Equilibria and Convergence in Fire Sale Games

Nils Bertschinger
Frankfurt Inst. for Advanced Studies
Frankfurt, Germany
bertschinger@fias.uni-frankfurt.de

Martin Hoefer
Goethe University
Frankfurt, Germany
mhoefer@em.uni-frankfurt.de

Simon Krogmann
Hasso Plattner Institute
Potsdam, Germany
simon.krogmann@hpi.de

Pascal Lenzner
Hasso Plattner Institute
Potsdam, Germany
pascal.lenzner@hpi.de

Steffen Schuldenzucker
Goethe University
Frankfurt, Germany
schuldenzucker@em.uni-frankfurt.de

Lisa Wilhelmi
Goethe University
Frankfurt, Germany
wilhelmi@em.uni-frankfurt.de

ABSTRACT

The complex interactions between algorithmic trading agents can have a severe influence on the functioning of our economy, as witnessed by recent banking crises and trading anomalies. A common phenomenon in these situations are *fire sales*, a contagious process of asset sales that trigger further sales. We study the existence and structure of equilibria in a game-theoretic model of fire sales. We prove that for a wide parameter range (e.g., convex price impact functions), equilibria exist and form a complete lattice. This is contrasted with a non-existence result for concave price impact functions. Moreover, we study the convergence of best-response dynamics towards equilibria when they exist. In general, best-response dynamics may cycle. However, in many settings they are guaranteed to converge to the socially optimal equilibrium when starting from a natural initial state. Moreover, we discuss a simplified variant of the dynamics that is less informationally demanding and converges to the same equilibria. We compare the dynamics in terms of convergence speed.

KEYWORDS

Equilibrium Existence; Best-Response Dynamics; Convergence; Financial Networks

ACM Reference Format:
Nils Bertschinger, Martin Hoefer, Simon Krogmann, Pascal Lenzner, Steffen Schuldenzucker, and Lisa Wilhelmi. 2023. Equilibria and Convergence in Fire Sale Games. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

On May 6, 2010, 2:45pm, one trillion dollars in stock market valuation disappeared. In an event known as a *flash crash*, the Dow Jones and many other stock indices collapsed by as much as 9%. The flash crash is generally seen as the product of a system of interacting agents, many of them computerized, that jointly exacerbated an initial shock.[1] While prices recovered after approximately 30 minutes (after trading was briefly halted), it is by no means guaranteed that this will always be the case in a flash crash. Recent price crashes in cryptocurrency markets, such as the Bitcoin crash on April 17, 2021 [33], may have resulted in permanently altered market conditions.

As prices deteriorate very quickly, it is important to understand the amplification of stock price declines through the interaction between different (electronic or human) trading agents. We study this aspect through the lens of algorithmic game theory.

Agents may amplify price decreases through many different processes. In this work, we focus on *price-mediated contagion due to leverage caps*. We study a collection of rational agents that interact with each other through overlapping portfolios. Each agent holds a share of the supply of one of several assets, where usually two or more agents hold a share in the same asset. Sales in any given asset depress its price, i.e., sales have *price impact*, which in turn may reduce the value of the asset holdings of another agent. We assume that agents are constrained by an upper bound on their *leverage*, i.e., the ratio between their (risky) asset holdings and their equity. The equity of an agent is the difference between the total value of her assets and her liabilities, and it includes cash proceeds from the liquidation of risky assets. Leverage caps can represent the agents' own desire to limit their risk, or they might be regulatory constraints.[2] Because of leverage caps, a price depression in some of the assets may cause agents to perform further sales to satisfy their leverage constraints. This can give rise to a contagious *fire sale* process, where a small initial price drop quickly leads to a large number of asset sales and corresponding price drop.

1.1 Related Work

Fire sales are a well-known phenomenon and have been studied both academically and by regulators (e.g., central banks). A *leverage cycle*, in which banks' reduction of leverage leads to price decline and further reduction of leverage, was first studied by Geanakoplos [19]. Aymanns and Farmer [2] and Aymanns et al. [1] studied a dynamic model of leveraged and unleveraged investors. Cont and Schaanning [11] studied a model of overlapping portfolios and leverage constraints. Agents react to an initial shock following a specific iterative liquidation process, where they proportionally sell

[1]The Joint CFTC-SEC Advisory Committee on Emerging Regulatory Issues [24] stated that the flash crash was triggered by a single market participant employing an (arguably simplistic) trading algorithm. Other trading agents at first absorbed the resulting price pressure, but then amplified and spread it. Other examples for flash crashes are extreme movements in currency markets in recent years [5, 10]. Anecdotal evidence suggests that smaller-scale flash crashes happen very frequently [10, 18].

[2]The Basel III regulatory framework stipulates leverage caps for banks [6].

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

off a fraction of their portfolio. Our paper re-interprets a generalization of their model in a game-theoretic setting. Importantly, Cont and Schaanning [11] did not study game-theoretic equilibria. In a later work [12], the same authors discuss risk indicators that help quantify the exposure of a given institution to price-mediated contagion. Baes and Schaanning [3] studied worst-case scenarios when agents respond to price depreciations in an individually optimal way. Banerjee and Feinstein [4] studied a fire sale process where liquidations occur at volume-weighted average prices. Price-mediated contagion has also received interest from regulators, e.g., in the European Central Bank's STAMP€ macro stress testing framework [13, Section 12.2.1].

Another related model for studying contagion effects in financial networks was introduced by Elliott et al. [16]. This model considers banks and assets and allows for banks owning shares of other banks. Banks in the network are connected by linear dependencies, i.e., cross-holdings. If any banks's value drops below a critical threshold, its value suffers an additional failure cost, potentially impacting the value of other banks. Hemenway and Khanna [21] study the sensitivity and computational complexity of this model.

Recently, there has been considerable interest in analyzing financial networks from an algorithmic and game-theoretic point of view. A number of works are based on a classical model for systemic risk in financial networks by Eisenberg and Noe [15] which studies the clearing problem, i.e., to determine which banks are in default and their exposure to systemic risk. A recent work by Bertschinger et al. [9] proposes a strategic version, in which firms are rational agents in a given directed graph of debt contracts. To clear its debt, every agent strategically decides on a ranking-based payment strategy. The authors study the existence and computational complexity of pure Nash and strong equilibria, and provide bounds on the (strong) prices of anarchy and stability. Kanellopoulos et al. [25] consider the same model but they focus on more general priority-proportional payments. Recently, Hoefer and Wilhelmi [22] analyzed minimal clearing and the impact of seniorities (i.e., priorities over debt contracts) on the existence of equilibria. Moreover, [26] study complexity questions regarding centralized bailouts and the forgiving of debts for banks in default, along with game-theoretic incentives that emerge in such scenarios.

The network model by Eisenberg and Noe [15] has also been augmented by considering credit-default swaps (CDS) [34]. The hardness of finding clearing payments with CDS is analyzed by Schuldenzucker et al. [35] and the approximation perspective was considered recently by Ioannidis et al. [23]. The impact of banks deleting or adding liabilities, donating to other banks, or changing external assets has been studied by [29]. These operations can be beneficial for the individual banks since the changes may enforce more favorable solutions. Still, after clearing some banks may end up in default. Papp and Wattenhofer [32] study the influence of the sequence of banks' defaults. In [31] several possible strategies for resolving default ambiguity, i.e., which banks end up in default and how much of their liabilities can these defaulting banks pay, are studied. In a different direction, in [30] risk mitigation via local network changes ("debt swapping") is investigated.

Frequent call markets represent another financial game which has sparked research recently [27, 28]. Here, the focus lies on preventing fraud while maintaining efficiency.

1.2 Our Contribution

In this paper, we innovate upon prior work by studying fire sales as a static *fire sale game* played by fully rational agents. Each agent decides on the fraction of her portfolio she sells and keeps, respectively.[3] The agent derives a utility equal to her equity, i.e., the shareholder value of her firm, as long as she satisfies the leverage constraint.[4] The equity depends on her own action (through sales choices and price impact) and the actions of others (through price impact). We study the Nash equilibria of this game under different restrictions on the strategy space, assumptions on the price impact function, and assumptions on how the price impact manifests while selling any given asset. We capture this latter dimension in a parameter $\alpha \in [0, 1]$.

In this paper, we discuss the case $\alpha = 1$, where price impact affects prices of assets sold as much as those of assets kept (i.e., *post-sale prices*). This yields the strongest price impact on the utility of the agents, so the scenario captures best the effects we intend to study. Conveniently, $\alpha = 1$ also simplifies some of our calculations.

We study the existence and the structure of equilibria. When $\alpha = 1$, we show that an equilibrium always exists and the set of equilibria further has a desirable *lattice structure*. In this case, a process of iterative best responses converges from above to the point-wise maximal equilibrium, which sets each agent up to be best off among all equilibria (Section 3).

Also, we study processes by which agents may actually converge to an equilibrium, where we focus on the case of $\alpha = 1$ and linear price impact. We consider two kinds of best-response dynamics: the regular best-response dynamics and a simplified best-response dynamics, where agents do not take the fact into account that their own response to changing market prices in turn generates price impact. One may argue that the simplified dynamics serves as a more realistic model of agent behavior since it is less complex to execute and requires less information about the precise shape of the price impact function. We show that, under the assumptions of our lattice structure result, both of these dynamics converge to the maximal equilibrium. However, we show via computational experiments that they can do so at vastly different speeds. Our experiments further suggest that the convergence speed depends on the overall diversification of agents across assets, with the worst case depending on the specific parameters (Section 4).

We provide further results in the full version [7]. For intermediate prices $\alpha \in (0, 1)$ equilibrium existence can be guaranteed only if the price impact function is convex. Then the equilibria have a lattice structure and best-response dynamics converges from above to the point-wise maximal equilibrium. In contrast, if $\alpha < 1$ and price impact is concave, an equilibrium need not even exist. Furthermore, we show that for $\alpha = 1$, the Pareto optima of the game form strong equilibria. This is not true for $\alpha < 1$ and convex price impact. In the latter case "bank-run" effects can manifest, where an agent is incentivized to sell more than strictly necessary to satisfy her

[3]The relative composition of each agent's portfolio is kept constant in all of this paper (we study a generalization in [7]). This is a standard assumption in the literature on price-mediated contagion [see, e.g., 11, 14] and supported by empirical evidence [e.g., 20]. Our main reason to discuss this case is that it helps to simplify the presentation. Most of the results in this paper can be extended in a straightforward way to general monotone sales, see our discussion in Section 5.

[4]We assume that agents that cannot satisfy the leverage constraint need to fully liquidate their position.

leverage constraint. Moreover, even under severe restrictions, there are games where a Nash equilibrium yields devastating utility for every agent in comparison to the social optimum.

2 PRELIMINARIES

2.1 The Model

Our market model is based on Cont and Schaanning [11]. We have sets of *agents* $N = \{1, \ldots, n\}$ and *assets* $M = \{1, \ldots, m\}$. Each agent i holds an amount $a_i^I > 0$ of *illiquid assets*, which cannot be sold and are not subject to price impact. Agent i holds an amount of x_{ij} in the (liquid) asset j. We assume w.l.o.g. the amounts to be normalized, so that $x_{ij} \in [0,1]$ for all i, j, if no agent sells any holdings in j.

Each agent i's strategic action consists of a number $y_i \in [0,1]$, which is the share of agent i's holdings in each asset j that agent i *keeps*. Agent i thus sells an amount of $1 - y_i$ of her total portfolio holdings on the market. The amount of asset j held by agent i after selling is $y_i x_{ij}$. Let $x_j(y) := \sum_{i \in N} y_i x_{ij}$ be the amount of asset j that has not been sold if agents act according to y. We assume that the price of each asset j decays when assets are sold according to a function $p_j(y) = p_j(x_j(y))$ such that $p_j(0) = 0$, $p_j(1) = p_j^0$, and $p_j(x_j)$ is continuous and increasing in x_j. We say that price impact is *linear* if for each $j \in M$, the function p_j simply linearly interpolates between the two given points, i.e., $p_j(x_j) = x_j p_j^0$. We say that price impact is *convex* if p_j is convex for each j and *concave* if p_j is concave for each j. If all agents sell according to y, the value of agent i's remaining asset holdings is therefore

$$a_i(y) = a_i^I + y_i \sum_{j \in M} x_{ij} p_j(y).$$

Selling assets redeems a certain amount of money for each agent i. As the agents sell their assets, the corresponding price impact would usually manifest over time: amounts of assets that are sold at the very beginning would typically not be subject to price impact, while sales that happen late would bear significant price impact. Since we consider a static game, we do not model this effect directly. Instead, we follow Cont and Schaanning [11] by capturing the effect using an *implementation shortfall parameter* $\alpha \in [0,1]$. Agents redeem a share of α of their sales according to the post-price-impact price and a share of $1 - \alpha$ according to the pre-price-impact price of an asset. If $\alpha = 0$, the market reacts very slowly to asset sales, so that price impact does not manifest in the effective price that agents receive when they sell their assets (but it does manifest for the post-sale values of the remaining assets). If $\alpha = 1$, then the full price impact manifests immediately; such a situation may arise when asset sales are conducted using an auction mechanism.[5] If all agents sell according to y, the total revenue that agent i derives from her asset sales is now

$$\Delta_i(y) = (1 - y_i) \sum_{j \in M} x_{ij}((1 - \alpha)p_j^0 + \alpha p_j(y)).$$

[5] An alternative approach is to compute the integral of of the price impact function with respect to a path of infinitesimal trades. This route was chosen by Banerjee and Feinstein [4]. For the purpose of this paper, where we are interested in the strategic implications of fire sales, our approach using the α parameter provides a simple way of capturing the extent to which price impact affects agents.

We assume that each agent i has *liabilities* of l_i to external creditors. Agent i's *equity* is the difference between her total assets and liabilities, where her total assets consist of her (illiquid and liquid) assets and the risk-free money she has redeemed from asset sales:

$$e_i(y) = a_i(y) + \Delta_i(y) - l_i.$$

Note that the equity of an agent is what would remain if the agent's (risky and risk-free) assets were used to pay off her liabilities. It therefore equals the agent's shareholder value.

If $e_i(y) > 0$, then agent i's *leverage* at y is the ratio between her risky assets and her equity:

$$\mathrm{lev}_i(y) = \frac{a_i(y)}{e_i(y)} = \frac{a_i(y)}{a_i(y) + \Delta_i(y) - l_i}.$$

Note how an agent that holds no risky assets (i.e., $a_i(y) = 0$) has a leverage of 0 (unless its equity is also 0) while an agent that holds high risky assets, only little risk-free assets and has high liabilities (i.e., $a_i(y)$ and l_i are large and $\Delta_i(y)$ is small) has a high leverage. This is why leverage is used as an instrument to gauge the riskiness of an agent. If $e_i(y) \le 0$, then $\mathrm{lev}_i(y)$ is not defined.

We assume that regulatory constraints limit the admissible leverage of an agent to a constant $\lambda > 1$, i.e., agent i needs to choose its action y_i such that

$$\mathrm{lev}_i(y) \le \lambda.$$

If no such y_i exists, we say that agent i is *illiquid* given the actions y_{-i} of the other agents. If no y_i exists for which $e_i(y) > 0$, we say that agent i is *insolvent* at y_{-i}. Insolvent or illiquid agents need to sell their whole asset holdings (otherwise, we assume that they receive utility $-\infty$); the other agents (which we call *liquid* agents) derive a utility equal to their equity. More in detail, we consider the following utility function. Define special strategies $y_i^0 := 0$ and $y_i^1 := 1$ and define the *utility* of agent i as

$$u_i(y) := \begin{cases} -\infty & \text{if } y_i \ne y_i^0 \text{ and} \\ & \quad (e_i(y) \le 0 \text{ or } \mathrm{lev}_i(y) > \lambda) \\ e_i(y) & \text{otherwise.} \end{cases}$$

Note that agents always have the option to sell everything (i.e., play $y_i = y_i^0$) and then receive a utility equal to their equity. This is motivated by the fact that it should always be possible to liquidate a firm; regulation must not prevent agents from exiting the market. Observe that, in any Nash equilibrium, insolvent or illiquid agents sell everything and liquid agents either play a y_i for which $e_i(y) > 0$ and $\mathrm{lev}_i(y) \le \lambda$ or sell everything (i.e., play $y_i = y_i^0$). We call a collection $(N, M, a^I, x, \alpha, \lambda)$ a *fire sale game*. See Figure 1 for an example instance.

2.2 Basic Properties

Our first proposition shows that sales of one agent destabilize other agents with overlapping portfolios, in the sense that their leverage increases. This effect introduces fire sales into our model. For technical reasons, we need to consider a lower bound of 1 on the leverage; note that this is irrelevant for our discussion since the leverage cap is always $\lambda > 1$.

PROPOSITION 2.1. *Let* $\overline{\mathrm{lev}}_i(y) = \max(1, \mathrm{lev}_i(y))$. *Then* $\overline{\mathrm{lev}}_i(y_i, y_{-i})$ *is monotonically decreasing in* y_{-i}. *More in detail, if* $y_i \in [0,1]$ *and* $y_{-i} \le y'_{-i}$ *point-wise, then* $\overline{\mathrm{lev}}_i(y_i, y_{-i}) \ge \overline{\mathrm{lev}}_i(y_i, y'_{-i})$.

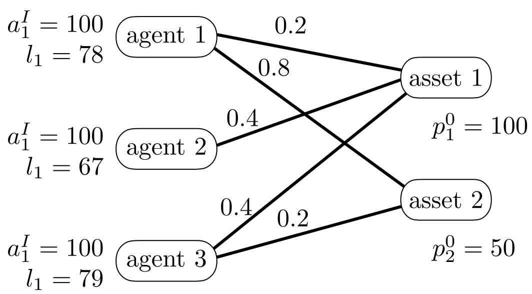

y_1	y_2	y_3	$\text{lev}_1(y)$	$u_1(y)$	$\text{lev}_2(y)$	$u_2(y)$	$\text{lev}_3(y)$	$u_3(y)$
1.0	1.0	1.0	1.95	82.0	1.92	73.0	2.11	$-\infty$
1.0	1.0	**0.7**	2.01	$-\infty$	1.98	68.18	2.0	65.58
0.95	1.0	0.7	2.0	75.51	1.99	67.81	2.01	$-\infty$
0.95	1.0	**0.65**	2.01	$-\infty$	1.99	67.02	2.0	63.94
0.89	1.0	0.65	2.0	72.36	2.01	$-\infty$	2.02	$-\infty$
0.89	**0.97**	0.65	2.01	$-\infty$	2.0	66.02	2.03	$-\infty$
0.89	0.97	**0.51**	2.04	$-\infty$	2.04	$-\infty$	2.0	59.88
0.0	0.97	0.51	2.64	$e_1(y)$	2.17	$-\infty$	2.46	$-\infty$
0.0	**0.0**	0.51	3.32	$e_1(y)$	2.43	$e_2(y)$	3.47	$-\infty$
0.0	0.0	**0.0**	4.54	$e_1(y)$	3.03	$e_2(y)$	4.76	$e_3(y)$

Figure 1: A fire sale game with three agents and two assets, $\alpha = 1$, $\lambda = 2.0$, and linear price impact, i.e., $p_j(y) = p_j^0 \cdot \sum_{i \in N} x_{ij} y_i$. Left: Asset holdings and initial prices. Right: Sequence of game states obtained by playing best responses starting from the state $(1, 1, 1)$. Newly chosen strategies are red, all values are rounded to two digits. When agent 3 starts selling to fulfill her leverage constraint, a fire sale starts that eventually forces all agents to sell everything.

Proof. By assumption and monotonicity of the functions p_j, we have $p_j(y_i, y_{-i}) \leq p_j(y_i, y'_{-i})$ for all j.

It immediately follows from the definition that $\text{lev}_i(y_i, y_{-i}) < 1$ if and only if $\Delta_i(y_i, y_{-i}) > l_i$. If this is the case, then also $\Delta_i(y_i, y'_{-i}) \geq \Delta_i(y_i, y_{-i}) > l_i$, where the first inequality immediately follows from the above statement about the prices p_j. Now also $\text{lev}_i(y_i, y'_{-i}) < 1$ and thus $\overline{\text{lev}}_i(y_i, y_{-i}) = 1 = \overline{\text{lev}}_i(y_i, y'_{-i})$.

Assume now that $\text{lev}_i(y_i, y_{-i})$, $\text{lev}_i(y_i, y'_{-i}) \geq 1$ and write short $a := a_i(y_i, y_{-i})$, $a' := a_i(y_i, y'_{-i})$, and likewise for Δ. Then

$$\overline{\text{lev}}_i(y_i, y_{-i}) = \text{lev}_i(y_i, y_{-i}) = \frac{a}{a + \Delta - l_i} \geq \frac{a'}{a' + \Delta - l_i}$$
$$\geq \frac{a'}{a' + \Delta' - l_i} = \text{lev}_i(y_i, y'_{-i}) = \overline{\text{lev}}_i(y_i, y'_{-i}),$$

where the first inequality holds since, by assumption, $\Delta - l_i \leq 0$ and $a \leq a'$, and the second inequality holds since $\Delta \leq \Delta'$, both of which immediately follow from monotonicity of the prices p_j. $\square$

Note that lev_i is *not* necessarily monotonically increasing in y_i, i.e., selling more does *not always* reduce leverage. Whether or not this is the case depends on the price impact functions, the α parameter, and y_{-i}.

The leverage function lev_i is also continuous in y, which follows directly from the definitions.

Proposition 2.2. *The leverage function* lev_i *is continuous in the strategy profile* y *in the region where agent* i *is solvent.*

Agents become illiquid before they become insolvent. This technical property will be useful in our proofs. It follows immediately from the definition $\text{lev}_i(y) = a_i(y)/e_i(y)$, the assumption $e_i(y^t) \to 0$, and the fact that $a_i(y^t) \geq a_i^I > 0$ for all t.

Proposition 2.3. *Let* i *be an agent and let* $(y^t)_t$ *be a sequence of strategy profiles such that* $e_i(y^t) > 0$ *for all* t *and* $e_i(y^t) \to 0$. *Then* $\text{lev}_i(y^t) \to \infty$. *In particular, there exists a* t *such that* $\text{lev}_i(y^t) > \lambda$.

2.3 Post-Sale Prices

For the results in the main body of the paper, we focus on the important special case when $\alpha = 1$. Here, agents liquidate their assets at an average price that is equal to the price after all assets have been sold. We also say that in this case the agents receive *post-sale prices*. Agents are highly affected by price devaluations in this case. Intuitively, agents will therefore avoid sales and only execute them to satisfy the leverage constraint. The following Proposition 2.4 formalizes this intuition. We will show that this further implies monotonicity of the best response of each agent, which will be an important building block towards our later results.

The best-response function $\Phi : [0,1]^N \to [0,1]^N$ is

$$\Phi_i(y) := \arg\max_{y_i \in [0,1]} u_i(y_i, y_{-i}).$$

Ties in the arg max are w.l.o.g. broken in favor of largest y_i.

It is easy to see that for $\alpha = 1$, the equity e_i simplifies to

$$e_i(y) = a_i^I - l_i + \sum_j x_{ij} p_j(y). \tag{1}$$

The equity is equal to the assets, net of liabilities, assuming that the agent has not actually sold anything, but is still exposed to price impact on her asset holdings. This drives the following result.

Proposition 2.4. *Let* $\alpha = 1$. *Then (1) each liquid agent* i *maximizes her utility by maximizing* y_i *subject to the leverage constraint and (2)* Φ *is monotonic.*

Proof. First, observe that the equity is monotonic in y_i since the functions p_j are monotonic and by (1). The best response of agent i thus minimizes sales (i.e., maximizes y_i) subject to the leverage constraint.

Now consider the best response y_i^* of i to a fixed y_{-i}. First, assume the special case where $y_i^* = 0$. Then, by monotonicity of the equity, agent i either exactly satisfies the leverage constraint or is illiquid or insolvent. In the latter scenario, the leverage is undefined, while the other cases imply $\text{lev}_i(y_i^*, y_{-i}) \geq \lambda > 1$. By monotonicity of the equity and the leverage (Proposition 2.1) in y_{-i}, the agent must maintain this strategy when the others increase their sales. Now assume that $y_i^* > 0$. To maximize utility, the agent chooses the highest value y_i as her strategy so that the leverage constraint is still satisfied. Whenever the agent fulfills the leverage constraint without liquidating any assets (i.e., $y_i^* = 1$), her leverage is at least 1, since $\text{lev}_i(y_i^*, y_i) = a_i(y_i^*, y_{-i})/(a_i(y_i^*, y_{-i}) - l_i) \geq 1$. Now assume the agent to sell a share of her assets, i.e., $y_i^* \in (0, 1)$. In particular, this means that $\text{lev}_i(y_i^*, y_{-i}) = \lambda > 1$ and $\text{lev}_i(\hat{y}_i, y_{-i}) > \lambda$ for all

$\hat{y}_i > y_i^*$. If the other agents now increase their sales, all $\hat{y}_i$ remain invalid (i.e., $y'_{-i} \leq y_{-i} \Rightarrow \text{lev}_i(\hat{y}_i, y'_{-i}) > \lambda$ for all $\hat{y}_i > y_i^*$, due to Proposition 2.1). Therefore, if y_i^* no longer satisfies the constraint, i must also liquidate more assets. □

3 EQUILIBRIUM EXISTENCE AND CONVERGENCE OF DYNAMICS

In this section, we discuss the existence of equilibria in a given fire sale game and the convergence of best-response dynamics to equilibrium. Our first main result is that, in the above-discussed cases where the best-response function is monotonic, the set of equilibria has a particularly desirable *lattice structure*. This in particular implies that (1) an equilibrium always exists and (2) there is always an equilibrium that minimizes the sales of each asset by each agent simultaneously and maximizes the equity of each individual agent among all equilibria. This equilibrium is in particular a Pareto optimum and maximizes social welfare among all equilibria. We first state this result in abstract terms in Lemma 3.1. The lemma together with Proposition 2.4 yields the main structural result.

Lemma 3.1. *Assume that the best-response function Φ is monotonic. Let E be the set of Nash equilibria. Then E is non-empty, and the pair $(E, \geq)$ forms a complete lattice.*

Proof. The set of all strategy profiles is $D = [0, 1]^{n\times m}$, and $(D, \geq)$ is a complete lattice. The map $\Phi : D \to D$ computes, for a given strategy profile y, the best response y_i^* for every agent i with respect to y_{-i}. Thus, $\Phi(y) = (y_1^*, y_2^*, \ldots, y_n^*)$ is the profile after a simultaneous best response of all agents. The fixed points are the Nash equilibria of the fire sale game. Using that Φ is monotonic, we apply the Knaster-Tarski Theorem, and the statement follows. □

Theorem 3.2. *Let E be the set of Nash equilibria. If $\alpha = 1$, then E is non-empty, and the pair $(E, \geq)$ forms a complete lattice.*

Under the assumptions of the theorem, iterating the best-response function further converges to the point-wise maximal equilibrium.

Theorem 3.3. *Let (y^t) be the iteration sequence defined as $y^1 = (1, 1, \ldots, 1)$ and $y^{t+1} = \Phi(y^t)$. If $\alpha = 1$, then (y^t) converges to the point-wise maximal equilibrium.*

Proof. The statement follows from the fact that Φ is monotonic and continuous from above. This is a standard technique and can be seen as special case of the Kleene fixed point theorem (see, e.g., [36, Lemma 3] for a direct proof). Monotonicity follows from Proposition 2.4. We show continuity from above.

Given y_{-i}, let $D_i(y_{-i}) = \{y_i \mid e_i(y_i, y_{-i}) > 0 \wedge \text{lev}_i(y_i, y_{-i}) \leq \lambda\}$. Let $E_i^\lambda = \{y_{-i} \mid D_i(y_{-i}) \neq \emptyset\}$. Proposition 2.3 implies that the set $\{y \mid e_i(y) > 0 \wedge \text{lev}_i(y) \leq \lambda\}$ is closed. This shows that $D_i(y_{-i})$ for any y_{-i} and E_i^λ are closed because they are projections of this set (all involved sets are bounded, so being closed and being compact are equivalent). $D_i(\,\cdot\,)$ is monotonic in the sense that if $y_{-i} \leq y'_{-i}$ point-wise, then $D_i(y_{-i}) \subseteq D_i(y'_{-i})$. This follows from monotonicity of e_i and inverse monotonicity of lev_i in y_{-i} (where $\text{lev}_i(y) > 1$, cf. Proposition 2.1). By the same argument, E_i is monotonic in the sense that if $y_{-i} \leq y'_{-i}$ and $y_{-i} \in E_i$, then $y'_{-i} \in E_i$.

Proposition 2.4 yields $\Phi_i(y) = \varphi_i(y_{-i}) := \max D_i(y_{-i})$ if $y_{-i} \in E_i^\lambda$ and 0 otherwise.

Let now $(y^t_{-i})_t$ be any point-wise decreasing sequence in $[0, 1]^{n-1}$ and let $y^+_{-i} = \lim_t y^t_{-i}$. By monotonicity and closedness of E_i, we know that, if $y^+_{-i} \notin E_i$, then $y^t_{-i} \notin E_i$ for almost all t, so $\Phi_i(y^t_{-i}) = 0$ for almost all t, and we have $\lim_t \Phi_i(y^t_{-i}) = \Phi_i(y^+_{-i})$. So assume $y^+_{-i} \in E_i$ and, thus, $y^t_{-i} \in E_i$ for all t, since $y^t_{-i} \geq y^+_{-i}$. It remains to prove that $\varphi_i(y^+_{-i}) = \lim_t \varphi_i(y^t_{-i})$.

We have $\text{lev}_i(\lim_t \varphi_i(y^t_{-i}), y^+_{-i}) = \lim_t \text{lev}_i(\varphi_i(y^t_{-i}), y^t_{-i}) \leq \lambda$ by continuity of lev_i, and the inequality by definition of φ_i. Thus, $\lim_t \varphi_i(y^t_{-i}) \in D_i(y^+_{-i})$ and thus, by choice of $\varphi_i(y^+_{-i})$, we have $\lim_t \varphi_i(y^t_{-i}) \leq \varphi_i(y^+_{-i})$. On the other hand, by monotonicity of lev_i in y_{-i}, $D_i(y^+_{-i}) \subseteq D_i(y^t_{-i})$ for all t. Thus, by choice of the $\varphi_i(\,\cdot\,)$, we have $\varphi_i(y^+_{-i}) \leq \lim_t \varphi_i(y^t_{-i})$ and obtain equality as required. □

In Theorem 3.3, agents *concurrently* deviate to best responses, i.e., $\Phi(y^t)$ applies best responses simultaneously to each component of the vector y. It is straightforward to observe that the result in the previous theorem can be shown also for any *sequential* best-response dynamics starting from $y^1 = (1, \ldots, 1)$, in which agents deviate one-by-one. We omit a formal adjustment of the proof.

In many game-theoretic scenarios, concurrent deviation gives rise to oscillation. The next example shows that fire sale games are no exception to this rule. We work out the example below for linear prices p_j and $\alpha = 1$. It is easy to see that similar examples exist for $\alpha = 1$ with monotonic prices p_j, or $\alpha \in (0, 1)$ with convex prices.

Example 3.4. Consider a game with two agents, one asset, linear price impact, where $p_1^0 = 1$, and $\alpha = 1$. The external assets and liabilities are such that $a_1^I = a_2^I = 1$ and $l_1 = l_2 = 5/4$. Moreover, both players hold half of the security, i.e., $x_{11} = x_{21} = 1/2$. We assume $\lambda = 6$. If both agents play $y_1 = y_2 = 1$, then the leverage is $1.5/(1/4) = 6 = \lambda$. If agent 1 plays $y_1 = 0$, then agent 2 is illiquid and, thus, her best response is $y_2 = 0$. Now suppose we start in state $y^1 = (1, 0)$ and let the agents deviate simultaneously. The process oscillates between $y^{2i} = (0, 1)$ and $y^{2i+1} = (1, 0)$.

If in the example the two agents deviate sequentially, then after one step an equilibrium is reached. More generally, we show below that for all fire sale games with two agents and monotonic best responses, there can be no cycle in sequential best-response dynamics, no matter from which initial state the dynamics starts.

Proposition 3.5. *Consider a fire sale game with two agents and assume that the best-response function is monotonic. Then every sequential best-response dynamics is acyclic.*

Proof. First consider the case that in some round t we have $\Phi(y^{(t)}) \geq y^{(t)}$, i.e., the best response is at least the current strategy for both agents. By monotonicity of the best-response function, the agents will only keep increasing their strategies, which makes a cycle impossible. A similar argument shows the result when in some round t we have $\Phi(y^{(t)}) \leq y^{(t)}$.

Now suppose that in some round t, the best response for agent 1 is at most $y_1^{(t)}$ and for agent 2 it is at least $y_2^{(t)}$. Suppose agent 1 moves in round $t + 1$. By monotonicity, this decreases the best response for agent 2. After round $t + 1$, the best response of agent 2 is either at most or at least $y_2^{(t+1)} = y_2^{(t)}$. Agent 1 is playing her exact best response (i.e., at most *and* at least the best response).

Hence, one of the previously considered cases applies. A symmetric argument applies if agent 2 moves first. □

We conjecture that there are fire sale games with two agents, for $\alpha = 1$ and monotonic prices or $\alpha \in (0, 1)$ and convex prices, in which the best-response function is not continuous from below. Then given a state where all agents want to increase their strategies, sequential best-response dynamics might not reach an equilibrium in the limit. In contrast, suppose that the agents are in a state where they all want to decrease their strategies. Then the proof of Theorem 3.3 can be applied to show that, for $\alpha = 1$ and monotonic prices or $\alpha \in (0, 1)$ and convex prices, the limit of the best-response dynamics is indeed an equilibrium, for any number of agents.

For three or more agents, if in the initial state there are agents above and below their best response, sequential best-response dynamics can again exhibit cyclical behavior. Below we discuss an example with three agents, $\alpha = 1$ and convex price impact. It is minimal in the sense that for two agents, sequential best-response dynamics are always acyclic in this case.

Example 3.6. Consider a fire sale game with three agents and three assets where $\alpha = 1$ and $\lambda = 6.2$. Furthermore, let $a_i^I = 100$ and $l_i = 90$ for all $i \in \{1, 2, 3\}$ and let the asset holdings be as follows:

$$\begin{array}{lll} x_{11} = 0.8 & x_{12} = 0.2 & \\ & x_{22} = 0.8 & x_{23} = 0.2 \\ x_{31} = 0.2 & & x_{33} = 0.8 \end{array}$$

Let $p_j(y) = 10 \cdot \left(\sum_i x_{ij} y_i\right)^2$ be the price function for all assets $j \in \{1, 2, 3\}$. Note that price impact is convex. A best-response cycle is given by the following table:

y_1	y_2	y_3	$u_1(y)$	$u_2(y)$	$u_3(y)$
1	1	0	$-\infty$	18.08	11.6
0	1	0	11.28	$-\infty$	10.32
0	1	1	11.6	$-\infty$	18.08
0	0	1	10.32	11.28	$-\infty$
1	0	1	18.08	11.6	$-\infty$
1	0	0	$-\infty$	10.32	11.28
1	1	0	$-\infty$	18.08	11.6
	$\vdots$			$\vdots$	

Every agent maximizes her utility by selling as little as possible. Due to convex price impact, each agent is particularly affected by price devaluation of an asset of which she holds a large share. These properties lead to the cyclic behavior shown in the table.

The pointwise maximal equilibrium is obtained when no assets are sold, i.e., y^1. Thus, all strategy profiles played in the best-response cycle are point-wise below the maximal equilibrium.

Also, note that the best-response cycle in this example is robust to different choices of α (e.g., using $\alpha = 0.5$).

4 CONVERGENCE SPEED OF DYNAMICS

In this section, we study two dynamics by which agents may reach an equilibrium: the standard best-response dynamics and a simplified dynamics where agents neglect their own price impact. We examine the convergence of these dynamics towards an equilibrium. We focus on $\alpha = 1$ and linear price impact.

4.1 Best-Response Dynamics

Theorem 3.3 shows that (in particular) for the case $\alpha = 1$, best-response dynamics starting at y^1 always converge to the point-wise maximal equilibrium (where asset sales are collectively minimized). If the sequence proceeds for a finite number of steps, we show that it reaches an *approximate equilibrium* quickly, assuming the numeric values of the game are reasonably large.

The traditional concept of an approximate Nash equilibrium is not appropriate for fire sale games because even a small violation of the leverage constraints leads to an infinite decrease in utility. We therefore define an *approximate equilibrium* as a strategy profile where agents can only improve their equity by a small amount and where the leverage constraints are approximately satisfied.

Definition 4.1. Let y be a strategy profile in a fire sale game and let $\varepsilon > 0$. Then y is an *ε-approximate equilibrium* iff the following hold for every agent i:

(1) If i is liquid, then for any y_i' such that i is liquid for (y_i', y_{-i}) we have $e_i(y) \geq e_i(y_i', y_{-i}) - \varepsilon$.
(2) If i is liquid, then $\mathrm{lev}_i(y) \leq \lambda + \varepsilon$.
(3) If i is insolvent or illiquid, then $y_i \leq \varepsilon$.

An approximate equilibrium might not be close to any exact equilibrium in terms of norm distance in strategy space, a common property of approximate equilibrium concepts (see, e.g., Etessami and Yannakakis [17] for a discussion of this phenomenon in the context of approximate Nash equilibria).

We show that best-response dynamics reach an ε-approximate equilibrium in pseudo-polynomial time.

Theorem 4.2. *Consider a fire sale game with $\alpha = 1$ and linear price impact. Let $x_{\max}$ be the maximum over all values x and x^{-1}, where x is a numeric value contained in the input. Let $\varepsilon > 0$. Then best-response dynamics, after n/ε steps, reaches a point y^* such that $\|\Phi(y^*) - y^*\|_\infty \leq \varepsilon$ and y^* is a $(\mathrm{poly}(x_{\max}) \cdot \varepsilon)$-approximate equilibrium.*

Proof. As long as $\|\Phi(y^*) - y^*\| > \varepsilon$, trivially, some y_i^* decreases by at least ε in every step. As y^* is bounded below by $(0, \ldots, 0)$, there can be at most n/ε such steps.

We now show that y^* is an ε-approximate equilibrium. First consider an insolvent or illiquid agent i. Then $\Phi_i(y^*) = 0$ and thus, since $\|\Phi_i(y^*) - y_i^*\| \leq \varepsilon$, we have $y_i^* \leq \varepsilon$ as required.

Let now i be liquid. Let $y' = \Phi(y^*)$ and assume that $y_i' < 1$ (otherwise, y^* even satisfies the requirements for an *exact* equilibrium at i). Then by choice of y_i' we have $\mathrm{lev}_i(y_i', y_{-i}^*) = \lambda$. We bound the difference $\mathrm{lev}_i(y_i^*, y_{-i}^*) - \mathrm{lev}_i(y_i', y_{-i}^*)$. To do this note that, as the sequence is descending, $y_i' \leq y_i^* \leq y_i' + \varepsilon$. Consider the derivative

$$\frac{\mathrm{d}\ \mathrm{lev}_i(y)}{\mathrm{d}y_i} = \left(\frac{\frac{\mathrm{d}a_i}{\mathrm{d}y_i} e_i - \frac{\mathrm{d}e_i}{\mathrm{d}y_i} a_i}{e_i^2}\right)(y) =: \frac{N(y)}{e_i^2(y)}.$$

As e_i is increasing in y_i, for any $y_i \in [y_i', y_i^*]$ we have $e_i(y_i, y_{-i}^*) \geq e_i(y_i', y_{-i}^*) = a_i(y_i', y_{-i}^*)/\lambda \geq a_i^I/\lambda \geq x_{\max}^{-2}$, where the first equality is because $\mathrm{lev}_i(y_i', y_{-i}^*) = \lambda$.

$N(y)$ is a polynomial (of degree 2) in y where values of the coefficients, but not the structure of the function, depend on the input. Since $y \in [0, 1]^n$, we have $N(y) \leq \mathrm{poly}(x_{\max})$ and thus

$$\frac{\mathrm{d}\ \mathrm{lev}_i(y)}{\mathrm{d}y_i} \leq \mathrm{poly}(x_{\max}).$$

By integration, we have $\text{lev}_i(y^*) \le \lambda + \text{poly}(x_{\max}) \cdot \varepsilon$.

Finally, we show that no liquid agent can improve her equity by more than $\text{poly}(x_{\max}) \cdot \varepsilon$. This follows using the same technique as above because we can bound the derivative $\frac{\mathrm{d}e_i(y)}{y_i}$ to receive $e_i(y^*) \ge e_i(y'_i, y^*_{-i}) - \text{poly}(x_{\max}) \cdot \varepsilon$ and y'_i maximizes $e_i(\,\cdot\,, y^*_{-i})$ by definition. □

4.2 Simplified Best-Response Dynamics

Computing a best response is relatively complex as the agent needs to take into account the price impact of the very sales she is about to decide on. For some agents, this may be unrealistic. Hence, we consider a simplified dynamics, where agents neglect their own price impact. Similar "best-response" dynamics have been considered by Cont and Schaanning [11]. Here, liquidations proceed in several rounds and agents consider prices as fixed during each round.

We first define simplified versions of the different components of each agent's wealth where the price impact of the agent's current choice of strategy is excluded. These functions take one parameter more compared to the full versions in Section 2.1 to differentiate between the current choice and a previous choice by the agent.

Definition 4.3. For $\tilde{y}_i \in [0,1]$ and $y \in [0,1]^N$, define the *simplified* assets, revenue, equity, and leverage as the respective terms from Section 2.1 where, however, the price impact is calculated based on y, but agent i's sales are calculated according to $\tilde{y}_i$. Formally, let

$$\begin{aligned}
\tilde{a}_i(\tilde{y}_i, y) &:= a_i^I + \tilde{y}_i \sum_j x_{ij} p_j(y) &&= a_i^I + \tilde{y}_i V_i(y) \\
\tilde{\Delta}_i(\tilde{y}_i, y) &:= \sum_j x_{ij}(1-\tilde{y}_i) p_j(y) &&= (1-\tilde{y}_i) V_i(y) \\
\tilde{e}_i(\tilde{y}_i, y) &:= \tilde{a}_i + \tilde{\Delta}_i - l_i &&= a_i^I - l_i + V_i(y) \\
\tilde{\text{lev}}(\tilde{y}_i, y) &:= \frac{\tilde{a}_i(\tilde{y}, y)}{\tilde{e}_i(y)} \\
\text{where} \quad V_i(y) &:= \sum_j x_{ij} p_j(y).
\end{aligned}$$

Observe that $V_i(y)$ is the value of i's liquid asset holdings if i sells nothing and price impact is given by y. Further observe that $\tilde{e}_i(\tilde{y}_i, y)$ is in fact independent of $\tilde{y}_i$ and (thus) we have $\tilde{e}_i(\tilde{y}_i, y) = e_i(y)$ for all $\tilde{y}_i$. This is because, under the assumption of the simplified best-response dynamics, sales according to $\tilde{y}_i$ do not generate any additional price impact and thus they transform assets (valued at market price) into risk-free assets at a rate of 1; these terms cancel out in the calculation of the equity. We extend our model by making the realistic assumption that agents still aim to sell as little as possible (i.e., maximize $\tilde{y}_i$) subject to meeting their leverage constraint. In this case, the best response of i according to the simplified dynamics is easily calculated:

Definition 4.4. For $y \in [0,1]^N$ we define

$$g_i(y) := \lambda - \frac{\lambda l_i - (\lambda - 1) a_i^I}{V_i(y)} \quad ,$$

and a *simplified best-response function* $\Psi : [0,1]^N \to [0,1]^N$ with

$$\Psi_i(y) := \begin{cases} \min(1, \max(0, g_i(y))) & \text{if } e_i(y) > 0 \\ 0 & \text{if } e_i(y) \le 0. \end{cases}$$

Lemma 4.5. *The value $\tilde{y}_i := \Psi_i(y)$ is the maximal $\tilde{y}_i$ such that $\tilde{e}_i(\tilde{y}_i, y) > 0$ and $\tilde{\text{lev}}_i(\tilde{y}_i, y) \le \lambda$, if such a $\tilde{y}_i$ exists. Otherwise, $\Psi_i(y) = 0$.*

Proof. Recall that $\tilde{e}(\tilde{y}_i, y) = e_i(y)\ \forall y$. If $e_i(y) \le 0$, then the statement is trivial. So assume $e_i(y) > 0$. It is easy to see that $g_i(y)$ is such that

$$0 = \tilde{a}_i(g_i(y), y) - \lambda \tilde{e}_i(g_i(y), y) = \tilde{a}_i(g_i(y), y) - \lambda e_i(y). \qquad (2)$$

If $e_i(y) \le 0$, then we must have $g_i(y) \le 0$ and thus $\Psi_i(y) = 0$ by definition. If $e_i(y) > 0$, then Equality (2) implies that (a) $g_i(y) \ge 0$ and (b) $\tilde{\text{lev}}_i(g_i(y), y) = \lambda$. Because $\tilde{\text{lev}}_i(\tilde{y}_i, y)$ is strictly monotonic in $\tilde{y}_i$ (because of linear price impact), this implies that $\Psi_i(y) = \min(1, g_i(y))$ has the properties of $\tilde{y}_i$ as needed. □

When agents act according to the simplified best-response function Ψ, they ignore their own price impact. However, the price impact resulting from these sales enters in the next round, where it affects all agents, including the ones who increased their sales in this round. Thus, no information is lost. As the following theorem shows, this implies that simplified best-response dynamics also converge to the maximal equilibrium, just like the best-response dynamics.

Theorem 4.6. *Consider a fire sale game with $\alpha = 1$ and linear price impact. Then the following hold:*

(1) Ψ is monotonic and continuous.
(2) Φ and Ψ have the same fixed points.
(3) Let $(\tilde{y}^t)$ be a sequence of strategy profiles defined by $\tilde{y}^0 = (1, \ldots, 1)$ and $\tilde{y}^{t+1} = \Psi(y^t)$. Then $(\tilde{y}^t)$ is point-wise monotonically decreasing and converges to the point-wise maximal equilibrium of the fire sale game.

Proof. (1) Monotonicity follows from the definition because the condition $e_i(y)$ is increasing in y (i.e., it can only switch from false to true as y increases point-wise, but not vice versa) and the function g_i is obviously monotonic. Towards continuity, first note that g_i is obviously continuous. For continuity between the two cases of the case distinction, it follows from the proof of Lemma 4.5 that, as $e_i(y) \to 0$, $g_i(y)$ converges to a value ≤ 0.

(2) It follows from Lemma 4.5 and $\tilde{e}_i(y_i, y) = e(y)$ and $\tilde{\text{lev}}_i(y_i, y) = \text{lev}_i(y)$. More in detail, if y is any strategy profile for which some agent i is insolvent or illiquid at y_{-i}, then $\Phi_i(y) = 0$ and by Lemma 4.5 also $\Psi_i(y) = 0$. Thus, $\Phi_i(y) = y_i$ iff $y_i = 0$ iff $\Psi_i(y) = y_i$. If agent i is liquid at y_{-i}, then $y_i = \Phi_i(y)$ iff y_i is maximal such that $e_i(y_i, y_{-i}) > 0$ and $\text{lev}_i(y_i, y_{-i}) \le \lambda$. Since $\text{lev}_i(y_i, y_{-i}) = \tilde{\text{lev}}_i(y_i, y)$, Lemma 4.5 implies equivalence to $y_i = \Psi_i(y)$.

(3) By part 1 and the same argument as in Theorem 3.3, the sequence converges to the maximal fixed point of Ψ, i.e., the maximal fixed point of Φ (by part 2) and point-wise maximal equilibrium. □

Φ and Ψ have the same fixed points (i.e., equilibria of the fire sale game) and both converge to the maximal equilibrium. The speed of convergence, however, can be vastly different, as we illustrate next.

4.3 Experiments and Diversification

We now study what properties of the asset holdings matrix x affect the stability of the financial system. Specifically, we are interested in the effect of *diversification*, i.e., to which degree each agent spreads

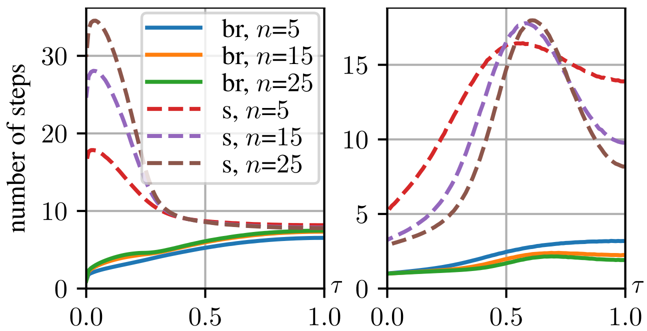

Figure 3: Comparison of the convergence time of the best-response (br) and the simplified dynamics (s) for varying number of banks and diversification τ. The two graphs use different sets of parameters.

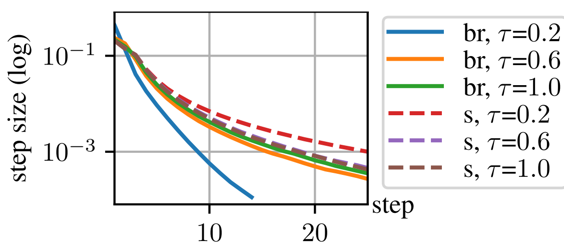

Figure 2: Comparison of the average step size over time of the best-response (br) and the simplified dynamics (s) for various values of diversification τ for games with 10 agents.

her investments across multiple assets. Higher diversification reduces the effect of each asset on each agent (increasing stability), but also increases the number of channels by which one agent may affect another (decreasing stability). We are interested in the convergence speed of both standard and simplified best-response dynamics. To that end, we perform a computational experiment with the required code available online [8]:

We consider games with an equal number of agents and assets $n = m$. For each agent i, we draw a_i^I uniformly in $[80, 120]$ and l_i in $[40, 60]$. Price impact is linear and we set $p_j^0 = 100$ for each asset j. We choose the asset holdings based on a parameter $\tau \in [0, 1]$, which measures diversification. We set $x_{ij} = \tau \cdot 1/n + (1 - \tau) \cdot \delta_{ij}$, where δ_{ij} is the Kronecker symbol.[6] Let λ^1 be the highest leverage of any agent at $y^1 = (1, \ldots, 1)$. We draw λ from $[0.6\lambda^1, 0.99\lambda^1]$ to ensure that y^1 is not a Nash equilibrium. We reject all instances with $\lambda \leq 1$. We let both dynamics converge from y^1 until all strategy changes in a step become smaller than 10^{-5}. Figure 2 depicts the average size of strategy changes over time, while Figure 3 displays the number of steps to convergence. For Figure 3 (right) we uniformly drew a_i^I from $[0, 100]$, l_i from $[40, 100]$, p_j^0 from $[50, 150]$, and λ from $[0.9\lambda^1, 0.99\lambda^1]$, with all other parameters unchanged. For each data point, we average over 10^6 runs.

In Figure 2, we see that the step size over time does not decrease exponentially for both dynamics, which suggests that convergence requires pseudo-polynomially many steps before reaching an approximate equilibrium on average. We also see that diversification has a significant effect on the best-response dynamics, since an agent is less exposed to other agents for low values and thus needs to make smaller corrections after the initial steps. The simplified dynamics, on the other hand, does not exhibit this behavior, as an agent is not accurately measuring her own price impact. Figure 3 shows that for our examples, the simplified dynamics converge more slowly, but for many values of diversification not by a large amount. Interestingly, a "hump" in convergence times appears at a certain diversification value, reminiscent of similar results on systemic risk [16]. However, a change of parameters may lead to wildly different effects of diversification on the convergence of the system for the simplified dynamics. Thus, to draw further conclusions about convergence times, there needs to be sufficient information on the asset holdings.

[6] For $\tau = 0$, each agent only holds a single (different) asset with no agent interactions; for $\tau = 1$, all agents hold all assets equally.

5 DISCUSSION AND CONCLUSIONS

We have studied price-mediated contagion from the perspective of algorithmic game theory. The existence and the shape of equilibria is heavily dependent on the assumptions regarding price impact. For $\alpha = 1$ or convex price impact, the set of equilibria forms a lattice and the Pareto optima form strong equilibria. However, agents face a twofold equilibrium coordination problem: (1) while the maximal equilibrium is the social optimum, the minimal equilibrium can be arbitrarily poor; (2) simplified best-response dynamics can take a long time to converge. This may lead to a delay in resolution and exacerbates a financial crisis. To help alleviate this problem, a regulator may estimate the maximal equilibrium and help guide agents towards it.

Our key insights can be established more generally beyond the case of even sales. Suppose for each asset j and each agent i there is a monotone and continuous *sales function* $f_{ij} : [0, 1] \longrightarrow [0, 1]$. In such a game, when agent i chooses a strategy $y_i \in [0, 1]$, it sells an amount of $f_{ij}(y_i) \cdot x_{ij}$ of asset j. Even sales represent the special case with $f_{ij}(y_i) = y_i$ for all i and j. Monotone sales functions allow to capture further natural behavior of an agent, e.g., selling the assets individually in a fixed priority order.

It is fairly straightforward to see that the central monotonicity results (e.g., Proposition 2.4) extend to this case. As such, our main results on existence, lattice structure, and convergence also hold for games with monotone sales.

For $\alpha \in (0, 1)$ and concave price impact, an equilibrium need not exist. It is an interesting open problem to study the computational complexity of deciding existence of equilibrium in a given fire sale game. Another open problem is characterizing existence and computation of equilibria for non-even-sales.

ACKNOWLEDGMENTS

This work was supported by DFG Research Unit ADYN under grant DFG 411362735.

REFERENCES

[1] Christoph Aymanns, Fabio Caccioli, J. Doyne Farmer, and Vincent Tan. 2016. Taming the Basel leverage cycle. *Journal of Financial Stability* 27 (2016), 263–277.
[2] Christoph Aymanns and J. Doyne Farmer. 2015. The dynamics of the leverage cycle. *Journal of Economic Dynamics and Control* 50 (2015), 155–179.

[3] Michel Baes and Eric Schaanning. 2020. Reverse Stress Testing: Scenario Design for Macroprudential Stress Tests. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3670916 Working Paper.

[4] Tathagata Banerjee and Zachary Feinstein. 2021. Price mediated contagion through capital ratio requirements with VWAP liquidation prices. *Europ. J. Oper. Res.* 295, 3 (2021), 1147–1160.

[5] Bank of International Settlements. 2017. The sterling 'flash event' of 7 October 2016. *BIS Markets Committee Paper* 9 (13 January 2017). https://www.bis.org/publ/mktc09.htm

[6] Basel Committee on Banking Supervision. 2011. Basel III: A global regulatory framework for more resilient banks and banking systems. http://www.bis.org/publ/bcbs189.pdf

[7] Nils Bertschinger, Martin Hoefer, Simon Krogmann, Pascal Lenzner, Steffen Schuldenzucker, and Lisa Wilhelmi. 2022. Equilibria and Convergence in Fire Sale Games. https://doi.org/10.48550/ARXIV.2206.14429

[8] Nils Bertschinger, Martin Hoefer, Simon Krogmann, Pascal Lenzner, Steffen Schuldenzucker, and Lisa Wilhelmi. 2023. Code for "Equilibria and Convergence in Fire Sale Games". https://doi.org/10.5281/zenodo.7661272

[9] Nils Bertschinger, Martin Hoefer, and Daniel Schmand. 2020. Strategic Payments in Financial Networks. In *Proc. 11th Symp. Innov. Theoret. Comput. Sci. (ITCS)*. 46:1–46:16.

[10] Saikat Chatterjee and Trevor Hunnicutt. 2019. Bump in the night: FX flash crashes put regulators on alert. *Reuters* (30 May 2019). https://www.reuters.com/article/us-markets-currencies-insight-idUSKCN1T00J0 News Item.

[11] Rama Cont and Eric Schaanning. 2017. Fire sales, indirect contagion and systemic stress testing. https://ssrn.com/abstract=2541114 Working Paper.

[12] Rama Cont and Eric Schaanning. 2019. Monitoring indirect contagion. *J. Bank. Financ.* 104 (2019), 85–102.

[13] Stéphane Dees, Jérôme Henry, and Reiner Martin (Eds.). 2017. *STAMP€: Stress-Test Analytics for Macroprudential Purposes in the euro area.* European Central Bank, Frankfurt.

[14] Fernando Duarte and Thomas Eisenbach. 2021. Fire-Sale Spillovers and Systemic Risk. *J. Finance* 76, 3 (2021), 1251–1294.

[15] Larry Eisenberg and Thomas Noe. 2001. Systemic risk in financial systems. *Management Science* 47, 2 (2001), 236–249.

[16] Matthew Elliott, Benjamin Golub, and Matthew Jackson. 2014. Financial Networks and Contagion. *Amer. Econ. Rev.* 104, 10 (2014), 3115–53.

[17] Kousha Etessami and Mihalis Yannakakis. 2010. On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.* 39, 6 (2010), 2531–2597.

[18] Maureen Farrel. 2013. Mini flash crashes: A dozen a day. *CNN* (20 March 2013). https://money.cnn.com/2013/03/20/investing/mini-flash-crash/index.html News Item.

[19] John Geanakoplos. 2009. The Leverage Cycle. *Cowles Foundation Discussion Paper* 1715 (2009). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1441943

[20] Giulio Girardi, Kathleen Hanley, Stanislava Nikolova, Loriana Pelizzon, and Mila Sherman. 2021. Portfolio similarity and asset liquidation in the insurance industry. *Journal of Financial Economics* 142, 1 (2021), 69–96.

[21] Brett Hemenway and Sanjeev Khanna. 2016. Sensitivity and computational complexity in financial networks. *Algorithmic Finance* 5, 3–4 (Mar 2016).

[22] Martin Hoefer and Lisa Wilhelmi. 2022. Seniorities and Minimal Clearing in Financial Network Games. In *Proc. 15th Symp. Algorithmic Game Theory (SAGT)*. 187–204.

[23] Stavros D. Ioannidis, Bart de Keijzer, and Carmine Ventre. 2022. Strong Approximations and Irrationality in Financial Networks with Derivatives. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. 76:1–76:18.

[24] Joint CFTC-SEC Advisory Committee on Emerging Regulatory Issues. 2010. Findings regarding the market events of May 6, 2010. https://www.sec.gov/news/studies/2010/marketevents-report.pdf Report.

[25] Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. 2021. Financial network games. In *ICAIF'21: 2nd ACM Int. Conf. AI in Finance*. 26:1–26:9.

[26] Panagiotis Kanellopoulos, Maria Kyropoulou, and Hao Zhou. 2022. Forgiving Debt in Financial Network Games. In *Proc. 21st Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1651–1653. https://doi.org/10.5555/3535850.3536065

[27] Buhong Liu, Maria Polukarov, Carmine Ventre, Lingbo Li, and Leslie Kanthan. 2021. Call Markets with Adaptive Clearing Intervals. In *Proc. 20th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1587–1589. https://doi.org/10.5555/3463952.3464168

[28] Buhong Liu, Maria Polukarov, Carmine Ventre, Lingbo Li, Leslie Kanthan, Fan Wu, and Michail Basios. 2022. The Spoofing Resistance of Frequent Call Markets. In *Proc. 21st Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 825–832. https://doi.org/10.5555/3535850.3535943

[29] Pál András Papp and Roger Wattenhofer. 2020. Network-Aware Strategies in Financial Systems. In *Proc. 47th Int. Colloq. Autom. Lang. Programming (ICALP)*. 91:1–91:17.

[30] Pál András Papp and Roger Wattenhofer. 2021. Debt Swapping for Risk Mitigation in Financial Networks. In *Proc. 22nd Conf. Econ. Comput. (EC)*. 765–784.

[31] Pál András Papp and Roger Wattenhofer. 2021. Default Ambiguity: Finding the Best Solution to the Clearing Problem. In *Proc. 17th Conf. Web and Internet Econ. (WINE)*. 391–409.

[32] Pál András Papp and Roger Wattenhofer. 2021. Sequential Defaulting in Financial Networks. In *Proc. 12th Symp. Innov. Theoret. Comput. Sci. (ITCS)*. 52:1–52:20.

[33] Jonathan Ponciano. 2021. Crypto Flash Crash Wiped Out $300 Billion In Less Than 24 Hours, Spurring Massive Bitcoin Liquidations. *Forbes* (April 2021). https://www.forbes.com/sites/jonathanponciano/2021/04/18/crypto-flash-crash-wiped-out-300-billion-in-less-than-24-hours-spurring-massive-bitcoin-liquidations/ News Item.

[34] Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. 2016. Clearing Payments in Financial Networks with Credit Default Swaps. In *Proc. 17th Conf. Econ. Comput. (EC)*. 759.

[35] Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. 2017. Finding Clearing Payments in Financial Networks with Credit Default Swaps is PPAD-complete. In *Proc. 8th Symp. Innov. Theoret. Comput. Sci. (ITCS)*.

[36] Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. 2020. Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks. *Manag. Sci.* 66, 5 (2020), 1783–2290.