

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Flow Allocation Games*

Nils Bertschinger

Frankfurt Institute of Advanced Studies, Goethe University Frankfurt, bertschinger@fias.uni-frankfurt.de

Martin Hoefer

Institute for Computer Science, Goethe University Frankfurt, mhoefer@em.uni-frankfurt.de

Daniel Schmand

Center for Industrial Mathematics, University of Bremen, schmand@uni-bremen.de

We study a game-theoretic variant of the maximum circulation problem. In a *flow allocation game*, we are given a directed flow network. Each node is a rational agent and can strategically allocate any incoming flow to the outgoing edges. Given the strategy choices of all agents, a maximal circulation that adheres to the chosen allocation strategies evolves in the network. Each agent wants to maximize the amount of flow through her node. Flow allocation games can be used to express strategic incentives of clearing in financial networks.

We provide a cumulative set of results on the existence and computational complexity of pure Nash and strong equilibria, as well as tight bounds on the (strong) prices of anarchy and stability. Our results show an interesting dichotomy: Ranking strategies over individual flow units allow to obtain optimal strong equilibria for many objective functions. In contrast, more intuitive ranking strategies over edges can give rise to unfavorable incentive properties.

Key words: Equilibrium Computation, Flow Game, Maximum Circulation Problem, Financial Networks

MSC2000 subject classification: 68Q25, 68W40, 91A43, 91G45

OR/MS subject classification: Primary: Networks/graphs: Flow Algorithms; secondary: Games/group decisions: Noncooperative

History: Received: December 21, 2022; revised: November 4, 2023.

1. Introduction. Flows and circulations in networks are a classic problem domain in combinatorial optimization. A network flow is called a feasible circulation in a graph if it maintains flow conservation at all nodes, and some given lower and maybe upper bounds on the flow on edges are fulfilled. Dinitz [11] and Edmonds and Karp [13] provided strongly polynomial-time algorithms for solving the existence problem of a circulation in a network. Tardos [45] even showed that the minimum-cost circulation problem can be solved in strongly polynomial time. There is a vast number of applications of flow problems, and flows give rise to beautiful and favorable mathematical and algorithmic properties. As a consequence, variants of flow and circulation problems have been investigated for decades.

In this paper, we explore a novel game-theoretic model for circulations in networks. In our model, each node belongs to a player who aims to maximize the flow through the node. We assume that

* An extended abstract of this paper has been published in the proceedings of the ITCS 2020 conference [4] with the title “Strategic Payments in Financial Networks”. While we receive motivation from financial networks, we study a fundamental game-theoretic approach in the context of classic flow and circulation problems that is not necessarily restricted to financial networks. To reflect this property, we changed the title to “Flow Allocation Games”.

players can strategically allocate the flow entering their node to outgoing edges. Given the strategy choices of all players, a maximal circulation that adheres to the chosen allocation strategies evolves in the network.

Flows and circulations have a broad range of applications to different areas. A fascinating area, where circulations play a key role, is in the analysis of financial networks and systemic risks. A very popular approach to model financial networks has emerged from the seminal work by Eisenberg and Noe [14]. In their model, the financial market can be seen as a directed graph $G = (V, E)$. The node set V corresponds to the set of institutions (or *firms*). The set of directed edges E expresses the debt relations among firms. Each edge $e = (u, v) \in E$ has a capacity c_e that corresponds to the nominal liability of firm u to firm v . In addition, each firm v has a non-negative supply b_v that corresponds to external assets. Given this networked scenario, the goal is to understand the properties of *clearing*, i.e., the resulting payments when firms have to clear their debt and “pay their bills” to their creditors. For this clearing task, we view the graph as a flow network, where payments constitute a flow of funds or assets (which we will call *money*). Each node has a non-negative supply and potentially receives additional money over its’ incoming edges. It then can use all this money to allocate flow towards its’ outgoing edges. As a consequence, money starts to circulate in the network. Eventually, all edges of a node become tight (and all debt is cleared), or the node runs out of funds.

In the majority of the literature, the circulation flow is governed by a static proportional (also called pro-rata) strategy for each node. The node must allocate the entire outgoing flow in proportion to the capacities of its outgoing edges. More recently, interest has emerged in more general, decentralized, and monotone allocation strategies. Notably, in an influential work, Csóka and Herings [9] consider flows based on an arbitrary integral, monotone allocation strategy for each node (which we term *unit ranking* below). This extension leads to a variety of interesting questions for the resulting money circulations.

In this work, our interest lies in the decentralized and, more concretely, incentive and stability properties of the circulation problem. In our model, nodes can individually choose the flow allocation strategy for the outgoing edges. Flow allocation games of this kind have been of interest recently – in particular, Guha et al. [21] explore games, in which each agent controls the strategies of one or more nodes and wants to maximize the flow routed to an agent-specific sink in the network. Our game is closely related, with the difference that each agent corresponds to a single node, but it might neither be a source nor a sink. Instead, the goal of the node in our game is to maximize the flow circulating through the node. In a financial context, this is equivalent to the natural goal of maximizing the *equity* (total assets minus total liabilities).

Each node can choose as strategy an allocation function that yields for each amount of available flow an assignment of this flow to the outgoing edges. Similar to Csóka and Herings [9] our interest lies in monotone, ranking-based strategies. Game-theoretic variants of maximum flows, even based on rankings, have been of recent interest in economics. Fleiner [16] introduced *stable flows*, which have been further developed by Király and Pap [32], Cseh et al. [8], and Cseh and Matuschke [7]. In stable flows, each node has an inherent preference ranking over the edges. A node v strives to maximize the amount of flow on its preferred edges. Thus, a stable flow can be seen as a flow that is immune to coalitional deviations of players, i.e., nodes of a non-saturated walk can jointly decide to add flow on this walk. In our work, the ranking is not externally given, but represents a strategic decision of the players.

To the best of our knowledge, flow allocation games studied in this paper have not been addressed before. We provide a cumulative analysis of the properties of equilibria in these games. We focus on pure Nash and strong equilibria. In these equilibria, nodes have no unilateral (pure Nash) or coalitional incentives (strong) to deviate from their chosen strategies. Depending on the set of strategies, the resulting games have different properties. If we assume that strategies are restricted

to priority orderings over edges, existence of a pure Nash or a strong equilibrium is not guaranteed and becomes strongly NP-hard to decide. Instead, if nodes can assign each unit of available flow in an arbitrary monotone fashion, a strong equilibrium always exists and can be computed in strongly polynomial time. Moreover, this strong equilibrium maximizes the total amount of flow circulating in the network (and, as such, Pareto optimizes the utility of all nodes). In addition, we show that for a diverse set of objective functions, there is a strong equilibrium that is optimal w.r.t. this objective (such as, e.g., minimizing the number of nodes or firms that are in default and unable to pay their liabilities).

This interesting technical dichotomy between games with different payment strategies (restricted edge-based vs. arbitrary integral and monotone) offers insights into the properties of financial networks. Our results show that a benevolent designer could realize a clearing mechanism with monotone unit-based payment strategies that leads to a socially optimal state, for many different notions of “social optimum”. It comes with the additional guarantee of giving no coalition of firms an incentive to pay their debts differently. In contrast, if clearing payments are determined in a decentralized fashion resulting in some arbitrary Nash or strong equilibrium, the total amount of flow in the system can deteriorate drastically (and similarly the social quality for many objectives).

Similar problems arise if a centralized mechanism is restricted to edge-based priorities. This can lead to non-existence of pure equilibria in the resulting games. Even if pure equilibria exist, they can be undesirable since, e.g., the total amount of circulating flow can be very small. This shows a marked contrast between centralized and decentralized approaches and highlights how the structure of permissible strategies impacts the structural properties of the resulting flows.

1.1. Contribution and outline. In Section 2 we introduce our formal model of a flow allocation game. We focus on natural classes of ranking-based payment strategies for the nodes as introduced by Csóka and Herings [9]. For an *edge-ranking strategy*, a node ranks its outgoing edges and assigns its incoming flow in the order of the ranking¹. As a superset of strategies, we consider *unit-ranking strategies*, where flow is considered in units. Instead of edges, each node ranks single units of each edge capacity. Edge- and unit-ranking strategies both are classes of *monotone strategies*, where the mapping of the available flow of a node to every outgoing edge is an arbitrary monotone function.

There can be several feasible flows for a given strategy profile. In fact, the feasible flows form a lattice with a partial order based on the total outgoing flow of each node [9]. There is a unique feasible flow that forms the supremum of the lattice – it pointwise maximizes the outgoing flow to each node (for the given profile of monotone strategies). We assume that this supremum flow is the clearing flow and defines the utility of each node in the game. After discussing some structural insights on the feasible flows in Section 3.1, we show in Section 3.2 that the supremum flow can be computed in strongly polynomial time (Proposition 5) for profiles of unit-ranking strategies.

In Section 4 we study unit-ranking games, in which all nodes use unit-ranking strategies. Our interest lies in the existence, computational complexity, and social quality of equilibria. We show that in every such game there exists a strategy profile that represents a strong equilibrium, in which no coalition of nodes has an incentive to deviate (Theorem 3). Furthermore, there are even strong equilibria that represent a state that maximizes the circulation flow. Moreover, for a large variety of natural notions of social objective functions (such as the sum, the minimum, or the geometric mean of all utilities, the number of fully saturated nodes, etc.) there is a strong equilibrium that maximizes the objective over all possible circulations (Corollary 1).

For simplicity, in the remainder of the paper we then focus on one standard objective in circulation problems, the total amount of flow circulating in the network. A strong equilibrium that

¹ Csóka and Herings called them priority rules.

maximizes the circulation can be computed in strongly polynomial time (Theorem 3). It can even be represented using a number of bits polynomial in the size of the input (i.e., the size of the graph and the size of all numbers in logarithmic encoding), even though every strategy shall rank all single flow units that a node might have available, and their number could be pseudo-polynomial (i.e., linear in the unary encoding size of the input). In contrast, it is strongly NP-hard to find a best-response strategy for a single node in a given arbitrary strategy profile of a unit-ranking game (Theorem 5).

For worst-case equilibria and the strong price of anarchy, we show that the deterioration of the circulation in a strong equilibrium compared to a max-circulation is tightly characterized by the min-max length of cycles in any max-circulation (Theorem 6). This implies that in networks, in which an optimal circulation is composed of small cycles, we see a small inefficiency in strong equilibria. In contrast, the circulation of a worst-case Nash equilibrium, which is stable only against unilateral deviations, can be arbitrarily worse than in an optimum, even in simple games with a constant number of nodes (Proposition 6).

In Section 5 we study a natural and interesting restriction on the strategies and analyze edge-ranking games, in which all nodes are restricted to edge-ranking strategies. Restricting the strategy space to rankings over edges can have devastating consequences for the existence of equilibria and the amount of circulating flow in an equilibrium. In edge-ranking games, pure Nash and strong equilibria can be absent, and deciding their existence is strongly NP-hard (Theorem 7). The same hardness applies for computing an optimal strategy profile, and for computing a pure Nash or strong equilibrium when it is guaranteed to exist. Even the best strong equilibrium can be a factor of $\Omega(n)$ worse than an optimum in terms of the circulating flow (Proposition 10). For pure Nash equilibria, even the best one can be arbitrarily worse than an optimum (Proposition 11).

We conclude in Section 6 with a summary of the main findings, a discussion of our results, and directions for future work.

1.2. Related work. Flow allocation games are based on circulations in financial network models that emerged from [14]. Rather than proportional payments, Csóka and Herings [9] analyze edge- and unit-ranking strategies as well as arbitrary monotone strategies. They analyze the structure of clearing flows and show that they constitute a complete lattice. For completeness, in Section 2 we recapitulate these findings using our notation.

Flow games. Our game-theoretic approach is related to a number of existing game-theoretic models based on flows in networks. In cooperative game theory, there are several notions of flow games based on a directed flow network. Existing variants include games, where edges are players [12, 27, 28, 20, 10, 2], or each player owns a source-sink pair [36, 34]. The total value of a coalition C is the profit from a maximum (multi-commodity) flow that can be routed through the network if only the players in C are present. There is a rich set of results on structural characterizations and computability of solutions in the core, as well as other solution concepts for cooperative games. In contrast to our work, these games are non-strategic. We consider each player as a single node with a strategic decision about flow allocation.

More recently, a class of strategic flow games has been proposed in [21, 33]. There is a capacitated flow network with a set of source nodes. At each source node, a given amount of flow enters the network. Each node of the network is owned by a single player. Each player always owns a designated sink node, as well as one or more additional nodes from the network. A player can choose a flow strategy for each of her nodes. The flow strategy specifies, for every node v and every $x \geq 0$, how an incoming flow of x at v is distributed onto the outgoing edges (if any). Each flow strategy needs to fulfill flow conservation constraints at every node, subject to capacity on the outgoing edges. Each player aims to maximize the incoming flow at its sink node.

For these games there exist a number of Σ_2^P -completeness results for, e.g., determining the value of a game in a two-player Stackelberg variant, or determining the existence of a pure Nash equilibrium in a multi-player variant. In the latter game, computing a best response can also be NP-hard. Our approach is related to these games. However, motivated by financial networks we assume each firm is a single node. The firm optimizes the incoming flow at its node (without it being a designated sink node). We study the computational complexity and social quality of equilibria. Moreover, strategic incentives arise mainly from cycles in the network (see Section 2.4 below) – a condition absent in the existing work on max-flow games [21, 33] where the network is assumed to be acyclic.

The problem of computing a clearing state for a given strategy profile in our games is closely related to the notion of a stable flow studied in [16, 7]. In the stable flow problem, each node is equipped with an intrinsic preference order over both incoming and outgoing arcs. The goal is to route as much flow as possible over most preferred arcs. There always exists a stable flow, where no group of agents all can benefit from rerouting the flow along a walk, and such a flow can be computed in polynomial time. The set of stable flows forms a lattice. The model has been extended to an over time setting [8] and to a multi-commodity variant [32].

Financial networks. We consider issues of strategic choice and computational complexity in flow allocation games. Flow allocation games have a strong connection to clearings in financial networks. There have been works addressing computational complexity of diverse issues in financial networks, such as pricing options with [1] and without information asymmetry [5], finding clearing payments with credit default swaps [42], or estimating the number of defaults when providing a shock in the financial system [23].

In addition, many extensions to the model by Eisenberg and Noe have been proposed in the literature on financial networks. However, even models including cross-holdings of equity [44], default costs [41], or debt contracts of different seniorities [15] follow the idea of the basic approach that all contracts have to be cleared consistently, i.e., clearing payments locally adhere to the rather mechanical clearing rule and constitute a fixed point solution globally. Indeed, Barucca et al. [3] have shown that many of the above models can be unified in terms of self-consistent network valuations. A well-known result of such models is the “robust-yet-fragile” property exhibited by financial networks, i.e., contagion arises in an all-or-nothing fashion akin to the formation of a giant connected component in random graph models [19]. This provides important insights into systemic risk and advises the need for macro-prudential regulation.

An extended abstract of this paper has appeared in the proceedings of the ITCS 2020 conference [4]. Subsequent to publication of the extended abstract, there has been a significant interest in extending our results and analyzing closely related computational problems in the context of financial networks. Kanellopoulos et al. [29] study flow allocation games with credit default swaps (CDS). They obtain a number of NP-hardness results for equilibrium existence. Instead of CDSes, Hoefler and Wilhelmi [24] consider extensions to seniorities (expressed by a class of threshold strategies) along with minimal clearing based on the infimum (rather than the supremum) of feasible flows. For endogenous seniorities, they extend our main existence result for strong equilibria; for exogenous seniorities, they observe NP-hardness results for equilibrium existence.

More fundamentally, Ioannidis et al. [25, 26] prove FIXP-completeness results for computing a strong approximation of clearing states with debts and CDSes, contrasting PPAD-completeness for weaker notions of approximation [42]. Further hardness results for computing optimal clearing states are provided by Papp and Wattenhofer [39]. They also study the question of computing the set of banks that default in such networks [40]. Interestingly, for networks with CDSes, the point in time at which a bank announces to default plays an important role. The same authors analyze incentives for adjusting the network structure to optimize individual funds in the clearing state by deleting an incoming edge or gifting some money to other banks [37]. Similar questions of optimally

removing debt contracts or bail-outs are analyzed by Kanellopoulos et al. [30]. In a related spirit, the effects of *debt swaps* of incoming payment obligations by two banks are analyzed in [38, 18].

In a broader context, strategic and financial aspects of networks have received substantial attention over the last decade. A related, yet orthogonal, body of work considers trading networks [22, 35]. These models are closely related to two-sided matching under preferences and the study of competitive equilibrium. Rather than circulation effects, agents strive to establish profitable trades with their neighbors. Depending on the model variant, this results in agents matching into pairs, exchanging goods, or establishing upstream/downstream relations with suppliers and customers. The analysis of these networks usually addresses similar issues as the ones we consider here, such as existence, structure, and computational complexity of equilibria. Equilibrium computation in trading networks has recently been studied in, e.g., [6, 17]. These works also contain a good overview of related work and pointers into the existing literature.

2. Flow allocation games.

2.1. Network model, monotone strategies, and utilities.

Network model. In a *flow allocation game* $\Gamma = (G, (b_v)_{v \in V}, (c_e)_{e \in E})$ we are given a graph $G = (V, E)$ with a node set V and a set of directed edges E . Each node $v \in V$ corresponds to a *player* and has a *fixed supply* $b_v \geq 0$. The *capacity* $c_e \geq 0$ of an edge $e = (u, v)$ is the maximum amount of flow that u can forward to v . In terms of financial networks, b_v is the amount of external assets of v , and e represents a liability of value c_e of firm u to firm v . We follow standard notation in graph theory and denote by $E^+(v) = \{(v, u) \in E\}$ and by $E^-(v) = \{(u, v) \in E\}$ the set of outgoing and incoming edges of $v \in V$, respectively. The *saturating output* c_v^+ of a node v is the maximum amount of flow v can send to other players, specified by the weighted outdegree

$$c_v^+ = \sum_{e \in E^+(v)} c_e.$$

We strive to analyze issues of computational complexity. As such, we will assume that all numbers in the input, i.e., all b_v and c_e , are integer numbers in binary encoding.

Allocation strategies. We analyze allocation strategies in flow allocation games. Each player $v \in V$ strategically allocates its *total supply*, i.e., the fixed supply b_v plus the total amount of incoming flow, to the outgoing edges $E^+(v)$. More formally, each player $v \in V$ chooses as a *strategy* a parametrized flow allocation function $a_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ for every outgoing edge $e \in E^+(v)$. This function specifies an amount of flow $a_e(y)$ that is forwarded to the edge e , for every $y \in \mathbb{R}_{\geq 0}$. Here, y represents a possible value of total supply. Intuitively, the strategy specifies for every possible value $y \geq 0$ of total flow that v might have available, how v will allocate this flow to the outgoing edges. The strategy $\mathbf{a}_v = (a_e)_{e \in E^+(v)}$ of player $v \in V$ must satisfy for every $y \geq 0$ and $e \in E^+(v)$

$$0 \leq a_e(y) \leq c_e, \quad (\text{capacity constraint}) \quad (1)$$

$$\sum_{e \in E^+(v)} a_e(y) \leq y, \quad (\text{weak flow conservation constraint}) \quad (2)$$

$$\sum_{e \in E^+(v)} a_e(y) = \min\{c_v^+, y\}. \quad (\text{no-fraud constraint}) \quad (3)$$

The capacity constraint ensures that no edge is overused, and the weak-flow conservation constraint ensures that v cannot generate additional flow. The no-fraud constraint ensures that each player forwards as much of its total supply as possible. No-fraud strategies are desirable in the application of financial networks, since they do not allow financial firms to hide or malversate assets while

having unpaid debt (hence the name “no-fraud”)². We assume this property in the strategies for simplicity. Note that even if we drop the no-fraud condition as a constraint, then in flow allocation games (with utilities resulting from monotone strategies defined below) it turns out that every player always has a best response that satisfies the no-fraud condition (see Section 3.1 below).


A *flow* is a vector of edge valuations $\mathbf{f} = (f_e)_{e \in E}$. Given some flow \mathbf{f} , we slightly abuse notation and denote the total supply of v by $f_v = b_v + \sum_{e \in E^-(v)} f_e$. For a *strategy profile* $\mathbf{a} = (\mathbf{a}_e)_{e \in E}$, a flow is called *feasible* if for all edges $(v, w) \in E$ it holds

$$f_{(v,w)} = a_{(v,w)}(f_v). \quad (\text{fixed point constraint}) \quad (4)$$

A node is called *fully saturated* if the feasible flow saturates all outgoing edges, i.e., $f_e = c_e$ for all $e \in E^+(v)$, or equivalently $f_v = c_v^+$. Given some feasible flow \mathbf{f} for a strategy profile \mathbf{a} , the *utility* of player v is defined by $u_v(\mathbf{a}, \mathbf{f}) = \sum_{e \in E^+(v)} f_e$, i.e., v 's goal is to choose a strategy to maximize the total outgoing flow. Capacity and weak flow conservation constraints are not sufficient for a consistent definition of utility. In particular, even if a strategy profile satisfies capacity and weak flow conservation constraints, it may not allow a feasible flow.

EXAMPLE 1. Consider a graph with three nodes $V = \{v_1, v_2, v_3\}$ and edges (v_1, v_2) , (v_2, v_1) and (v_1, v_3) . Let $c_e = 2$ for all edges $e \in E$. For the strategy profile \mathbf{a} , we assume

$$\begin{aligned} v_2 \text{ plays } a_{(v_2,v_1)}(y) &= y \\ v_1 \text{ plays } a_{(v_1,v_2)}(y) &= \begin{cases} y & \text{for } y \leq 1 \\ 0 & \text{for } y > 1 \end{cases} \\ a_{(v_1,v_3)}(y) &= y - a_{(v_1,v_2)}(y) \text{ for } y \geq 0. \end{aligned}$$



Informally, v_1 forwards the total supply to v_2 if it is at most 1. Otherwise, v_1 forwards all supply to v_3 . The fixed supply is $b_{v_1} = 1$ and 0 for the other nodes.

There is no feasible flow for strategy profile \mathbf{a} . Suppose the total supply of v_1 is 1, then v_1 routes all supply to v_2 , who forwards it back to v_1 . The total supply of v_1 must be at least 2, a contradiction. Suppose the total supply of v_1 is more than 1, then v_1 forwards all supply to v_3 , so v_1 only has the fixed supply. The total supply of v_1 must be 1, a contradiction. ■

The main problem with feasible flows in the example is that the strategies are not *monotone*.

Monotonicity. Monotonicity is a natural condition for payment strategies. In a flow allocation game with *monotone strategies*, each player $v \in V$ forwards the flow in a monotone fashion. Monotone strategies are characterized by capacity and weak-flow conservation constraints, and, for every $y, y' \in \mathbb{R}_{\geq 0}$ with $y \geq y'$

$$a_e(y) \geq a_e(y'). \quad (\text{monotonicity constraint}) \quad (5)$$

Monotone strategies have been proposed and studied before by Csóka and Herings [9]. In the following theorem, we recapitulate their main structural insight on feasible flows. Consider a flow-allocation game and a strategy profile \mathbf{a} of monotone strategies. Let \mathcal{F} be the set of feasible flows for \mathbf{a} . We observe that \mathcal{F} is non-empty, i.e., a feasible flow always exists. Moreover, (\mathcal{F}, \leq) forms a lattice with the coordinate-wise comparison. Formally, $\mathbf{f} \leq \mathbf{f}'$ iff $f_e \leq f'_e$ for all $e \in E$; and $\mathbf{f} < \mathbf{f}'$ iff $f_e \leq f'_e$ for all e and $f_e < f'_e$ for at least one edge e .

² For example, Eisenberg and Noe [14] restrict the choice of each player to a specific no-fraud strategy with *pro-rata* payments: Each node v distributes its total supply in proportion to the edge capacities. Formally, for every edge $e \in E^+(v)$ the strategy is fixed to $a_e(y) = \min \left\{ c_e, y \cdot \frac{c_e}{c_v^+} \right\}$.

THEOREM 1 ([9]). *For every strategy profile \mathbf{a} in a flow-allocation game with monotone strategies, the pair (\mathcal{F}, \leq) is a non-empty complete lattice.*

Csoka and Herings state the result only for *integral* monotone strategy profiles \mathbf{a} . We here reiterate the proof to show that it works for all monotone ones. This theorem motivates the study of flow allocation games with monotone strategies. It follows using the Knaster-Tarski theorem.

THEOREM 2 ([46]). *Let (L, \leq) be any complete lattice. Suppose $g : L \rightarrow L$ is order-preserving, i.e., for all $x, y \in L$ we have that $x \leq y$ implies $g(x) \leq g(y)$. Then the set of all fixed points of g is a non-empty, complete lattice with respect to \leq .*

Proof of Theorem 1. Consider $F = \{\mathbf{f} \mid 0 \leq f_e \leq c_e, \forall e \in E\}$, a compact superset of all possible flow vectors. Obviously, (F, \leq) forms a complete lattice with the coordinate-wise comparison defined above. For a given strategy profile \mathbf{a} , the map $g : F \rightarrow F$ with

$$g(\mathbf{f})_{(u,v)} = a_{(u,v)} \left(\sum_{e' \in E^-(u)} f_{e'} + b_u \right) \quad \text{for every } (u, v) \in E \quad (6)$$

is an order-preserving function for every edge $e \in E$, since the strategies are monotone. Obviously, the set of feasible flows \mathcal{F} is the set of fixed-points of g . The result follows by applying the Knaster-Tarski theorem. \square

Clearing states and utilities. We consider feasible flows that arise due to the strategic flow allocation decisions. We determine the utility $u_v(\mathbf{a}, \mathbf{f})$ of a player v by using a feasible flow \mathbf{f} for the given strategy profile \mathbf{a} . We focus on monotone strategies in order to guarantee the existence of at least one feasible flow. This is a necessary condition to make the game well-defined. In many cases, for a fixed strategy profile \mathbf{a} , there is a unique feasible flow \mathbf{f} . However, even in very special cases, there might be infinitely many feasible flows for the same strategy profile³ \mathbf{a} .

Based on these properties, and similarly to the vast majority of the literature, we concentrate on the supremum of the lattice. For every profile \mathbf{a} , we focus on the unique feasible flow that maximizes the total flow in the network. We denote this feasible flow by $\hat{\mathbf{f}}$ and call it the *clearing state* of strategy profile \mathbf{a} . The clearing state $\hat{\mathbf{f}}$ determines the utility of every player in \mathbf{a} , i.e.,

$$u_v(\mathbf{a}) = u_v(\mathbf{a}, \hat{\mathbf{f}}) = \sum_{e \in E^+(v)} \hat{f}_e .$$

Clearly, the clearing state $\hat{\mathbf{f}}$ and the resulting utilities significantly depend on the strategy choices of the nodes in the profile \mathbf{a} .

2.2. Ranking-based strategies. In this paper we are interested in classes of intuitive, expressive, and meaningful monotone strategies. We concentrate on strategies that can be derived via *rankings* [9].

Unit-ranking strategies. In unit-ranking games, we rely on integrality of all values for c_e and b_v , and choose the strategies such that the feasible flows will be integral. We define the parametrized flow functions $a_e(y)$ for every outgoing edge $e \in E^+(v)$ on the non-negative integer numbers $a_e(y) : \mathbb{N}_0 \rightarrow \mathbb{N}_0$. Note that Theorem 1 can also be shown for $a_e(y) : \mathbb{N}_0 \rightarrow \mathbb{N}_0$. The proof is analogous, where we define F on integrals and replace *compact* by *finite*. This ensures that unit-ranking strategies are well-defined in our model. We interpret the flow as being discretized into unsplittable “units” or “particles” of size 1.

³ For example, consider a simple cycle with two nodes v and w . The capacities $c_{(v,w)} = c_{(w,v)} = 1$, the fixed supplies $b_v = b_w = 0$, and the strategies $a_{(v,w)}(y) = a_{(w,v)}(y) = y$. Every flow with $f_{(v,w)} = f_{(w,v)} \in [0, 1]$ is a feasible flow.

Unit-ranking strategies are relevant for the application of flow-allocation games to financial networks. Usually, currencies have some smallest indivisible amount of money. In this way, unit-ranking strategies provide a rich and powerful class of strategies that can be used to express payments strategies in this context.

Edge-ranking strategies. In a flow-allocation game with edge-ranking strategies, each player $v \in V$ forwards its total supply according to a strict and total order over $E^+(v)$, which we represent by a permutation $\pi_v = (e_1, e_2, \dots)$. Player v first allocates the maximum possible flow to edge $e_1 = \pi_v(1)$, then $e_2 = \pi_v(2)$, etc. until all edges are at their capacity or v has no supply left. Formally, $a_{e_i}(y) = \min\{c_{e_i}, \max\{0, y - \sum_{j < i} c_{e_j}\}\}$. The *edge-ranking strategy*⁴ of v is fully described by the ranking π_v , hence we denote a strategy profile in edge-ranking games by $\boldsymbol{\pi} = (\pi_v)_{v \in V}$.

Edge-ranking strategies are a special case of unit-ranking strategies in the sense that every edge-ranking strategy can be written as a unit-ranking strategy. Maybe counterintuitively, every unit-ranking game is also a special edge-ranking game – replacing each edge e with capacity c_e many multi-edges of unit capacity expands a unit-ranking game into an equivalent edge-ranking game. There is a one-to-one correspondence between unit-ranking strategies in the original game and edge-ranking strategies in the expanded game. Intuitively, for a unit-ranking strategy in the original game, a player v assigns the first particle of flow to the multi-edge $\pi_v(1)$, the second particle to $\pi_v(2)$, etc. in the expanded edge-ranking game until all outgoing edges are saturated or v runs out of supply. The expansion of the game implies a pseudo-polynomial blowup in representation size. Nevertheless, the structural equivalence turns out to be very useful for characterizing and analyzing feasible flows and equilibria in unit-ranking games.

Note that the trivial representation of a unit-ranking strategy might be pseudo-polynomial in the size of the original non-expanded game, since edge capacities are given in binary encoding. We will address this issue briefly in Theorem 3 when we discuss polynomial-time computation of equilibria. It turns out that there always exist equilibria with a representation that is polynomial in the input size of the game.

2.3. Equilibria and Social Quality

Equilibrium concepts. We study pure Nash and strong equilibria of flow allocation games. A (*pure*) *Nash equilibrium* in a flow allocation game is a strategy profile \mathbf{a} such that no player v has an incentive to unilaterally deviate from the strategy \mathbf{a}_v . More formally, in a pure Nash equilibrium \mathbf{a} we have $u_v(\mathbf{a}) \geq u_v(\mathbf{a}'_v, \mathbf{a}_{-v})$ for every player $v \in V$ and every strategy \mathbf{a}'_v . Here \mathbf{a}_{-v} denotes the reduced profile composed of all entries of \mathbf{a} except the entries for player v .

For the definition of a *strong equilibrium*, we first define the notion of a profitable deviation of a coalition. A coalition $C \subseteq V$ of nodes has a *profitable deviation* $\mathbf{a}'_C = (\mathbf{a}'_v)_{v \in C}$ if upon joint deviation of C to \mathbf{a}'_C , the resulting utility in the new profile $(\mathbf{a}'_C, \mathbf{a}_{-C})$ is strictly better for every player in C , i.e., $u_v(\mathbf{a}'_C, \mathbf{a}_{-C}) > u_v(\mathbf{a})$ for every $v \in C$. Here \mathbf{a}_{-C} denotes the reduced profile composed of all entries of \mathbf{a} except the ones for players $v \in C$. A strategy profile \mathbf{a} is a *strong equilibrium* if no coalition $C \subseteq V$ has any profitable deviation. Thus, by definition, a strong equilibrium is a Nash equilibrium.

In general, pure Nash or strong equilibria might not exist in a flow allocation game. If they are guaranteed to exist, they might not be unique.

⁴ These strategies have also been termed *singleton liability priority lists* in [26].

Prices of anarchy and stability. In addition to existence and computational complexity, we also quantify the performance of a feasible flow in equilibrium in terms of natural notions of social quality. For any non-negative objective function $Q(\mathbf{a})$ measuring the quality of a strategy profile, we rely on standard notions of *price of anarchy* and *price of stability* to relate the quality in equilibrium to the one that could be obtained in a strategy profile that maximizes Q .

The *price of anarchy* for an equilibrium concept in a game Γ is given by the ratio

$$\max_{\mathbf{a} \in \mathcal{A}_{\text{Eq}}(\Gamma)} \frac{Q(\mathbf{a}^*)}{Q(\mathbf{a})} = \frac{Q(\mathbf{a}^*)}{\min_{\mathbf{a} \in \mathcal{A}_{\text{Eq}}(\Gamma)} Q(\mathbf{a})} . \quad (7)$$

Here $\mathcal{A}_{\text{Eq}}(\Gamma)$ is a set of equilibria (e.g., the set of all Nash equilibria, or the set of all strong equilibria) of the game Γ , and \mathbf{a}^* is a strategy profile maximizing Q . The price of anarchy for a class of games is the largest price of anarchy of any game in the class. The *price of stability* for an equilibrium concept in a game Γ is defined by replacing max with min and vice versa in (7). The price of stability for a class of games is the largest price of stability in any game in the class. Note that both prices of anarchy and stability are at least 1.

Intuitively, an upper bound of ρ on the price of anarchy implies that *every equilibrium* has a quality of at least $Q(\mathbf{a}^*)/\rho$ in *every game* of the class. A lower bound of ρ implies that for *some game* there is *some equilibrium* in that game with quality at most $Q(\mathbf{a}^*)/\rho$. Similarly, an upper bound of ρ on the price of stability implies that *at least one equilibrium* has a quality of at least $Q(\mathbf{a}^*)/\rho$ in *every game* of the class. A lower bound of ρ implies that for *some game* it holds that *every equilibrium* in that game with quality at most $Q(\mathbf{a}^*)/\rho$.

Our main result for unit-ranking games shows the existence of optimal strong equilibria. More in detail, the price of stability is 1 for pure Nash and strong equilibria for every quality function from a large class of so-called *flow-monotone functions* Q (for details see Corollary 1 below).

For the remaining results on prices of anarchy and stability in this paper, we concentrate on the total amount of flow, which is a standard objective for circulation and flow allocation problems. Formally, the *total (amount of) flow* for a strategy profile \mathbf{a} is

$$\text{FLOW}(\mathbf{a}) = \sum_{e \in E} \hat{f}_e = \sum_{v \in V} \sum_{e \in E^+(v)} \hat{f}_e = \sum_{v \in V} u_v(\mathbf{a}) , \quad (8)$$

since the clearing state $\hat{\mathbf{f}}$ determines the utilities for all players. Hence, $\text{FLOW}(\mathbf{a})$ also represents the *utilitarian welfare*.

2.4. Games on DAGs Flow allocation games are designed to analyze incentives in networks with circulation flows. In contrast to previous work [21, 33], incentives in our games are inherently connected to cycles. To see this, we briefly discuss games on directed acyclic graphs (DAGs).

PROPOSITION 1. *In a flow allocation game on a DAG G , every strategy profile is a strong equilibrium.*

Proof. We prove the statement inductively. Consider an arbitrary profile \mathbf{a} . We first discuss a natural algorithm to construct a feasible flow. A DAG contains a node v without incoming edges. Clearly, the incoming assets of v are fixed to b_v . Thus, every strategy \mathbf{a}_v allocates an amount of $\min\{c_v^+, b_v\}$ to the outgoing edges. This is independent of the behavior of v or any other player, so \mathbf{a}_v represents a best response. The algorithm then constructs an equivalent network for the remaining players by (1) removing v and all its outgoing edges, and (2) increasing b_w by $a_{(v,w)}(b_v)$, for every $(v, w) \in E^+(v)$. The new network is again a DAG and contains a node without incoming edges. In this way, the algorithm proceeds until the remaining network contains no edges.

To show that \mathbf{a} is a strong equilibrium, consider any coalition C and the first player $v_1 \in C$ that was processed by our algorithm above. v_1 and all players being processed after v_1 cannot increase

the incoming assets of v_1 by changing their strategies. As such, v_1 has no incentive to deviate from \mathbf{a}_{v_1} . C has no profitable deviation. Therefore, \mathbf{a} is a strong equilibrium. \square

Our arguments do not rely on monotonicity – the algorithm constructs a feasible flow even for non-monotone strategies \mathbf{a} . Moreover, by the same induction, for any given strategy profile \mathbf{a} the feasible flow is unique. The proposition holds even for flow allocation games with non-monotone strategies. Since *every* strategy profile is a strong equilibrium, it is stable against arbitrary deviations of coalitions. Hence, the result continues to hold even when we restrict to games on DAGs with monotone unit- or edge-ranking strategies.

3. Properties of feasible flows. We observe a useful *circulation representation* of feasible flows in flow allocation games and some preliminary results that will be used in the subsequent sections. For edge-ranking games, the circulation representation can be used to describe a polynomial time algorithm that computes the clearing state $\hat{\mathbf{f}}$ in polynomial time.

3.1. Preliminaries on feasible flows and clearing states.

Circulation structure. Given a game Γ on G with strategies \mathbf{a} , we build a *circulation network* G' with strategies \mathbf{a}' and consider an extended game in G' as follows. We add to G an auxiliary node s . For every $v \in V$, we add an auxiliary edge (v, s) with capacity $c_{(v,s)} = \infty$. For every $v \in V$ with $b_v > 0$ we add an auxiliary edge (s, v) with $c_{(s,v)} = b_v$, and choose $b'_v = 0$ in the new game. In this way, the supply of v becomes an incoming flow to v on edge (s, v) . The strategy \mathbf{a}'_s of player s is arbitrary. For every strategy profile \mathbf{a} of the original game, we modify \mathbf{a}_v such that flow not forwarded by some player v will now be a flow on edge (v, s) under \mathbf{a}'_v .

PROPOSITION 2. *For every feasible flow \mathbf{f} for a strategy profile \mathbf{a} , the corresponding flow for the modified strategy profile \mathbf{a}' in the extended game G' can be decomposed and represented as a circulation. The auxiliary node s has an incoming flow of $\sum_{v \in V} b_v$, and all auxiliary edges (s, v) are saturated.*

Proof. The proposition is a simple consequence of fixed point constraint and no-fraud constraint. Non-forwarded flow at firm $v \in V$ in the original game exists only if v saturates all outgoing edges

$$\sum_{e \in E^+(v)} a_e(f_v) = \min\{f_v, c_v^+\} .$$

Moreover, the total sum of non-forwarded flow in the original game is exactly the sum of all fixed supply:

$$\begin{aligned} \sum_{v \in V} b_v &= \sum_{v \in V} b_v + \sum_{v \in V} \sum_{e=(u,v) \in E^-(v)} a_e(f_u) - \sum_{v \in V} \sum_{e=(u,v) \in E^-(v)} a_e(f_u) \\ &= \sum_{v \in V} \left(f_v - \sum_{e \in E^+(v)} a_e(f_v) \right) = \sum_{v \in V} \max\{0, f_v - c_v^+\} . \end{aligned}$$

The non-forwarded flow in the original game gets routed to the auxiliary node s in the extended game. This constitutes the incoming flow of s , i.e., $f_s = \sum_{v \in V} \max\{0, f_v - c_v^+\} = \sum_{v \in V} b_v$, and all auxiliary edges (s, v) are saturated. Overall, by routing the non-forwarded flow in the original game to the auxiliary node s in the extended game, we obtain exact flow conservation at every node. As such, the flow can be decomposed and represented as a circulation. \square

Fully saturated nodes. In flow allocation games Γ , it turns out that the clearing state $\hat{\mathbf{f}}$ is fixed as long as all nodes that are not fully saturated stick to their strategies. Since every strategy satisfies capacity and no-fraud constraints, the forwarded flow of fully saturated nodes remains the same if they have the same total supply, and vice versa. Consequently, strategies of fully saturated nodes have no impact on the clearing state $\hat{\mathbf{f}}$. For any fully saturated node v , every strategy constitutes a best response.

PROPOSITION 3. *For a given flow allocation game, consider any profile \mathbf{a} , the corresponding clearing state $\hat{\mathbf{f}}$, and any fully saturated node v with $\hat{f}_v \geq c_v^+$. Every strategy \mathbf{a}'_v is a best response for v against the other strategies \mathbf{a}_{-v} and results in the same clearing state $\hat{\mathbf{f}}$.*

Proof. Firm v is fully saturated under $\hat{\mathbf{f}}$, thus $\hat{f}_v \geq \sum_{e \in E^+(v)} c_e$ and $\hat{f}_e = c_e$ for all $e \in E^+(v)$. Consider a deviation \mathbf{a}'_v , the resulting profile $\mathbf{a}' = (\mathbf{a}'_v, \mathbf{a}_{-v})$. It suffices to show that $\hat{\mathbf{f}}$ is feasible under \mathbf{a}' . Note that capacity, weak flow conservation and no-fraud constraints ensure that $a'_e(\hat{f}_v) = c_e$ for all $e \in E^+(v)$, which immediately implies the feasibility of $\hat{\mathbf{f}}$ under \mathbf{a}' . \square

No-Fraud Property. We observe that violating the no-fraud constraint is never in the interest of any player in any game.

PROPOSITION 4. *Suppose \mathbf{a} is a strategy profile of monotone strategies that do not necessarily fulfill the no-fraud constraint (3), and strategy \mathbf{a}_v of player v is such that for a value y of total supply we have $\sum_{e \in E^+(v)} a_e(y) < \min\{c_v^+, y\}$. There is a no-fraud strategy \mathbf{a}'_v such that $u_w(\mathbf{a}) \leq u_w(\mathbf{a}'_v, \mathbf{a}_{-v})$ for every player $w \in V$.*

Proof. Consider any no-fraud strategy \mathbf{a}'_v arising from \mathbf{a}_v by increasing the functions a_e arbitrarily such that Condition (3) holds. Consider the clearing states $\hat{\mathbf{f}}$ for \mathbf{a} and $\hat{\mathbf{f}}'$ for $(\mathbf{a}'_v, \mathbf{a}_{-v})$. For the same total supply y of v we know $a_e(y) \leq a'_e(y)$ for every $e \in E^+(v)$. Now consider $\hat{\mathbf{f}}$. If $\hat{\mathbf{f}}$ is a feasible flow for $(\mathbf{a}'_v, \mathbf{a}_{-v})$, then $\hat{\mathbf{f}}' \geq \hat{\mathbf{f}}$, so the utility of every player in $(\mathbf{a}'_v, \mathbf{a}_{-v})$ is weakly improved and we are done.

Otherwise, $\hat{\mathbf{f}}$ is not a feasible flow for $(\mathbf{a}'_v, \mathbf{a}_{-v})$. We apply the straightforward monotone fixed-point iteration based on the map g in Equation (6). Since the space F of possible flow vectors is compact, the iteration converges to a feasible flow for $(\mathbf{a}'_v, \mathbf{a}_{-v})$ that is coordinate-wise at least $\hat{\mathbf{f}}$. Hence $\hat{\mathbf{f}}' \geq \hat{\mathbf{f}}$, so the utility of every player in $(\mathbf{a}'_v, \mathbf{a}_{-v})$ is weakly improved. \square

3.2. Structure and computation of feasible flows in edge-ranking games. For edge-ranking games, we provide a more detailed analysis of the structure of feasible flows in a strategy profile $\boldsymbol{\pi}$. We observed above that every feasible flow $\mathbf{f} \in \mathcal{F}$ is in a one-to-one correspondence to a circulation in the circulation network. Here, we observe that the circulation follows a partial order. We use this structural insight to show that the clearing state $\hat{\mathbf{f}}$ can be computed in polynomial time.

Before we describe the algorithm, we will define the notion of an *active edge*. For a given strategy profile $\boldsymbol{\pi}$, consider the circulation network $G' = (V', E')$ with auxiliary node s and the corresponding auxiliary edges described in the previous section. We also extend $\boldsymbol{\pi}$ to $\boldsymbol{\pi}'$ and define the auxiliary edge (v, s) to be the least preferred edge by node v . For completeness, we treat s as a node with a fixed strategy π'_s over its outgoing auxiliary edges. Due to Proposition 2, s exactly saturates all outgoing edges in every feasible flow. As such, the strategy π'_s has no impact on the set of feasible flows.

Let \mathbf{f} be a feasible flow in a circulation network G' . In games with ranking-based strategies, we can observe the following. Independent of the total supply f_v of some node v , there is a uniquely defined outgoing edge for the next unit of v 's supply. We call this edge the *active edge* of v . If $\mathbf{f} \equiv \mathbf{0}$, the active edge of a node v is $\pi'_v(1)$. If $\sum_{i=1}^{k-1} c_{\pi'_v(i)} \leq f_v < \sum_{i=1}^k c_{\pi'_v(i)}$ for some k , the active edge is $\pi'_v(k)$.

Algorithm 1: TOPCYCLEINCREASE

Input : The circulation network G' , strategy profile π' , auxiliary node s

Output: The circulation \mathbf{f} representing the clearing state for π in G

$f_e \leftarrow 0$ for all $e \in E'$.

while there is a cycle C of active edges under \mathbf{f} **do**

Choose an arbitrary cycle C .

$\delta \leftarrow \min_{e \in C} \{c_e - f_e\}$

$f_e \leftarrow f_e + \delta$ for all $e \in C$.

return \mathbf{f}

In the following, we will describe the TOPCYCLEINCREASE algorithm. For a formal description see Algorithm 1. In the description we distinguish necessary and optional cycles. This classification builds on the lattice structure of feasible flows (c.f. Theorem 1). Since the algorithm computes the supremum of the lattice, it raises flow along all – necessary and optional – cycles.

Necessary cycles. Consider the set of all active edges for $\mathbf{f} \equiv \mathbf{0}$. Every node $v \neq s$ has exactly one outgoing active edge $\pi'_v(1)$. The set of active edges form disjoint cycles with attached trees. Each tree attached to a cycle is rooted in a node from the cycle, and directed towards the cycle C . We define the *orbit* of C by

$$o(C) = \{v \in V \mid \exists v\text{-}u\text{-path of active edges, for some } u \in C\},$$

i.e., the set of nodes v from which we can reach C over active edges. Now, consider the auxiliary node s . Note that $\mathbf{f} \equiv \mathbf{0}$ is a feasible circulation in G' , but $\sum_v f_{(v,s)} = c_s^+ = \sum_v b_v$ is a necessary constraint to ensure that \mathbf{f} represents a feasible flow in the original graph G . However, as long as $\sum_v f_{(v,s)} < \sum_v b_v$ there is an active outgoing edge of s and all nodes (including s) belong to some orbit. Hence, there is a cycle C with $s \in o(C)$. Flow conservation and the monotonicity of strategies implies that some amount of flow of s must eventually reach C . Due to flow conservation in C , a flow of at least $\delta_C = \min\{c_e \mid e \in C\}$ must thus be present on every edge of C . This is a necessary condition in every feasible flow $\mathbf{f} \in \mathcal{F}$.

It is straightforward to inductively apply this argument, thereby obtaining a sequence of necessary cycles C that must be filled with flow $\delta_C = \min\{c_e - f_e \mid e \in C\}$. In particular, when a flow of $f_e = \delta_C$ has been assigned to every edge $e \in C$, the active edge of at least one of the vertices in C changes. This implies that the orbits change, i.e., the orbit $o(C)$ partitions into new suborbits, or parts that get attached to other orbits. Note that once a vertex v becomes fully saturated and all regular outgoing edges are filled, the active edge becomes (v, s) .

Orbits present at the same time are always mutually disjoint. Thus, for two existing orbits $o(C_1)$ and $o(C_2)$, pushing flow along C_1 can never change the active edges of $o(C_2)$. Hence, it is necessary that all the cycles C , where s eventually appears in the orbit, must get assigned a flow increase δ_C in order to reach $\sum_v f_{(v,s)} = \sum_v b_v$. Once we reach a flow \mathbf{f} where $\sum_v f_{(v,s)} = \sum_v b_v$, an “orbit” $o(s)$ emerges composed of a tree rooted in s . At this point, we have indeed constructed a feasible flow, which by induction is the unique minimal circulation in G' that represents a feasible flow in G .

Optional cycles. In the following, we characterize the structure of all other feasible flows $\mathbf{f} \in \mathcal{F}$ by applying similar observations. Fix some flow \mathbf{f} in the circulation network that is also feasible in G , i.e., $\sum_v f_{(s,v)} = \sum_v b_v$ and consider the set of all active edges.

Suppose there is a cycle C with some orbit $o(C)$, i.e., a set of nodes v that are not attached to the tree rooted in s . In this case, one can push a non-zero amount of flow along C , i.e., increase f_v by a strictly positive amount for every $v \in C$. This obviously yields a new feasible clearing state.

When an edge becomes saturated, the set of active edges changes and the orbit $o(C)$ disappears, i.e., gets split up as explained above (new suborbits, parts attached to other orbits, parts attached to the tree rooted in s).

Note that there might exist multiple orbits at the same time and, thus, multiple possibilities to extend \mathbf{f} by increasing flow along a cycle. However, as observed above, orbits present at the same time are always mutually disjoint, and pushing flow along cycle C_1 can never change the active edges in an orbit $o(C_2)$ present at that time. Now consider some vertex $v \in o(C)$. In order to create some feasible flow \mathbf{f}' with $f'_v > f_v$ it is necessary to push flow along C until f_v is reached (if $v \in o(C) \cap C$) or $o(C)$ disappears and splits up (if $v \in o(C) \setminus C$). Since the flow adjustments monotonically increase all flow values, there is a one-to-one correspondence between sets of cycles with flow increase and the feasible flows.

Note that the cycles chosen for flow increase form a partial order: A cycle C' might not be present in the beginning – there might be predecessor-cycles C that have to be filled up to δ_C to break an existing orbit $o(C)$, change some of the active edges, and make C' appear. In the argumentation above, it can be seen that the set of predecessor cycles $\text{pred}(C')$ for some cycle C' is uniquely defined resulting from the ranking of edges in π .

The arguments above imply a natural algorithm to compute $\hat{\mathbf{f}}$, which is similar in spirit to the classic Top-Trading-Cycles algorithm for house allocation [43]. The TOPCYCLEINCREASE algorithm iteratively raises flow along cycles among the active edges in the circulation network G' . Thereby it computes the unique maximal feasible flow $\hat{\mathbf{f}}$ from the lattice. The algorithm runs in polynomial time – in every round it increases the flow along a cycle C by δ_C . At this point at least one edge (from G or auxiliary) becomes saturated. In terms of the original network $G = (V, E)$, the algorithm needs at most $O(|V| + |E|)$ rounds. Each round can easily be implemented in strongly polynomial time.

PROPOSITION 5. *In every edge-ranking game, the TOPCYCLEINCREASE algorithm computes the clearing state $\hat{\mathbf{f}}$ in strongly polynomial time.*

Since unit-ranking games can be cast as edge-ranking games with unit-capacity multi-edges, the algorithm can be applied in unit-ranking games as well. Note that the running time does not necessarily remain polynomial due to the pseudo-polynomial blowup in representation size.

4. Unit-ranking games.

4.1. Existence and computation of equilibria. In this section we consider equilibria in unit-ranking games. Our first result is that in every unit-ranking game there is a strong equilibrium that can be computed in polynomial time. Moreover, for a large class of quality functions Q , we can guarantee the existence of a strong equilibrium that is optimal with respect to Q . In particular, every such equilibrium profile of unit-ranking strategies can be represented compactly. In the next theorem, we prove the existence and representation result, and we show that the price of stability is 1 for FLOW. The extension to more general quality functions Q is discussed subsequently.

Consider an arbitrary flow allocation game and a circulation that maximizes the total flow in the circulation network G' . We show that this circulation can be expressed as a clearing state of a strong equilibrium in unit-ranking strategies.

THEOREM 3. *For every unit-ranking game, there is a strong equilibrium that maximizes the total amount of flow. The strong equilibrium can be computed in polynomial time, and the equilibrium profile can be represented in polynomial space.*

Proof. Consider the circulation network $G' = (V, E')$. For a moment, assume this is a standard flow network without strategic flow allocation. Consider an optimal circulation \mathbf{f}^* that maximizes the total flow value, i.e., it maximizes the sum of flow on all edges. This implies, in particular, that

it saturates all outgoing auxiliary edges from s . Clearly, \mathbf{f}^* yields an upper bound on the achievable total amount of flow (denoted by $\text{FLOW}(\mathbf{f}^*)$) in any strategy profile \mathbf{a} of the game

$$\sum_{e \in E'} f_e^* = 2 \sum_{v \in V} b_v + \sum_{v \in V} \sum_{e \in E^+(v)} f_e^* = 2 \sum_{v \in V} b_v + \text{FLOW}(\mathbf{f}^*) \geq 2 \sum_{v \in V} b_v + \text{FLOW}(\mathbf{a}) .$$

\mathbf{f}^* can be computed in strongly polynomial time [45]. Since all edge capacities are integral, we can assume all f_e^* are integral.

We now turn this circulation into a clearing state for a carefully chosen strategy profile \mathbf{a}^* of unit-ranking strategies. We will choose \mathbf{a}^* such that it can be compactly represented by *threshold-ranking* strategies.

In a threshold-ranking strategy, every firm v chooses a permutation π_v over $E^+(v)$ and thresholds τ_e . The interpretation of threshold-ranking strategies is that node v first assigns τ_e particles to every edge $e \in E^+(v)$, sequentially in the order given by π_v . Then, it assigns the remaining $c_e - \tau_e$ particles to every edge in the order given by π_v . That is, v first considers edge $\pi_v(1)$ and forwards the first $\tau_{\pi_v(1)}$ particles to this edge. The next $\tau_{\pi_v(2)}$ particles are forwarded to edge $\pi_v(2)$ etc. until $\sum_{j=1}^{|E^+(v)|} \tau_{\pi_v(j)}$ particles are sent to the edges (or v runs out of flow). Then, the remaining $c_{\pi_v(1)} - \tau_{\pi_v(1)}$ particles are forwarded to edge $\pi_v(1)$, then the next $c_{\pi_v(2)} - \tau_{\pi_v(2)}$ particles to $\pi_v(2)$ etc. Clearly, threshold-ranking strategies are more general than edge-ranking strategies. They constitute a special class of unit-ranking strategies with compact representation.

We choose \mathbf{a}^* as follows. Every firm v chooses an arbitrary permutation π_v over $E^+(v)$ and sets $\tau_e = f_e^*$. It is easy to see that in \mathbf{a}^* , the optimal circulation \mathbf{f}^* corresponds to the clearing state. Let us prove that \mathbf{a}^* is a strong equilibrium, i.e., that no coalition $C \subseteq V$ has a profitable deviation.

Suppose for contradiction that there is a coalition C with a profitable deviation. Examine the new profile $(\mathbf{a}'_C, \mathbf{a}^*_{-C})$ and assume $\hat{\mathbf{f}}'$ is the clearing state. Consider a node $v \in C$. Since $u_v(\mathbf{a}'_C, \mathbf{a}^*_{-C}) > u_v(\mathbf{a}^*)$, there must be strictly more outgoing flow from v in the new profile. Due to the no-fraud condition, this can only happen if v also has strictly more *incoming flow* in the new profile. Hence, there is an incoming edge $e = (w, v) \in E^-(v)$ with $\hat{f}'_e > f_e^*$. Now consider node w . If $w \in C$, then $u_w(\mathbf{a}'_C, \mathbf{a}^*_{-C}) > u_w(\mathbf{a}^*)$, so by the same reasoning there is again some incoming edge in $E^-(w)$ that has strictly more flow in $\hat{\mathbf{f}}'$. Otherwise, if $w \notin C$, then w still plays the strategy \mathbf{a}^*_w . Due to monotonicity, a higher flow on (w, v) can only occur if w has larger total supply. Thus, there is again some incoming edge in $E^-(w)$ that has strictly more flow in $\hat{\mathbf{f}}'$.

We can repeat this argument indefinitely. As such, there must be a cycle of edges that all have more flow in $\hat{\mathbf{f}}'$ than in \mathbf{f}^* . Such a cycle can be used to increase the circulation, which contradicts that \mathbf{f}^* is an optimal circulation in G' . \square

REMARK 1. For the profitable deviation, we can even allow arbitrary continuous strategies and any choice of clearing state for the deviation profile. This applies even for games with non-monotone and fraud strategies, as long as a feasible flow exists and the clearing state is chosen arbitrarily among the feasible flows that are not weakly dominated in terms of coordinate-wise comparison.

REMARK 2. If we consider deviations that weakly improve the coalition (i.e., $u_v(\mathbf{a}'_C, \mathbf{a}^*_{-C}) \geq u_v(\mathbf{a}^*)$ for all $v \in C$ and $u_w(\mathbf{a}'_C, \mathbf{a}^*_{-C}) > u_w(\mathbf{a}^*)$ for at least one $w \in C$), it is a simple exercise to see that there are unit-ranking games, in which no such (often termed “super-strong”) equilibrium exists.

As mentioned before, the result can be generalized quite substantially beyond FLOW to a large class containing various quality functions $Q(\mathbf{a})$. Let \mathcal{A} be the set of all possible strategy profiles of a unit-ranking game. Let $\mathbf{a}, \mathbf{a}' \in \mathcal{A}$ be two strategy profiles and $\hat{\mathbf{f}}$ and $\hat{\mathbf{f}}'$ the clearing states in \mathbf{a} and \mathbf{a}' , respectively. Recall the coordinate-wise comparison of flows used in Theorem 1. The function $Q: \mathcal{A} \rightarrow \mathbb{R}$ is *flow-monotone* if $\hat{\mathbf{f}} > \hat{\mathbf{f}}'$ implies $Q(\mathbf{a}) \geq Q(\mathbf{a}')$.

COROLLARY 1. *For every unit-ranking game and every flow-monotone social welfare function Q , there is a strong equilibrium that maximizes Q . The equilibrium profile can be represented in polynomial space.*

Proof. The proof uses the same argument as above. Consider the optimal value $Q_{\max} = \max_{\mathbf{a} \in \mathcal{A}} Q(\mathbf{a})$ and the set $\mathcal{A}_{\max} = \{\mathbf{a} \mid Q(\mathbf{a}) = Q_{\max}\}$ of optimal strategy profiles w.r.t. Q . Consider a profile $\mathbf{a}^* \in \mathcal{A}_{\max}$ such that the clearing state $\hat{\mathbf{f}}$ is not coordinate-wise dominated by a clearing state of another optimal profile, i.e., there is no $\mathbf{a}' \in \mathcal{A}_{\max}$ with $\hat{\mathbf{f}}' > \hat{\mathbf{f}}^*$. We show that \mathbf{a}^* is a strong equilibrium.

For contradiction, suppose there is a coalitional deviation in \mathbf{a}^* . By the same arguments as in the proof above, the deviation implies that there is a cycle that allows to increase the flow. Thus, there is a circulation $\mathbf{f} > \hat{\mathbf{f}}^*$ in the flow network G' . By constructing a strategy profile \mathbf{a} using threshold-ranking strategies, \mathbf{f} can be turned into the clearing state of \mathbf{a} . Since Q is flow-monotone, $Q_{\max} = Q(\mathbf{a}^*) \leq Q(\mathbf{a})$. Hence, $\mathbf{a} \in \mathcal{A}_{\max}$ and $\mathbf{f} > \hat{\mathbf{f}}^*$, a contradiction to the choice of \mathbf{a}^* .

Finally, by using threshold-ranking strategies the equilibrium profile has a compact representation. \square

The corollary implies that the price of stability for strong equilibria is 1 for a very wide range of natural quality functions. For example, instead of $\text{FLOW}(\mathbf{a})$ (representing utilitarian welfare (8)) we might prefer to express the quality of a strategy profile by

- *egalitarian welfare* $\text{EW}(\mathbf{a}) = \min_{v \in V} u_v(\mathbf{a})$, i.e., the minimum utility of any player in the network, or
- *Nash social welfare* $\text{NSW}(\mathbf{a}) = \left(\prod_{v \in V} u_v(\mathbf{a})\right)^{1/n}$, i.e., the geometric mean of player utilities, or
- the *number of fully saturated nodes* $\text{FSN}(\mathbf{a}) = |\{v \in V \mid \hat{f}_v = c_v^+\}|$, which in financial networks corresponds to the number of solvent firms, or
- any monotone transformation or combination of the above functions.

Corollary 1 implies the existence of an optimal strong equilibrium for all these functions.

While we can compute any strong equilibrium in polynomial time, for some functions Q an *optimal* strong equilibrium can be NP-hard to compute. A simple reduction shows that for each of the objective functions in the list above it is NP-hard to compute an optimal strategy profile and, consequently, also an optimal strong equilibrium.

THEOREM 4. *In unit-ranking games, it is strongly NP-hard to compute a strategy profile that maximizes EW, NSW, or FSN.*

Proof. We start by considering the problem of optimizing FSN. This task can be at least as hard as deciding EXACT COVER BY 3-SETS (X3C). In an instance of X3C, we have a set R of $3k$ elements for an integer $k \geq 1$ and a set $\mathcal{S} \subseteq 2^E$ of m triplets (i.e., $|S| = 3$ for each $S \in \mathcal{S}$). The goal is to decide if there are k non-overlapping sets in \mathcal{S} .

For the reduction, we construct a game containing a source node v with $b_v = 3k$. Each $S \in \mathcal{S}$ is a node, and for each $r \in R$ there are three nodes r^1, r^2, r^3 . We have edges (v, S) of capacity $3k + 1$, for every $S \in \mathcal{S}$, edges (S, r^1) of capacity 1, for every $r \in R$, and two edges $(r^1, r^2), (r^2, r^3)$ of capacity 1 for every $r \in R$. Note that all r^3 are fully saturated, and v can never be fully saturated. By routing the flow from v via the sets to the elements, we can fully saturate at least the $9k$ nodes corresponding to elements in R . If nodes of overlapping sets are fully saturated, we route a flow of $x \geq 2$ to some element node r^1 , which implies that $2(x - 1)$ nodes for $x - 1$ other elements $r' \in R$ are not fully saturated. This does not happen in an optimal profile if (and only if) the X3C instance is a yes-instance. More formally, an optimal profile saturates at least $9k + k$ nodes for elements and sets if and only if there are k non-overlapping sets in \mathcal{S} .

To prove the result for objectives EW and NSW, we introduce two auxiliary nodes S_a and r_a . For each $S \in \mathcal{S}$, we add edges (S, S_a) and (S_a, S) with capacity 1. W.l.o.g. we can assume that all

nodes S give first priority to the edge (S, S_a) . Consequently, each S has a positive utility of at least 1. S_a has utility $|S|$. Similarly, for each node r^3 , we add edges (r^3, r_a) and (r_a, r^3) with capacity 1. As such, all nodes r^3 can be assumed to have a utility of at least 1. r_a has utility $3k$.

Now the problem of optimizing EW or NSW reduces to deciding whether all nodes r^1, r^2 can simultaneously obtain positive utility, since otherwise the objective value will be 0. This property, however, is equivalent to choosing k non-overlapping sets to route the total demand of $3k$ from v to the $3k$ nodes r^1 . As such, a positive value for EW or NSW can be obtained if and only if the X3C instance is a yes-instance. \square

Irrespective of the quality function, even computing a *best-response strategy* for a single node v in a strategy profile can be strongly NP-hard, since best responses can provide answers to computationally hard decision problems. This holds even in games without fixed supply and with edge capacities in $\{0, 1\}$.

THEOREM 5. *For a given strategy profile \mathbf{a} of a unit-ranking game with $b_v = 0$ for all $v \in V$ and $c_e \in \{0, 1\}$ for all $e \in E$, deciding whether a given node v has a best response resulting in utility at least k is strongly NP-complete.*

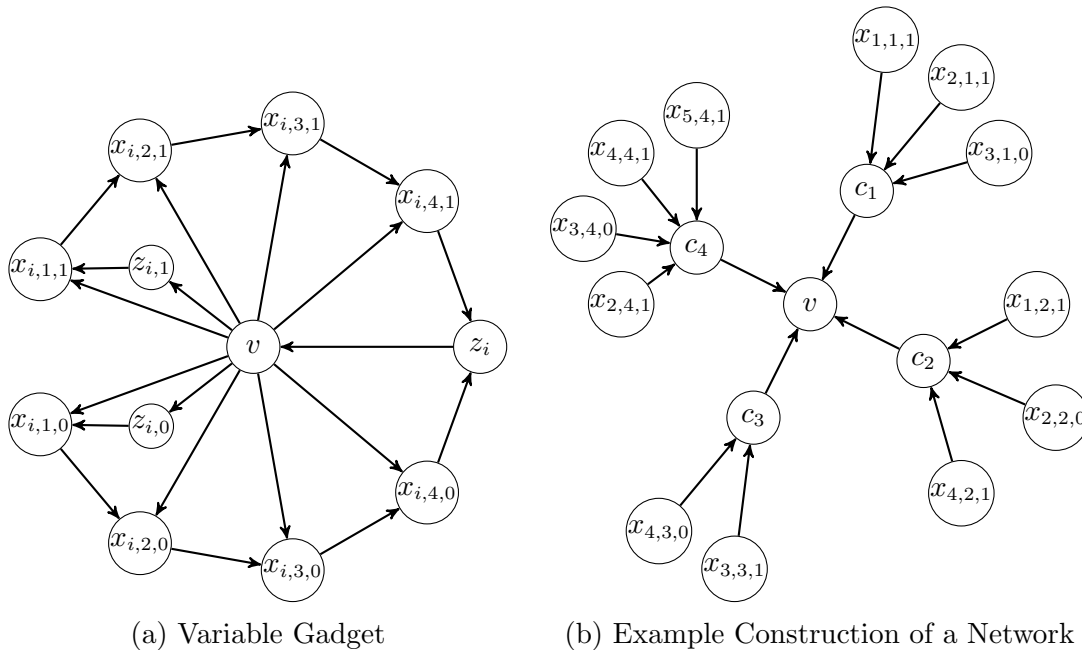


FIGURE 1. Structures used in the proof of Theorem 5.

Proof. For unit-ranking games with edge capacities in $\{0, 1\}$ and fixed supplies equal to 0, the decision problem is obviously contained in NP: We can represent every unit-ranking strategy as a ranking over edges. Then, our TOPCYCLEINCREASE algorithm to compute $\hat{\mathbf{f}}$ discussed in Section 3.2 runs in polynomial time and we can efficiently verify the utilities for all players.

For strong NP-hardness, suppose we are given an instance I of SATISFIABILITY in conjunctive normal form with n variables and m clauses. We construct a unit-ranking game in edge-ranking representation with a node v and a strategy profile π_{-v} for the other players such that the following holds: There is a strategy π_v with utility $u_v(\pi_v, \pi_{-v}) \geq k' + n$ if and only if I has a variable assignment that fulfills at least k' clauses.

We construct the game as follows. We denote the variables of I by x_1, \dots, x_n and the clauses by C_1, \dots, C_m . For each variable x_i we create nodes $x_{i,j,0}$ and $x_{i,j,1}$ for all $j \in \{1, \dots, m\}$, as well as a

node z_i . For each clause C_j we add a clause node c_j . In addition, there is a separate node v , for which we strive to find a best response.

For each clause C_j , we add a unit-capacity edge from $x_{i,j,0}$ to c_j if x_i appears as $\neg x_i$ in C_j and from $x_{i,j,1}$ to c_j if it appears as x_i in C_j . Flow incoming to c_j will eventually indicate a literal that fulfills the clause C_j . There is an edge (c_j, v) for all $j \in \{1, \dots, m\}$. We will show below that this edge ensures that satisfying clause C_j adds exactly one unit to the total supply of v .

For each variable x_i , we add a variable gadget. It consists of nodes $v, x_{i,j,0}$ and $x_{i,j,1}$ for all $j \in \{1, \dots, m\}$, as well as auxiliary nodes $z_i, z_{i,0}$ and $z_{i,1}$. There are unit-capacity edges $(v, x_{i,j,0})$ and $(v, x_{i,j,1})$ for all $j \in \{1, \dots, m\}$, edges $(x_{i,j,0}, x_{i,j+1,0}), (x_{i,j,1}, x_{i,j+1,1})$ for $j \in \{1, \dots, m-1\}$, and edges $(v, z_{i,0}), (z_{i,0}, x_{i,1,0})$ and $(v, z_{i,1}), (z_{i,1}, x_{i,1,1})$. Firm z_i has edges $(x_{i,m,0}, z_i), (x_{i,m,1}, z_i)$, and (z_i, v) . An example for the gadget that is constructed for a variable x_i and $m = 4$ is depicted in Fig. 1a. Note that for every strategy of v , there is at most one cycle emerging in this gadget, since all cycles must include the outgoing edge of z_i .

In Fig. 1b, we show an example of the network without the variable gadgets for the SATISFIABILITY instance $I = (x_1 \vee x_2 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_4) \wedge (x_3 \neg x_4) \wedge (x_2 \vee \neg x_3 \vee x_4 \vee x_5)$.

We construct a strategy profile π_{-v} as follows. Observe that nodes c_1, \dots, c_m and x_1, \dots, x_n each have a single outgoing edge, their strategies are trivial. If a node $x_{i,j,0}$ or $x_{i,j,1}$ has multiple outgoing edges, it always prioritizes the edge to nodes $x_{i,j+1,0}$ and $x_{i,j+1,1}$, respectively, or to z_i if $j = m$.

In the following, we argue that there is a best response of v with utility $k' + n$ if and only if there is a variable assignment such that k' clauses are fulfilled in I . Suppose there is a variable assignment such that k' clauses are fulfilled. Fix this assignment, and for every satisfied clause c choose a single literal l_c that evaluates to true in the clause. Let v choose the following strategy: First, prioritize edges $(v, x_{i,1,0})$ if $x_a = \text{false}$ in the assignment and $(v, x_{i,1,1})$ if $x_a = \text{true}$ in the assignment. All these edges will close a cycle $(v, x_{i,1,0}, x_{i,2,0}, \dots, x_{i,m,0}, x_i, v)$ or $(v, x_{i,1,1}, x_{i,2,1}, \dots, x_{i,m,1}, x_i, v)$. After that, for all clause-fulfilling literals l_c prioritize the edges $(v, x_{i,j,0})$ if $c = C_j$ and $l_c = \neg x_i$ in any order. Prioritize the edges $(v, x_{i,j,1})$ if $c = C_j$ and $l_c = x_i$. All these edges close a cycle via the clause node c_j , leading to a total inflow of $k' + n$.


For showing the other direction, we observe the following structural property for all variables x_i . If there is some flow on an edge $(x_{i,m,0}, z_i)$ there cannot be any flow on edge $(x_{i,m,1}, z_i)$, and vice versa. We conclude that flow on some edge $(x_{i,j,0}, c_j)$ implies flow on edge $(x_{i,j,0}, x_{i,j+1,0})$ (since it has a higher priority), and $(x_{i,m,0}, z_i)$, and thus no flow on $(x_{i,j',1}, c_{j'})$ for all j' . Analogously, we observe that any flow on some edge $(x_{i,j,0}, c_j)$ implies no flow on edges $(x_{i,j',1}, c_{j'})$.

We observe that if there is a best response of node v with inflow equal to $k' + n$, the node has $k' + n$ incoming edges that carry flow. At most n of these edges can be (z_i, v) -edges, so there are at least k' clause-edges (c_j, v) that carry flow. Thus, all these k' clause nodes receive incoming flow. If this flow for some clause c_j comes from a node $x_{i,j,0}$, we know by the observation above that no edge $x_{i,j',1}$ carries flow. Thus we can set the variable x_i to false, which fulfills clause C_j . Applying this and analogous operations for flow on an edge $(x_{i,j,1}, v)$ yields a variable assignment which fulfills at least k' clauses. \square

4.2. Total flow in equilibrium. In this section, we analyze the quality of pure Nash and strong equilibria in unit-ranking games, mostly in terms of the FLOW objective. In the last section we observed that the prices of stability for Nash and strong equilibria in unit-ranking games are both 1. We here bound the prices of anarchy for Nash and strong equilibria. The total flow depends crucially on the emergence of cycles in the strategy profile. This requires an effort that is inherently coalitional. As such, it might be unsurprising that there are games in which the worst Nash equilibrium may fail to provide any reasonable fraction of the optimal total flow.

PROPOSITION 6. *The price of anarchy for Nash equilibria in terms of FLOW is unbounded, even in unit-ranking games without fixed supplies.*

Proof. Consider the game depicted on the right. All edges have capacity 1, fixed supplies are $b_v = 0$ for all nodes v . Consider π with $\pi_1 = (e_1, e_3)$ and $\pi_2 = (e_2, e_4)$. It is a pure Nash equilibrium with $\text{FLOW}(\pi) = 0$. No unilateral deviation can close a cycle and increase the value of $\hat{\mathbf{f}}$. The optimal solution π^* with $\pi_1^* = (e_3, e_1)$ and $\pi_2^* = (e_4, e_2)$ has $\text{FLOW}(\mathbf{a}^*) = 2$. \square



To analyze the quality of strong equilibria, we again consider the unit-ranking game in the form of unit-capacity multi-edges. Consider an optimal circulation \mathbf{f}^* of maximum social welfare in the circulation network G' . Since we have unit-capacity edges, we can assume that the optimal circulation has binary flows on each edge. Let $\mathcal{C}(\mathbf{f}^*) = \{C_1, \dots, C_k\}$ be a decomposition of \mathbf{f}^* into cycles of unit flow. We denote by

$$d = \min_{\mathbf{f}^*, \mathcal{C}(\mathbf{f}^*)} \max_{C_i \in \mathcal{C}(\mathbf{f}^*)} |C_i|$$

the min-max size of any cycle, in any decomposition $\mathcal{C}(\mathbf{f}^*)$ of any optimal circulation \mathbf{f}^* .

THEOREM 6. *In unit-ranking games, the price of anarchy for strong equilibria in terms of FLOW is at most d .*

Proof. Consider an optimal circulation \mathbf{f}^* and a decomposition $\mathcal{C}(\mathbf{f}^*)$ such that all flow cycles $C_i \in \mathcal{C}(\mathbf{f}^*)$ have size at most $|C_i| \leq d$. As observed in the proof of Theorem 3, this circulation yields the total flow of an optimal strategy profile \mathbf{a}^* , i.e.,

$$\text{FLOW}(\mathbf{a}^*) = \sum_{v \in V} \sum_{e \in E^+(v)} f_e^* = \sum_{C_i \in \mathcal{C}(\mathbf{f}^*)} |C_i| - 2 \sum_{v \in V} b_v \leq \sum_{C_i \in \mathcal{C}(\mathbf{f}^*)} d.$$

Now consider any strong equilibrium \mathbf{a} in the unit-ranking game with clearing state $\hat{\mathbf{f}}$. The flow $\hat{\mathbf{f}}$ can be assumed to have binary edge flows. Suppose there is a cycle $C_i \in \mathcal{C}(\mathbf{f}^*)$ such that $a_e(\hat{\mathbf{f}}_v) = 0$ for all $e = (v, w) \in C_i$. Then the nodes in this cycle have an incentive to jointly deviate and place the edges of C_i on the first position in their ranking. Then the clearing state $\hat{\mathbf{f}}$ will emerge as before, adding a flow of 1 along the cycle C_i . This is a profitable deviation for the nodes of C_i .

Consequently, for every cycle $C_i \in \mathcal{C}(\mathbf{f}^*)$ there must be at least one edge $e = (u, v) \in C_i$ such that $a_e(\hat{\mathbf{f}}_u) = 1$. Thus, the total flow in the strong equilibrium \mathbf{a} is

$$\text{FLOW}(\mathbf{a}) \geq \sum_{C_i \in \mathcal{C}(\mathbf{f}^*)} 1$$

and, hence, the ratio is at most d . \square

PROPOSITION 7. *For every $d \geq 2$, there is a unit-ranking game in which the price of anarchy for strong equilibria in terms of FLOW is at least $d - 1$.*

Proof. The game is given by a graph G with $d + (d - 1)(d - 2)$ nodes. G is constructed as follows. The nodes v_1, \dots, v_d are called *central nodes* and they form a cycle of length d . For each $i = 1, \dots, d - 1$, there are nodes $(v_{i,j})_{j=1, \dots, d-2}$ that form additional cycles of length d with the edge (v_i, v_{i+1}) . Thus, the set of edges is given by

$$E = \{(v_i, v_{i+1}) \mid i \in \{1, \dots, d-1\}\} \cup \{(v_d, v_1)\} \\ \cup \bigcup_{i=2, \dots, d} ((v_i, v_i^1) \cup \{(v_i^j, v_i^{j+1}) \mid j = 1, \dots, d-3\} \cup (v_i^{d+2}, v_{i-1})).$$




FIGURE 2. A unit-ranking game with $d=5$ and a price of anarchy for strong equilibria of $d-1=4$.

All nodes have fixed supply of 0. All edges have unit capacity. An example of the instance with $d=5$ is depicted in Fig. 2. Observe that only nodes v_i , $i=2, \dots, d$ have multiple outgoing edges and, thus, these nodes are the only ones with a non-trivial strategy choice.

Since all edges have capacity 1, we will view this game equivalently as an edge-ranking game. We claim that the strategy profile $\pi_i = ((v_i, v_{i+1}), (v_i, v_i^1))$, for $i=2, \dots, d-1$, and $\pi_d = ((v_d, v_1), (v_d, v_d^1))$ is a strong equilibrium in the game. In order to see this, let $\hat{\mathbf{f}}$ be the clearing state of π . Note that

$$\hat{\mathbf{f}}_v = \begin{cases} 1 & \text{if } v = v_1, v_2, \dots, v_d, \\ 0 & \text{otherwise,} \end{cases}$$

that is, $\text{FLOW}(\pi) = d$. Now suppose there is a non-empty coalition of nodes $C \subseteq (v_2, \dots, v_d)$ that all strictly increase their utility by a joint deviation. Note that v_d only has a single incoming edge that is saturated in π , so $v_d \notin C$. Thus, the cycle $v_d, v_d^1, \dots, v_d^{d-2}, v_{d-1}$ cannot carry any flow. We conclude that v_{d-1} has only a single edge that can carry flow. Iterating this argument yields $C = \emptyset$, a contradiction.

In contrast, the strategy profile with $\pi_i = ((v_i, v_i^1), (v_i, v_{i+1}))$ for $i=2, \dots, d-1$ and $\pi_d = ((v_d, v_d^1), (v_d, v_1))$ has social welfare of $(d-1)d$. Thus, the price of anarchy for strong equilibria in this instance is at least $d-1$. \square

For completeness, let us also provide a cumulative characterization of prices of anarchy for pure Nash and strong equilibria for the three other quality functions mentioned in the previous section.

PROPOSITION 8. *In unit-ranking games, the prices of anarchy for pure Nash and strong equilibria in terms of*

- EW are unbounded,
- NSW are unbounded,
- FSN are exactly $n-1$.

The lower bounds apply even in games with $d=3$.

Proof. Consider the game depicted on the right.





FIGURE 3. An edge-ranking game without a pure Nash equilibrium.

All edges have capacity 1, fixed supplies are $b_v = 0$ for all nodes. Only v_1 has a non-trivial strategy choice, $\pi_1 = (e_1, e_2)$ or $\pi_2 = (e_2, e_1)$. Both choices yields a utility of 1 for v_1 , so both are a best response for v_1 . Since v_1 is the only player with a choice, both options represent a strong equilibrium. Both EW and NSW are 0 in the strong equilibrium when v_1 plays π_1 and 1 in the optimal state with π_2 . Hence, the prices of anarchy for both pure Nash and strong equilibria are unbounded for both these objectives. In the optimal state, the unique flow cycle has length 3.



Towards the FSN objective, we note that in every strategy profile \mathbf{a} of every unit-ranking game there must be at least one node that is fully saturated. Otherwise, every node has an outgoing active edge, and hence there must be cycle of active edges – a contradiction to the property that the clearing state is maximal. Clearly, since there can be at most n saturated nodes, the prices of anarchy for pure Nash and strong equilibria are at most n .

More precisely, suppose there is a state \mathbf{a}^* with $\text{FSN}(\mathbf{a}^*) = n$. In Proposition 3 we proved that the strategy choices of saturated nodes have no impact on the clearing state. As such, if *all* nodes are fully saturated, then for each node her utility is independent of the entire strategy profile. In this case, prices of anarchy for pure Nash and strong equilibria are 1. The prices are non-trivial only if every strategy profile has at least one non-saturated node. Hence, they can be at most $n - 1$.

To show the lower bound of $n - 1$, we slightly extend the construction above. We add nodes v_4, \dots, v_n , each with two edges (v_1, v_i) and (v_i, v_2) of capacity 1. Edges (v_1, v_2) and (v_2, v_1) get an increased capacity of $n - 2$. Again, v_1 is the only player with a non-trivial strategy choice. It is easy to see that for any strategy, v_1 obtains a utility of $n - 2$. As such, every strategy is a best response, and every state represents a pure Nash and a strong equilibrium. In the worst equilibrium \mathbf{a} , v_1 exchanges $n - 2$ units of flow with v_2 . v_2 gets saturated and $\text{FSN}(\mathbf{a}) = 1$. In the optimal state \mathbf{a}^* , v_1 routes $n - 2$ units of flow in single units via v_3, \dots, v_n to v_2 . Then v_2, \dots, v_n get saturated and $\text{FSN}(\mathbf{a}^*) = n - 1$. Note that in \mathbf{a}^* every cycle has length 3. \square

5. Edge-ranking games.

5.1. Existence and computation of equilibria. With unit-ranking strategies we assume that nodes have flexibility in allocation of single particles. In this section, we focus on strategies, in which nodes simply rank their outgoing edges and allocate flow in order of this ranking until they run out of supply or all edges are saturated. In contrast to unit-ranking games, the restriction to rankings over edges (with different capacity) can destroy the existence of (optimal) stable states. In fact, there are even games without a single pure Nash equilibrium.

PROPOSITION 9. *There is an edge-ranking game without a pure Nash equilibrium.*

Proof. Consider the game in Fig. 3. The capacities of the edges are depicted next to the edges. Nodes v_2 and v_3 each have fixed supply of 2, the other nodes have fixed supply 0. Nodes v_1, v_2 and

v_3 are the only ones with multiple outgoing edges and, thus, a non-trivial strategy choice. Due to the symmetry of the graph, we can assume w.l.o.g. $\pi_{v_1} = ((v_1, v_4), (v_1, v_6))$. There are two possible strategy choices for each of the nodes v_2 and v_3 . Checking all four resulting strategy profiles yields the following utility matrix for nodes v_2 and v_3 :

	π_{v_3}	$((v_3, v_1), (v_3, v_7))$	$((v_3, v_7), (v_3, v_1))$
π_{v_2}			
$((v_2, v_1), (v_2, v_5))$		4	3
		4	4
$((v_2, v_5), (v_2, v_1))$		2	3
		5	3

Inspecting the utilities, we see that there is no pure Nash equilibrium. \square

The next theorem shows that a number of natural decision and optimization problems in edge-ranking games are indeed computationally intractable. Note that for unit-ranking games, these problems are either trivial (a strong equilibrium always exists) or can be solved in polynomial time (compute a strong equilibrium that represents a profile with maximum total flow).

THEOREM 7. *Given an edge-ranking game the following problems are strongly NP-hard:*

1. *Deciding whether a pure Nash equilibrium exists or not.*
2. *Deciding whether a strong equilibrium exists or not.*
3. *Computing a pure Nash equilibrium, when it is guaranteed to exist.*
4. *Computing a strong equilibrium, when it is guaranteed to exist.*
5. *Computing a strategy profile π that maximizes $\text{FLOW}(\pi)$.*
6. *For a given strategy profile π , deciding whether a given node has a best response resulting in utility at least k .*

Proof. We start by proving hardness for computing a pure Nash or a strong equilibrium when it is guaranteed to exist.

Hardness of computing Nash or strong equilibria. Consider an instance I of the problem 3-DIMENSIONAL-MATCHING. I is given by a finite set T with $|T| = 3k$ and a set $U \subseteq T \times T \times T$. [31] proved it is strongly NP-complete to decide whether there is a subset $W \subseteq U$ such that $|W| = k$ and no two elements of W have a non-empty intersection. The existence of such a set $|W|$ would be an exact cover of T . Given an instance I of 3-DIMENSIONAL-MATCHING, we construct an edge-ranking game as follows. Suppose there is a central node v . For each set $u \in U$, add a node u and connect v to u by an edge (v, u) with capacity $c((v, u)) = 3$. For each pair t, u with $t \in T$, $u \in U$ and $t \in u$ add a node t and add an edge (u, t) with capacity $c((u, t)) = 1$. Finally, connect each element node $t \in T$ by an edge (t, v) with $c((t, v)) = 1$ to v .

The idea is that computing any pure Nash or strong equilibrium reduces to finding a best response for player v . This best response, however, gives utility of $3k$ if and only if I has a solution. Hence, by computing a pure Nash or strong equilibrium, we obtain a best response for v and thereby a certificate as to whether I is solvable or not (and vice versa).

Let us first argue that there is always a strong equilibrium in this edge-ranking game. First, we note that all nodes $t \in T$ have a single outgoing edge. We fix an arbitrary feasible strategy vector π' . We will argue that best-response dynamics yield a strong equilibrium. For every strategy of v , there is at most one node u_i with $a_{u_i} \in \{1, 2\}$. For all other nodes $u \in U \setminus \{u_i\}$, every strategy π_u is a best response for u due to Proposition 3. We conclude that if π' is not a strong equilibrium, there is a coalition S of players with an improvement move such that all players in S strictly increase their utility. If $v \notin S$, there is a player $u_i \in S \cup U$. The improving move of u_i also increases the

utility of v . Thus, the utility of v increases in every step of the dynamics. This shows that the dynamics terminate with a strong equilibrium.

Let π be any Nash equilibrium in the edge-ranking game. We show the following claim. There is a subset W such that $|W| = k$ and no two elements of W have a non-empty intersection in I if and only if $u_v(\pi) = |T|$ in π .

First, let us assume that π is a Nash equilibrium with $u_v(\pi) = |T|$. Since each outgoing edge from v has capacity 3 and $|T| = 3k$, we know by the definition of edge-ranking games that $|\{w \in U \mid u_w(\pi) = 3\}| = k$. We denote these vertices by w_1, \dots, w_k . Thus, $u_{w_i}(\pi) = 3$ for $i \leq k$ and $u_{w_i}(\pi) = 0$ for all $i > k$. We will show that the sets corresponding to vertices w_1, \dots, w_k form a solution to I . Suppose there are two sets w_1, w_2 that have a non-empty intersection. If this is the case, there is an element $t \in T$ with $t \in w_1 \cap w_2$. There are edges $(w_1, t), (w_2, t)$ in the edge-ranking game that carry flow. This is a contradiction to the fact that $u_v(\pi) = |T|$.

Now, let us assume there is a solution w_1, \dots, w_k to I . We will show that every pure Nash equilibrium π in the edge-ranking game yields utility $u_v(\pi) = 3k$. The total capacity of all incoming edges of v is exactly $3k$, so $u_v(\pi) \leq 3k$ clearly holds for all strategy profiles. We will now argue that independent of the strategy choices of all other players, node v can always obtain $u_v(\pi) = 3k$. Let $\pi_v = ((v, w_1), \dots, (v, w_k), \dots)$. Since w_1, \dots, w_k exactly cover all elements $t \in T$, this induces $3k$ cycles in the TOPCYCLEINCREASE algorithm discussed in Section 3.2. This is independent of the strategy choices of all other nodes since they always have the property that either all outgoing edges are fully saturated, or there is no flow at all.

This shows that even in a class of games with guaranteed existence, computing a pure Nash equilibrium or a strong equilibrium is strongly NP-hard.

Hardness of computing social optima. Consider the previous construction. Every simple cycle in the network involves v and exactly two other nodes. There are no fixed supplies. As such, the social welfare in the system is exactly $3u_v(\pi)$. Hence, by computing a best-response for v , we also obtain a strategy profile with maximum social welfare. This proves strong NP-hardness for the computation of optimal strategy profiles. Note that this result is in contrast to unit-ranking games, where an optimal strategy profile can be computed in strongly polynomial time.

Hardness of deciding existence of Nash and strong equilibria. For hardness of the decision version, we adjust the construction from the first part of this proof and combine it with the game without pure Nash equilibrium in Fig. 3. Observe that fixed supplies of 1 for node v_7 in Fig. 3 would lead to the dominant strategy $\pi_{v_3} = ((v_3, v_1), (v_3, v_7))$ for v_3 . Given this strategy, it is straightforward to derive the utilities for all strategy choices of v_1 and v_2 .

		π_{v_2}	
		$((v_2, v_1), (v_2, v_5))$	$((v_2, v_5), (v_2, v_1))$
π_{v_1}	$((v_1, v_4), (v_1, v_6))$	9	3
	$((v_1, v_6), (v_1, v_4))$	9	5

Thus, given that v_7 has fixed supply of 1 and v_3 plays the dominant strategy $\pi_{v_3} = ((v_3, v_1), (v_3, v_7))$, the strategy profile with $\pi_{v_1} = ((v_1, v_6), (v_1, v_4))$ and $\pi_{v_2} = ((v_2, v_1), (v_2, v_5))$ is a strong equilibrium.

We use this insight to design a class of edge-ranking games based on instances of 3-DIMENSIONAL-MATCHING with the following property. If instance I has a solution, the game has a strong equilibrium (and hence a pure Nash equilibrium); if I has no solution, the game has no pure Nash equilibrium (and hence no strong equilibrium). We adjust the construction from the first part of this proof and combine it with $|T|$ copies of the example of Fig. 3: In the construction from the

first part of the proof above, delete all edges (t, v) for all $t \in T$ and add a fixed supply of $b_v = 3k$ for node v . We know that each $t \in T$ receives an inflow of 1 if and only if there is a solution to the instance I of 3-DIMENSIONAL-MATCHING. The instance depicted in Fig. 3 is copied $|T|$ times. We denote the $|T|$ copies of node v_7 by $v_7^1, v_7^2, \dots, v_7^{|T|}$. Add an edge with capacity 1 from each $t_i \in T$ to the corresponding copy of v_7 , i.e., edges (t_i, v_7^i) for $i = 1, \dots, |T|$. If there is a solution to I , all nodes $t_i \in T$ receive a flow of 1. This flow is forwarded to nodes v_7^i and can be seen as their fixed supply. Hence, there is a strong equilibrium in all copies and also in the game as a whole. On the other hand, if there is no solution to I , there is always some $t_i \in |T|$ with $u_{t_i}(\pi) = 0$, say, w.l.o.g. $u_{t_1}(\pi) = 0$. The corresponding node v_7^1 does not receive any inflow from the remaining network. Hence, there is no Nash equilibrium in the respective copy and no Nash equilibrium in the game.

Hardness of deciding the existence of a best response with revenue at least k . The instance constructed in the proof of Theorem 5 is an edge-ranking game and the proof immediately carries over. \square

REMARK 3. It is unclear whether the problem of deciding existence of a pure Nash equilibrium in an edge-ranking game is in NP or not, due to NP-hardness of verification that a node plays a best response (see Theorem 5). It is easy to see that the decision problem is in Σ_2^P . The problem could be Σ_2^P -complete, similar to related decision problems in strategic max-flow games [33, 21]. Proving such a result is an interesting open problem.

5.2. Total flow in equilibrium. For edge-ranking games, the lower bound on the price of anarchy for Nash equilibria observed in Proposition 6 applies, i.e., the price of anarchy for Nash equilibria in terms of FLOW can be unbounded. The restriction to edge-ranking strategies can have a drastic effect even on the total flow in the best equilibrium in case it exists. In particular, in edge-ranking games the price of stability for strong equilibria in terms of FLOW can be as high as $\Omega(n)$, and the price of stability for Nash equilibria might even be unbounded.

PROPOSITION 10. *For every $\varepsilon > 0$, there is an edge-ranking game with price of stability for strong equilibria in terms of FLOW of at least $n/2 - \varepsilon$.*

Proof. We construct an edge-ranking game that consists of a single cycle plus one additional edge. More formally, we have firms $V = \{v_1, \dots, v_n\}$ and edges $E = \left\{ \{(v_i, v_{i+1}) \mid i \in \{1, \dots, n-1\}\} \cup (v_n, v_1) \cup (v_1, v_n) \right\}$. The edges $(v_1, v_n), (v_n, v_1), (v_1, v_2)$ have a capacity $M + 1$ and all other edges a capacity of M . The only node with more than a single outgoing edge is v_1 . If $\pi_{v_1} = ((v_1, v_n), (v_1, v_2))$, player v_1 gets a total supply of $M + 1$, which is optimal. Observe that $\text{FLOW}(\pi) = 2M + 2$, and that π is the only Nash equilibrium and the only strong equilibrium.

In contrast, for profile π' with $\pi'_{v_1} = ((v_1, v_2), (v_1, v_n))$, firm v_1 only gets a utility of M , but $\text{FLOW}(\pi') = nM$. Thus, the strong price of stability is at least $nM/(2M + 2) = n/2 - n/(2M + 2)$, which is at least $n/2 - \varepsilon$ for $M \geq n/(2\varepsilon) - 1$. \square

PROPOSITION 11. *The price of stability for Nash equilibria in terms of FLOW is unbounded in edge-ranking games.*

Proof. Consider the game in Fig. 4, which extends the game without pure equilibrium from Fig. 3. We add three nodes. w_1 has fixed supply 1, w_2 and w_3 no fixed supply. These nodes are involved in a cycle C of edges (w_1, w_2) and (w_2, w_3) with capacity $M \gg 2$, as well as edge (w_3, w_1) with capacity $M - 2$. In addition, there are edges (w_1, v_9) and (w_2, v_9) of capacity 2.

In an optimal circulation, w_1 and w_2 prioritize the edges of C , leading to utilities of $\Theta(M)$. In contrast, a pure Nash equilibrium can only exist if the w -nodes ensure that the fixed supply of w_1 is forwarded to v_7 , in which case a Nash equilibrium can exist (as observed in the proof of Theorem 7). Clearly, both w_1 and w_2 have an incentive to deviate towards C . Hence, if either w_1




FIGURE 4. An edge-ranking game with unbounded price of stability for Nash equilibria.

or w_2 places the edge to v_7 in first rank and the other does not, a unilateral deviation suffices to close C – thereby leaving the v -nodes with instability. However, if both w_1 and w_2 play strategies $\pi_{w_1} = ((w_1, v_7), (w_1, w_2))$ and $\pi_{w_2} = ((w_2, v_7), (w_2, w_3))$, no unilateral deviation can lead to flow along C . In this case, a pure Nash equilibrium evolves. Obviously the total flow in this equilibrium is at most a constant. Hence, the price of stability is as large as $\Omega(M)$. \square

6. Conclusions. In this paper, we have proposed and analyzed flow allocation games. Our main results show that in these games, if firms are following priority rankings over units of flow (i.e., unit-ranking strategies), there is always a strong equilibrium. Moreover, it can be computed in strongly polynomial time and represented in polynomial space. More generally, for a large class of *flow-monotone* quality functions Q , there is even an *optimal* strong equilibrium, i.e., the price of stability for Nash and strong equilibria in terms of Q is 1. In terms of computational complexity, the properties of such optimal strong equilibria depend highly on the quality function Q . While for some objectives such as FLOW the optimization problem of finding an optimal strong equilibrium can be solved in polynomial time, for other objectives such as FSN the optimization problem can become strongly NP-hard. Alternatively, when restricting the strategy spaces to priorities over edges (i.e., edge-ranking strategies), pure Nash and strong equilibria can be absent, and even deciding their existence is a NP-hard problem.

As a concrete example of a quality function, our results shed further light on the performance of equilibria in terms of FLOW. When considering decentralized clearing and arbitrary strong equilibria in unit-ranking games, the price of anarchy for strong equilibria depends on the length of cycles in the money circulation of an optimum profile. For pure Nash equilibria, the deterioration in FLOW can be severe due to the lack of coordination among firms. For edge-ranking strategies, pure Nash and strong equilibria can have very poor quality in terms of FLOW when they exist.

Our work provides a game-theoretic perspective on clearing in financial networks. In the context of financial networks, unit-ranking functions allow a centralized market regulator to obtain a clearing state, in which a desired social quality is maximized and no coalition of firms gets an incentive to deviate. As such, our work reveals an interesting alignment of incentives – firms (even *groups* of firms) share an intrinsic interest to implement these proposed clearing payments. This *strategic robustness* represents an elegant game-theoretic complement to standard *legal enforcement* by regulators in financial networks.

More generally, while we have initiated the analysis of flow allocation games, numerous recent follow-up works in the context of financial networks (c.f. our discussion in Section 1.2) show that our work has sparked interest in a variety of directions. Understanding these issues continues to be of vital importance to improve the financial system and to inform the discussion about financial regulation from a computational and game-theoretic perspective.

Acknowledgments. The authors thank Pascal Lenzner and Steffen Schuldenzucker for valuable discussions and feedback on the results of this paper. The authors acknowledge financial support by DFG Research Group ADYN under grant DFG 411362735. NB thanks Dr. h. c. Maucher for funding his position.

References

- [1] Arora S, Barak B, Brunnermeier M, Ge R (2011) Computational complexity and information asymmetry in financial products. *Comm. ACM* 54(5):101–107.
- [2] Bachrach Y, Rosenschein J (2009) Power in threshold network flow games. *Auton. Agents and Multi-Agent Syst.* 18(1):106–132.
- [3] Barucca P, Bardoscia M, Caccioli F, D’Errico M, Visentin G, Caldarelli G, Battiston S (2020) Network valuation in financial systems. *Math. Finance* 30(4):1181–1204.
- [4] Bertschinger N, Hoefler M, Schmand D (2020) Strategic payments in financial networks. *Proc. 11th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, 46:1–46:16.
- [5] Braverman M, Pasricha K (2014) The computational hardness of pricing compound options. *Proc. 5th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, 103–104.
- [6] Candogan O, Epiropou M, Vohra R (2021) Competitive equilibrium and trading networks: A network flow approach. *Oper. Res.* 69(1):114–147.
- [7] Cseh Á, Matuschke J (2019) New and simple algorithms for stable flow problems. *Algorithmica* 81(6):2557–2591.
- [8] Cseh A, Matuschke J, Skutella M (2013) Stable flows over time. *Algorithms* 6(3):532–545.
- [9] Csóka P, Herings PJ (2018) Decentralized clearing in financial networks. *Manag. Sci.* 64(10):4681–4699.
- [10] Deng X, Ibaraki T, Nagamochi H (1999) Algorithmic aspects of the core of combinatorial optimization games. *Math. Oper. Res.* 24(3):751–766.
- [11] Dinitz Y (1970) Algorithm for solution of a problem of maximum flow in networks with power estimation. *Soviet Math. Dokl.* 11:1277–1280.
- [12] Dubey P, Shapley L (1984) Totally balanced games arising from controlled programming problems. *Math. Prog.* 29(3):245–267.
- [13] Edmonds J, Karp RM (1972) Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM* 19(2):248–264.
- [14] Eisenberg L, Noe T (2001) Systemic risk in financial systems. *Manag. Sci.* 47(2):236–249.
- [15] Fischer T (2014) No-arbitrage pricing under systemic risk: Accounting for cross-ownership. *Math. Finance* 24(1):97–124.
- [16] Fleiner T (2014) On stable matchings and flows. *Algorithms* 7(1):1–14.
- [17] Fleiner T, Jankó Z, Schlotter I, Teytelboym A (2023) Complexity of stability in trading networks. *Int. J. Game Theory* 52(3):629–648.
- [18] Froese H, Hoefler M, Wilhelmi L (2023) The complexity of debt swapping. *CoRR* abs/2302.11250.
- [19] Gai P, Kapadia S (2010) Contagion in financial networks. *Proc. Royal Soc. London A: Math. Phys. Eng. Sci.* 466(2120):2401–2423.
- [20] Granot D, Granot F (1992) On some network flow games. *Math. Oper. Res.* 17(4):792–841.
- [21] Guha S, Kupferman O, Vardi G (2019) Multi-player flow games. *Auton. Agents Multi Agent Syst.* 33(6):798–820.
- [22] Hatfield J, Kominers S, Nichifor A, Ostrovsky M, Westkamp A (2013) Stability and competitive equilibrium in trading networks. *J. Political Econ.* 121(5):966–1005.
- [23] Hemenway B, Khanna S (2016) Sensitivity and computational complexity in financial networks. *Algorithmic Finance* 5(3-4):95–110.
- [24] Hoefler M, Wilhelmi L (2022) Seniorities and minimal clearing in financial network games. *Proc. 15th Symp. Algorithmic Game Theory (SAGT)*, 187–204.

- [25] Ioannidis S, de Keijzer B, Ventre C (2022) Strong approximations and irrationality in financial networks with financial derivatives. *Proc. 49th Int. Colloq. Autom. Lang. Programming (ICALP)*, 76:1–76:18.
- [26] Ioannidis S, de Keijzer B, Ventre C (2023) Financial networks with singleton liability priorities. *Theor. Comput. Sci.* 963:113965.
- [27] Kalai E, Zemel E (1982) Generalized network problems yielding totally balanced games. *Oper. Res.* 30(5):998–1008.
- [28] Kalai E, Zemel E (1982) Totally balanced games and games of flow. *Math. Oper. Res.* 7(3):476–478.
- [29] Kanellopoulos P, Kyropoulou M, Zhou H (2021) Financial network games. *Proc. 2nd ACM Int. Conf. AI in Finance (ICAIF)*, 26:1–26:9.
- [30] Kanellopoulos P, Kyropoulou M, Zhou H (2022) Forgiving debt in financial network games. *Proc. 31st Int. Joint Conf. Artif. Intell. (IJCAI)*, 335–341.
- [31] Karp R (1972) Reducibility among combinatorial problems. Miller R, Thatcher J, eds., *Complexity of Computer Computations* (Plenum Press, New York).
- [32] Király T, Pap J (2013) Stable multicommodity flows. *Algorithms* 6(1):161–168.
- [33] Kupferman O, Vardi G, Vardi M (2017) Flow games. *Proc. 37th Conf. Found. Software Tech. Theor. Comput. Sci. (FSTTCS)*, 38:1–38:16.
- [34] Markakis E, Saberi A (2005) On the core of the multicommodity flow game. *Decis. Support Syst.* 39(1):3–10.
- [35] Ostrovsky M (2008) Stability in supply chain networks. *Amer. Econ. Rev.* 98(3):897–923.
- [36] Papadimitriou C (2001) Algorithms, games and the internet. *Proc. 33rd Symp. Theory Comput. (STOC)*, 749–753.
- [37] Papp PA, Wattenhofer R (2020) Network-aware strategies in financial systems. *Proc. 47th Int. Colloq. Autom. Lang. Programming (ICALP)*, 91:1–91:17.
- [38] Papp PA, Wattenhofer R (2021) Debt swapping for risk mitigation in financial networks. *Proc. 22nd Conf. Econ. Comput. (EC)*, 765–784.
- [39] Papp PA, Wattenhofer R (2021) Default ambiguity: Finding the best solution to the clearing problem. *Proc. 17th Conf. Web and Internet Econ. (WINE)*, 391–409.
- [40] Papp PA, Wattenhofer R (2021) Sequential defaulting in financial networks. *Proc. 12th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, 52:1–52:20.
- [41] Rogers LCG, Veraart LAM (2013) Failure and rescue in an interbank network. *Manag. Sci.* 59(4):882–898.
- [42] Schuldenzucker S, Seuken S, Battiston S (2017) Finding clearing payments in financial networks with credit default swaps is PPAD-complete. *Proc. 8th Symp. Innov. Theoret. Comput. Sci. (ITCS)*, 32:1–32:20.
- [43] Shapley L, Scarf H (1974) On cores and indivisibility. *J. Math. Econ.* 1(1):23–37.
- [44] Suzuki T (2002) Valuing corporate debt: The effect of cross-holdings of stock and debt. *J. Oper. Res. Soc. Japan* 2.
- [45] Tardos É (1985) A strongly polynomial minimum cost circulation algorithm. *Combinatorica* 5(3):247–256.
- [46] Tarski A (1955) A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* 5(2):285–309.