compression. In *Proc. 2nd Int. Conf. AI in Finance (ICAIF)*, pages 20:1–20:8, 2021.
16 Marios Papachristou and Jon Kleinberg. Allocating stimulus checks in times of crisis. In *Proc. 31th World Wide Web Conf. (WWW)*, pages 16–26, 2022.
17 Pál András Papp and Roger Wattenhofer. Network-aware strategies in financial systems. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *Proc. 47th Int. Colloq. Autom. Lang. Programming (ICALP)*, pages 91:1–91:17, 2020.
18 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.
19 Pál András Papp and Roger Wattenhofer. Default ambiguity: Finding the best solution to the clearing problem. In *Proc. 17th Conf. Web and Internet Econ. (WINE)*, pages 391–409, 2021.
20 Leonard Rogers and Luitgard Veraart. Failure and rescue in an interbank network. *Manag. Sci.*, 59(4):882–898, 2013.
21 Steffen Schuldenzucker and Sven Seuken. Portfolio compression in financial networks: Incentives and systemic risk. In *Proc. 21st Conf. Econ. Comput. (EC)*, page 79, 2020.
22 Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Clearing payments in financial networks with credit default swaps. In *Proc. 17th Conf. Econ. Comput. (EC)*, page 759, 2016.
23 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.
24 Luitgard Veraart. When does portfolio compression reduce systemic risk? *Math. Finance*, 32(3):727–778, 2022.