ECAI 2024
U. Endriss et al. (Eds.)

doi:10.3233/FAIA240890

Selfishly Cancelling Debts Can Reduce Systemic Risk

Jinyun Tong[a,*], **Bart De Keijzer**[a,**] **and Carmine Ventre**[a,***]

[a]King's College London
ORCID (Jinyun Tong): https://orcid.org/0009-0002-9155-8730, ORCID (Bart De Keijzer):
https://orcid.org/0000-0001-9465-0837, ORCID (Carmine Ventre): https://orcid.org/0000-0003-1464-1215

Abstract. The exposure of banks to systemic risk in financial networks usually requires large bailouts of taxpayer money with long-lasting and damaging societal consequences. We examine whether the banking network can reduce systemic risk from within by selfishly cancelling the debts of banks in distress. This operation can in principle reduce losses and prevent default cascades. We define an abstract model to simulate the ensuing strategic game on randomly generated financial networks, where each systemically important bank independently decides how likely it is to cancel some debts of insolvent banks. We compute the equilibrium of the induced empirical game with the empirical game-theoretic analysis and analyse its efficiency by measuring the price of anarchy. Our results show that selfish debt cancellation can reduce systemic risk when adopting the equilibrium strategy profile. However, our results also indicate that the efficiency of the equilibrium can be low and relatively few banks cancel debts at equilibrium, and we explain the reason for this through analysis of the banks' incentives and game dynamics.

1 Introduction

Since the 2007-2008 financial crisis, systemic risk in financial systems has become one of the most significant concerns for regulators. Recently, the collapses of Silicon Valley Bank and Credit Suisse have raised concerns about financial instability caused by complex and nontransparent interconnections among firms in the financial system. For example, the inability of a bank to repay its inter-bank loans can lead to shocks to its creditor banks, in turn causing the creditors to default, producing a chain reaction in the banking system.

Reducing systemic risk often requires saving some banks, with governments taking steps in this direction in many instances over the past decades. As such bailouts are tedious and use taxpayers' money, in the present paper, we consider instead whether it is possible for a banking system to "self-heal" in an incentive-compatible way. Our financial systems are modelled in a standard way, through directed arc-weighted and node-weighted networks, where a node represents a financial firm (a *bank*, in short) with its weight representing the bank's a priori available (liquid) *external assets*, whereas an arc represents a debt contract between a debtor and a creditor bank and the arc's weight is the amount owed. If a bank does not have enough assets available to repay its debts, it is *in default* in which case a fraction of the bank's assets is charged to cover the default costs, after which the remaining part of the bank's insufficient funds are distributed to the bank's creditors [19].

The main operation we consider is for creditors to rescue debtors by cancelling some of their debts, in cases where this reduces their own losses to prevent financial contagion. For example, in the single-arc financial network of Figure 1(a), bank A defaults as it cannot repay its liability. Considering default costs of 20% (i.e., A can only call in 80% of its assets when in default), the actual payment of A to B will be 0.8. If B reduces A's liability to 1, A will not default, so B can receive the incoming payment of 1 instead of 0.8. This example shows that debt cancellation can be effective. However, cancelling debts also implies giving up some claims of the debts, which may lead to a decrease in incoming payments if more than one creditor is involved. For example, in Figure 1(b), if A defaults, B and C will receive 0.4, respectively, with the same default cost rate. If both B and C reduce A's liabilities to them to 0.5, A will not default, so they will receive 0.5, which seems beneficial. However, if one of them quits, the other bank has to cancel all the debts it claims to rescue A. In this case, the bank cancelling debts will lose all while the other bank will receive full payment. It is easy to recognise that this game is a prisoner's dilemma if there is no communication. These two toy examples illustrate the dual character of debt cancellation.

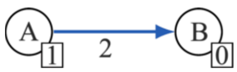

(a) One creditor

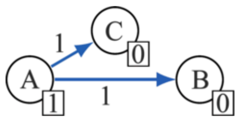

(b) Two creditors

Figure 1. Examples of debt cancellations where the numbers in boxes represent banks' external assets and numbers with which the arcs are labelled represent the notionals of inter-bank loans.

Beyond the above examples, it is important to notice that although in this work we only consider non-cooperative games without a binding coalition, self-enforcing cooperation can also be achieved with proper payoff conditions. Nevertheless, the financial market in reality is much more complicated than our model, including various business alliances and coalitions, and we attempt to propose an abstract model to explain some phenomena from a specific perspective.

In addition, the structural complexity of a banking network contributes to the difficulty of analysing debt cancellation games. Some studies pointed out the scale-free characteristics and core-periphery structures of inter-bank networks, which emphasises the importance of "core" banks in debt cancellation games because core banks as money centres have more inter-bank connections and influence than

* Corresponding Author. Email: jinyun.tong@kcl.ac.uk
** Corresponding Author. Email: bart.de_keijzer@kcl.ac.uk
*** Corresponding Author. Email: carmine.ventre@kcl.ac.uk

peripheral banks [7, 12, 14, 3]. The difference between banks, such as the heterogeneous financial data, generates complicated structures, making it difficult to define a game and compute the equilibrium of such games. Empirical Game-Theoretic Analysis (EGTA) [25, 26] provides a framework for solving this problem. In our work, we use the Eisenberg-Noe model [6], generalised with default costs [19] to model a banking system and perform simulations to calculate empirical payoffs and the equilibrium strategies of agents in these games.

A large body of research has been conducted with similar frameworks based on the Eisenberg-Noe model. For example, Allouch et al. defined a default game and showed that the introduction of a central clearing counterparty allows banks which play different strategies at different Nash equilibria to coordinate on the best equilibrium at no additional cost [2]. Schuldenzucker and Seuken analysed the impact of portfolio compression (i.e., removing liability cycles) on financial networks, showing that portfolio compression may be detrimental to social welfare and banks' incentives may be misaligned with social welfare [20]. Mayo and Wellman also studied portfolio compression at equilibrium using EGTA, finding that portfolio compression may be beneficial when banks strategically remove liability cycles [13]. Kanellopoulos et al. analysed the optimisation and computational complexity problems when financial authorities wish to maximise the total liquidity by injecting cash or removing debts, observing that forgiving some incoming debts (i.e., edges in financial networks) might be in the best interests of banks [8]. They also defined and studied debt transfer games in financial networks by game-theoretic analysis and empirical experiments, proving results about the existence and quality of game equilibrium in which banks maximise their utilities [9]. Papp and Wattenhofer demonstrated the feasibility of bank-to-bank rescue by donations, debt removal and swapping in some specific networks [16, 17].

Our main contributions can be summarised as follows:

1. we present an agent-based model allowing banks to choose to cancel any part of the debts owed to them by insolvent banks, which can lead to self-healing a banking network from within;
2. we provide empirical validation of the existence of equilibria in these debt cancellation games and perform sensitivity analyses based on the model parameters;
3. we analyse the game equilibria by studying incentives, game dynamics and the price of anarchy with a similar definition in [13].

2 Model

In this section, we define a banking network based on the Eisenberg-Noe model [6] and introduce a debt cancellation game, after which we define a meta-game with the heuristic payoff table.

2.1 Banking Network

We consider a set of *banks* $N = \{1, \cdots, n\}$. The *liabilities matrix* $\boldsymbol{L}$ contains liability data in the banking network, where the entry l_{ij} represents the inter-bank liability of bank i to bank j. We assume that each liability is non-negative and a bank does not have a liability to itself, i.e., $l_{ij} \geq 0$ and $l_{ii} = 0$ for all $i, j \in N, i \neq j$. Bank i's total liability, i.e., the sum of its debts, is given by $L_i = \sum_{j \in N} l_{ij}$.

The *relative liabilities matrix* is denoted by $\boldsymbol{R}$ where the entry r_{ij} represents the proportion of the liability of i to j relative to i's total liability. It is given by

$$r_{ij} = \begin{cases} \frac{l_{ij}}{L_i} & \text{if } L_i > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Each bank also has some *external assets*, denoted by e_i, and the external asset vector is denoted by $\boldsymbol{e}$. We assume that the external assets of each bank are sufficient to repay all its debts initially:

$$e_i = \max\left\{c \cdot L_i, L_i - \sum_{j \in N} l_{ji}\right\},$$

where $c \in [0, 1]$ is a constant which we call the *external asset parameter*, and which globally controls the size of the external assets. In our experiments, described in Section 3 below, we then simulate a financial shock by removing the external assets of a bank in the network. This may leads to a default cascade, i.e., insolvency of multiple banks in the network.

A *clearing vector* $\boldsymbol{\Lambda}$, representing the actual payments of debtor banks to creditor banks, is a vector in $\mathbb{R}^n_{\geq 0}$ such that $\boldsymbol{\Lambda} = \Phi(\boldsymbol{\Lambda})$ where function Φ is given by

$$\Phi(\boldsymbol{\Lambda})_i = \begin{cases} L_i & \text{if } e_i + \sum_{j \in N} r_{ji}\Lambda_j \geq L_i, \\ a\left(e_i + \sum_{j \in N} r_{ji}\Lambda_j\right) & \text{otherwise.} \end{cases}$$

In the above function, $a \in (0, 1)$ is a constant called the *recovery parameter*, representing the fraction of the nominal value of assets for liquidation. It is implicitly assumed that a $(1 - a)$ fraction of the total assets of a defaulting bank are spent on default costs. We use the Greatest Clearing Vector Algorithm (GA) [19] to calculate the greatest clearing vector for a given banking network $(\boldsymbol{e}, \boldsymbol{L}, a)$.

Based on the definition of clearing vector, the *value* of a bank is defined as the difference between its actual total assets and its total liabilities, given by

$$v_i = e_i + \sum_{j \in N} r_{ji}\Lambda_j - L_i.$$

A bank is considered *insolvent* if its value is negative, that is $v_i < 0$, which is equivalent to $\Lambda_i < L_i$.

2.2 Debt Cancellation

We allow banks to reduce or remove other banks' debts, as a form of debt restructuring, to rescue them. In principle, each bank can independently decide whether to cancel any debts owed to it by any other bank. Cancellation of debts of an insolvent debtor bank may lead to that bank becoming solvent, which may thus comprise a rescue strategy. The debt cancellation strategy of a bank is invisible to others, resulting in the problem that an insolvent bank may receive excessive resources if more than one creditor bank reduces its debts with its net liabilities. To solve this problem, a bank's *atomic strategy* is defined as whether it joins in the debt cancellation of one particular insolvent bank, given by

$$f_{ij} = \begin{cases} 1 & \text{if } i \text{ joins in the debt cancellation of } j, \\ 0 & \text{otherwise.} \end{cases}$$

The atomic strategies of all banks are organised in a *strategy matrix* $\boldsymbol{F}$.

We assume that the participating banks in cancelling an insolvent bank j's debts will cancel no more debts than the amount $-v_j$ necessary to render j solvent, and furthermore we assume that those banks all cancel an equal fraction of the debt that j owes them: Thus, if a set of banks K manage to rescue insolvent bank j, each $k \in K$ reduces j's liabilities with the factor

$$d_{kj} = \frac{-v_j}{\sum_{k \in K} l_{jk}}.$$

From the above formula, we know that all participant banks will reduce j's debts by the same fraction. We can put all cancellation proportions into a *cancellation matrix* $\boldsymbol{D}(\boldsymbol{F})$ where the (i, j)-entry is given by

$$d(\boldsymbol{F})_{ij} = \begin{cases} \min\left\{1, \dfrac{f_{ij} \cdot (-v_j)}{\sum_{k \in N} f_{kj} l_{jk}}\right\} & \text{if } v_j < 0, \\ 0 & \text{if } v_j \geq 0. \end{cases} \quad (1)$$

Note that here we also assume that the cancellation proportion is bounded by 1 because the banks can at best remove all the liabilities of the insolvent bank. During the debt cancellation process, the liabilities matrix $\boldsymbol{L}$ will be updated according to the cancellation matrix, resulting in a new liabilities matrix $\boldsymbol{L}(\boldsymbol{D})'$, given by:

$$\forall i, j \in N, \quad l(\boldsymbol{D})'_{ji} = l_{ji}(1 - d(\boldsymbol{F})_{ij}). \quad (2)$$

Due to the update of the liabilities matrix, a new banking network $(\boldsymbol{e}, \boldsymbol{L}(\boldsymbol{D})', a)$ is defined. Again its bank value vector $\boldsymbol{v}(\boldsymbol{F})'$ will be calculated by the GA algorithm. The *payoff* of each bank adopting the strategy matrix, denoted by $p(\boldsymbol{F})_i$, is defined by the change in its value, that is

$$p(\boldsymbol{F})_i = v(\boldsymbol{F})'_i - v_i.$$

2.3 Meta-Game

As banks can independently decide whether to join in the debt cancellations by choosing atomic strategies from the atomic strategy space $\{0, 1\}$, there will be $(2^{n-1})^n$ different strategy profiles in total, and it can be difficult to analyse a game of this size.

To solve this problem, we replace this underlying game with a *meta-game* [18, 22]. A meta-game is a simplified game where players apply *meta-strategies* instead. In our work, we assume that banks determine their strategies probabilistically. Specifically, we replace the strategy space from atomic strategies $\{0, 1\}$ to an interval $[0, 1]$; each number in the interval represents a probability of choosing atomic strategy 1. We split this interval into three sub-intervals, representing three meta-strategies $\boldsymbol{S} = \{S_1, S_2, S_3\}$ where $S_1 = \left[0, \frac{1}{3}\right], S_2 = \left(\frac{1}{3}, \frac{2}{3}\right], S_3 = \left(\frac{2}{3}, 1\right]$, referring to a low, medium and high probability of participating in debt cancellations, respectively. Instead of selecting atomic strategies, every bank picks a meta-strategy before debt cancellations, thus composing a *meta-strategy profile*, $\boldsymbol{s} = \{s_1, \cdots, s_n\}$ where $s_i \in \boldsymbol{S}$. In each simulation, a probability is chosen uniformly at random for each bank according to its meta-strategy, determining strategy matrix $\boldsymbol{F}(\boldsymbol{s})$:

$$f_{ij}(\boldsymbol{s}) = \begin{cases} 1 & \text{w.p. } \xi_i \\ 0 & \text{w.p. } (1 - \xi_i) \end{cases}, \quad \xi_i \sim U(s_i) \quad (3)$$

where $U(s_i)$ represents the uniform distribution on the interval defined by the meta-strategy s_i.

In this way, we shift the research object from specific bank-to-bank actions to how likely they are to participate in debt cancellation actions. This reduces the size of the game in terms of the number of strategy profiles to 3^n, but even in this case it may be still challenging to collect adequate payoff data and compute the equilibrium. To further reduce the size of the game, we replace payoff matrix with the *heuristic payoff table* [24, 23] in which the payoffs of meta-strategies are stored as a function of only *the number of banks* using them. Let $\mathcal{N}$ be a table where each entry corresponds to a *meta-strategy cardinality profile* ℓ, that describes the number of banks assigned to each strategy of the meta-strategy space $\boldsymbol{S}$. In other words, ℓ is a profile $(\ell_1, \ell_2, \ell_3) \in \mathbb{N}_0^3$, $\ell_1 + \ell_2 + \ell_3 = n$, showing how many banks use each meta-strategy. For a meta-strategy $S_j, j \in \{1, 2, 3\}$, the value $\mathcal{N}(\ell, j)$ in the table represents the payoff to a bank using meta-strategy S_j when for all $k \in \{1, 2, 3\}$, the number of banks playing meta-strategy S_k equals ℓ_k. Using such a heuristic payoff table, we further reduce the size of the game to the number of rows in $\mathcal{N}$, that is $C^n_{n+|\boldsymbol{S}|-1}$.

Heuristic payoff tables have been proven effective for symmetric games [23], but they may not be directly applicable to our model, as our game is not symmetric: payoffs may vary if the banks interchange their strategies, due to their heterogeneous financial data. Furthermore, even within the set of banks using the same meta-strategy, the payoff of each bank generally differs because of the heterogeneity, so we cannot precisely determine the payoffs for each entry in $\mathcal{N}$. Consequently, we let $\mathcal{N}$ denote *average* payoffs instead, and to generate $\mathcal{N}$ we run simulations to generate a number of $M = 100$ estimations $\mathcal{N}^1, \ldots, \mathcal{N}^M$ of $\mathcal{N}$, and finally set $\mathcal{N}$ to the average of them. For each $\mathcal{N}^k, k \in \{1, \ldots, M\}$, we set each entry ℓ selecting a random profile of meta-strategies $\boldsymbol{s}^k$ such that $|s_j^k(\ell)| = \ell_j$ for each $j \in \{1, 2, 3\}$, where $s_j^k(\ell)$ denotes the set of banks choosing meta-strategy S_j under $\boldsymbol{s}^k$. We set the payoff of the entry $\mathcal{N}^k(\ell, j)$ to the average payoff of all the banks in $s_j^k(\ell)$:

$$\mathcal{N}^k(\ell, j) = \begin{cases} \dfrac{1}{\ell_j} \displaystyle\sum_{i \in s_j^k(\ell)} \Big(v(\boldsymbol{F}(\boldsymbol{s}^k(\ell)))'_i - v_i\Big) & \text{if } \ell_j > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

As stated above, we set M to 100 in our setup. Although an increase in sample size can further reduce the standard error of the mean, hypothesis testing shows that the total variation distance between the stationary distributions computed by 100 samples and by 2048 samples is identical to 0 with a significance level of 0.05. Therefore, it can be reasonable to use the stationary distribution computed by 100 samples to approximate the real stationary distribution.

3 Method

Based on the game model defined above, the experimental procedure is introduced in this section, including the values of the model parameters. Then, we also briefly introduce the equilibrium solver used in our research, the α-Rank algorithm.

3.1 Experimental Procedure

We work with randomly generated networks and a range of choices of model parameters. The external asset parameter c is taken from the set $\{0, 0.05, \cdots, 0.95, 1\}$ while we choose the recovery parameter a from $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. At the beginning of the experimental procedure, we generate a directed scale-free network with $n = 100$ banks according to the algorithm introduced in [5] (with fixed parameters $\alpha = 0.1, \beta = 0.8, \gamma = 0.1, \delta_{in} = \delta_{out} = 4.5$ as in [3]). Then, we consider the 10 banks with the highest in-degrees, and mark them as the *systemically important banks*. We assume that only they can decide whether to cancel the (partial) debts they claim and other banks merely join in the clearing process. We do this to keep the computation time of our experience within feasible bounds, and it can furthermore be considered a reasonable choice informed by practice, as according to published information, there are around 100 UK-headquartered banks [21] and circa 10 systemically important banks in the UK [4]. This setup yields generated networks that are realistic as they have the typical core-periphery structure that has been observed in practice. Each directed edge in the scale-free network represents an inter-bank loan in the banking network.

The notional amount of an inter-bank loan is the absolute value of a random number following the standard normal distribution, i.e., $l_{ij} = |z|, z \sim N(0,1), \forall i,j \in N, i \neq j$. After the liabilities matrix $\boldsymbol{L}$ is established, the external asset vector $\boldsymbol{e}$ is determined by $\boldsymbol{L}$ and c, and this yields our initial banking network $(\boldsymbol{e}, \boldsymbol{L}, a)$.

For each generated banking network, we then run Algorithm 1, which triggers a default cascade by removing external assets of a random bank whose out-degree is greater than 5 but less than 10. This is intended to simulate the collapse of a mid-sized bank (e.g., the Silicon Valley Bank's collapse in 2023 [1]). The algorithm then computes an equilibrium of the debt cancellation game on the resulting network, with the player set restricted to the 10 systemically important banks, using the α-Rank algorithm [15]. The computation of this equilibrium uses the heuristic payoff table which is computed by Algorithm 1 prior to running α-Rank.

The algorithm outputs various quantities we are interested in for our analysis: i) the number I of insolvent banks without debt cancellation; ii) the average reduction n_r^* in the number of insolvent banks as a result of debt cancellation (at equilibrium); iii) a vector $\mathcal{N}^*$ of average numbers of players choosing each of the meta-strategies; iv) the maximum reduction n_r^+ in the number of insolvent banks among all meta-strategy cardinality profiles ℓ; and v) the meta-strategy cardinality profile $\mathcal{N}^+$ where n_r^+ is attained.

The experimental procedure captured by Algorithm 1 will then be repeated 100 times for each combination of the two parameters, with different random banking networks, to further reduce the impact of randomness.

Algorithm 1 Experimental procedure

Input: banking network $(\boldsymbol{e}, \boldsymbol{L}, a)$
Parameter: meta-strategy space $\boldsymbol{S}$
Output: $I, n_r^*, \mathcal{N}^*, n_r^+, \mathcal{N}^+$

Randomly remove external assets of a mid-sized bank, yielding updated external asset vector $\boldsymbol{e}'$.
Clear $(\boldsymbol{e}', \boldsymbol{L}, a)$ using GA, and denote the number of insolvent banks by I.
for each entry ℓ of $\mathcal{N}$ **do**
  for k in $\{1, 2, \cdots, 100\}$ **do**
    Randomly select a meta-strategy profile $\boldsymbol{s}^k(\ell)$, such that $|s_j^k(\ell)| = \ell_j$ for all $j \in \{1,2,3\}$
    Calculate strategy matrix $\boldsymbol{F}(\boldsymbol{s})$ using (3)
    Calculate cancellation matrix $\boldsymbol{D}(\boldsymbol{F})$ using (1)
    Cancel debts according to $\boldsymbol{D}(\boldsymbol{F})$
    Update $\boldsymbol{L}$ to $\boldsymbol{L}(\boldsymbol{D})'$ using (2)
    Clear $(\boldsymbol{e}', \boldsymbol{L}(\boldsymbol{D})', a)$ using GA and obtain $\mathcal{N}^k(\ell, j)$ for all $j \in \{1,2,3\}$ using (4). Denote the number of insolvent banks by $I^k(\ell)$
  end for
  Set $\mathcal{N}(\ell, j) = \sum_{k=1}^{100} \mathcal{N}^k(\ell, j)/100, \forall j \in \{1,2,3\}$
  Set $I(\ell) = \sum_{k=1}^{100} I^k(\ell)/100$
end for
Compute equilibrium π of the game induced by $\mathcal{N}$ using the α-Rank algorithm.
Calculate $n_r^* = \sum_\ell \pi_\ell (I - I(\ell))$
Calculate $\mathcal{N}^* = \sum_{\mathcal{N}_\ell} \pi_\ell \mathcal{N}(\ell, \cdot)$
Calculate $n_r^+ = \max_\ell (I - I(\ell))$
Calculate $\mathcal{N}^+ = \arg\max_\ell (I - I(\ell))$
return $I, n_r^*, \mathcal{N}^*, n_r^+, \mathcal{N}^+$

3.2 α-Rank

In principle, the α-Rank algorithm presents a discrete-state imitation dynamics, where the logistic selection function defines the imitation protocol and the transition probabilities between pure strategy profiles are given based on the imitation protocol. Specifically, the imitation protocol is defined as the fixation probability $\rho^i_{\sigma,\tau}(\boldsymbol{s})$ which describes in the evolutionary population i, how likely all the members playing the original strategy σ copy the mutant strategy τ so that τ takes over the whole population, while the strategies played by other evolutionary populations $\boldsymbol{s}_{-i}$ are fixed. The fixation probability is given by

$$\rho^i_{\sigma,\tau}(\boldsymbol{s}) = \begin{cases} \frac{1-e^{-\alpha(u(\tau,\boldsymbol{s}_{-i})-u(\sigma,\boldsymbol{s}_{-i}))}}{1-e^{-m\alpha(u(\tau,\boldsymbol{s}_{-i})-u(\sigma,\boldsymbol{s}_{-i}))}} & u(\tau,\boldsymbol{s}_{-i}) \neq u(\sigma,\boldsymbol{s}_{-i}), \\ \frac{1}{m} & u(\tau,\boldsymbol{s}_{-i}) = u(\sigma,\boldsymbol{s}_{-i}), \end{cases}$$

where $u(\sigma, \boldsymbol{s}_{-i})$ and $u(\tau, \boldsymbol{s}_{-i})$ capture i's payoffs playing strategy σ and τ, respectively, m is the number of members in each population, and α is the ranking intensity which controls the selection strength (see [15] for more details). Based on this fixation probability, we define a Markov chain where a transition probability is assigned to each pair of pure strategy profiles $(\sigma, \boldsymbol{s}_{-i})$ and $(\tau, \boldsymbol{s}_{-i})$. In principle, the stationary distribution of this Markov chain forms predictions about the very long-run dynamics of the evolutionary populations.

However, in practice, it may not be straightforward to determine a high enough ranking intensity to guarantee a single stationary distribution of the Markov chain in some cases. To solve this problem, we replace the above definition of fixation probability with the following one from OpenSpiel [11]:

$$\rho^i_{\sigma,\tau}(\boldsymbol{s}) = \begin{cases} \epsilon & \text{if } u(\tau,\boldsymbol{s}_{-i}) < u(\sigma,\boldsymbol{s}_{-i}) \\ 0.5 & \text{if } u(\tau,\boldsymbol{s}_{-i}) = u(\sigma,\boldsymbol{s}_{-i}) \\ 1-\epsilon & \text{if } u(\tau,\boldsymbol{s}_{-i}) > u(\sigma,\boldsymbol{s}_{-i}) \end{cases} \quad (5)$$

where $\epsilon \in (0, 0.5]$ represents the minimal fixation probability, and we set $\epsilon = 0.01$ in our work. The modified definition can be considered an approximation of the original definition where ranking intensity $\alpha \to \infty$. Note that these fixation probabilities need to be multiplied by a normalising factor μ to yield valid transition probabilities. More importantly, as the transition probability between a strategy profile and each of its mutant strategy profiles is not less than $\mu\epsilon$, the Markov chain is irreducible (i.e., every state has a positive probability to be visited). Therefore, there will be a unique stationary distribution which captures the evolutionary stable probability distribution over all strategy profiles (i.e., the equilibrium of the game). In addition, (5) has the specific property that the fixation probability between two strategies depends on the rank difference between their payoffs rather than the payoff difference. Therefore, small deviations in payoffs will not alter the fixation probabilities, which provides a strong argument in favour of our approach of using the heuristic payoff table to approximate the real stationary distribution.

4 Results and Evaluation

In this section, we present the experimental results and evaluate the equilibria of the empirical games by analysing the strategy choice of banks, game dynamics, the effects of debt cancellations and the price of anarchy.

4.1 Strategy Choice

We analyse the equilibria computed through the α-Rank algorithm in Algorithm 1, resulting from our experiments. Figure 2 shows the

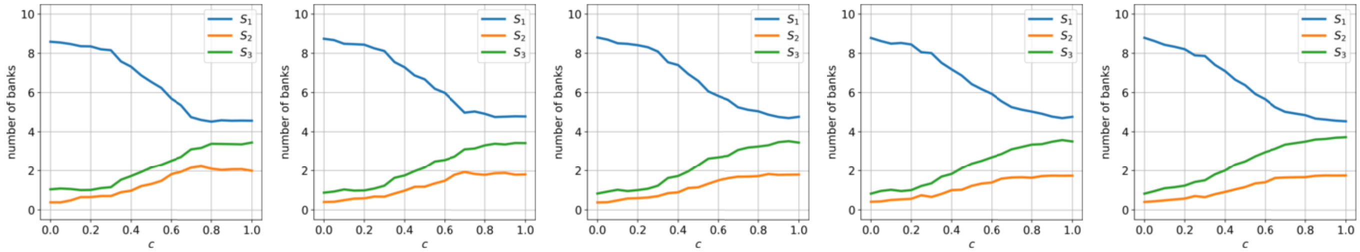

(a) $a = 0.1$. (b) $a = 0.3$. (c) $a = 0.5$. (d) $a = 0.7$. (e) $a = 0.9$.

Figure 2. Average usage of all meta-strategies at equilibrium under varying choices of external asset parameter c and recovery parameter a.

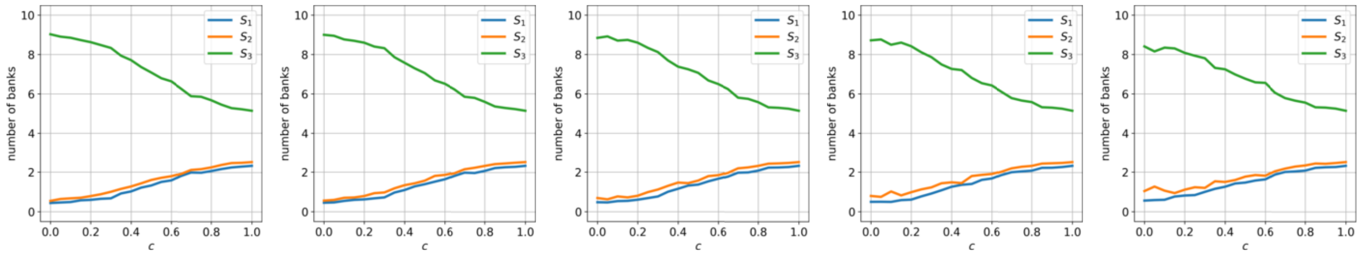

(a) $a = 0.1$. (b) $a = 0.3$. (c) $a = 0.5$. (d) $a = 0.7$. (e) $a = 0.9$.

Figure 3. Average usage of all meta-strategies in the optimal strategy profile under varying choices of external asset parameter c and recovery parameter a.

average number of systemically important banks (hereinafter banks) choosing each of the meta-strategies at equilibrium in the banking networks generated, under varying choices of external asset parameter c and recovery parameter a. From Figure 2, we observe that the usage of each of the meta-strategies is dependent on c but seems to be essentially independent of a. In general, the number of banks participating in debt cancellations with low probability (i.e., S_1) is always greater than the sum of those reducing debts with medium and high probability (i.e., S_2 and S_3).

On the other hand, Figure 3 shows the average number of banks choosing each of the meta-strategies at the optimal solution, i.e., the strategy profile minimising the number of insolvent banks. Differently from the equilibrium strategy profiles, here the banks are more likely to reduce debts to rescue more insolvent banks. However, in both situations, with the increase in c, the numbers of banks using each of the meta-strategies get closer together. The main reason is that with the growth of c, the size of the default cascade triggered by the default of a single bank becomes increasingly smaller (see Figure 4(a) below). For $c = 1$, there even is no such cascade at all, as all banks other than the one explicitly targeted have sufficient external assets to repay their liabilities. Note that there are no differences between the three meta-strategies for a bank that has no arc coming in from an insolvent bank. As a result, if c is large enough, some banks will have no insolvent debtors and thus will choose any of the meta-strategies equiprobably, leading to an evenly distributed state.

Considering the difference in strategy choice at equilibrium and the optimal situation, we are interested in understanding better why the banks are not always willing to cooperate in rescuing insolvent banks at equilibrium, even if it benefits the whole banking system. To answer this question, we analyse the banks' incentives and game dynamics in the next subsection.

4.2 Dynamics Analysis

Since the inter-bank liabilities are independent, creditor banks can address the liabilities of different debtor banks separately. Therefore, we can divide a complicated banking network into independent games, each containing only one insolvent bank. To further simplify the analysis, we focus on two typical cases.

Homogeneous banking network. First of all, we consider a simple homogeneous banking network. Let i denote the only insolvent bank, and its total assets and total liabilities are denoted by A_i and L_i, respectively, satisfying $0 < A_i < L_i$. Let m denote the number of creditor banks of i, each of which can potentially cancel some of i's debts. Among these banks, m_1 is the number of banks which intend to rescue i by removing part of its liabilities and $m_2 = m - m_1$ is the number of banks which do not cancel any debts. In a homogeneous banking network, the liabilities of i to all creditors are identical, denoted by l_{im}, thus $\sum_{k=1}^{m} l_{ik} = ml_{im} = L_i$. If i is solvent after the debt cancellation process, we will get the following clearing equation

$$m_1\lambda_{im_1} + m_2\lambda_{im_2} = A_i$$

where λ_{im_1} and λ_{im_2} represent the actual payments of i to the banks reducing debts and to the banks which do not cancel any debts, respectively. Clearly, $m_1 \geq 1$ because there must be at least one bank which removes part of debts to rescue i. Let d denote the cancellation proportion, we have $\lambda_{im_1} = l_{im}(1-d)$ and $\lambda_{im_2} = l_{im}$, so

$$\begin{aligned} m_1 l_{im}(1-d) + (m - m_1) l_{im} &= A_i \\ d &= \frac{L_i - A_i}{m_1 l_{im}}. \end{aligned}$$

Therefore, creditor banks' incoming payments are

$$\lambda_{im_1} = l_{im} - \frac{L_i - A_i}{m_1}, \qquad \lambda_{im_2} = l_{im}.$$

Clearly, $\lambda_{im_1} < \lambda_{im_2}$, independent of m_1, which indicates that the cancellation strategy is dominated by the non-cancellation strategy, so banks naturally prefer not to cancel any debts. Thus, in equilibrium, m_1 will be the least value such that i is saved even if the banks choosing the cancellation strategy give up all the debts they claim (i.e., $d = 1$). The critical value of the above clearing equation is

$$m_1^* = m - \left\lfloor \frac{A_i}{l_{im}} \right\rfloor, \qquad m_2^* = \left\lfloor \frac{A_i}{l_{im}} \right\rfloor, \tag{6}$$

where m_1^* indicates the minimum number of banks reducing debts for successful rescue of i.

Consider now a solution (m_1', m_2') to the above system where $m_1' < m_1^*$ and $m_2' > m_2^*$. Even though the banks choosing the cancellation strategy remove all i's liabilities to them, bank i is still insolvent. In this case, the equations turn into

$$\begin{aligned} m_1'\lambda_{im_1}' + m_2'\lambda_{im_2}' &= aA_i \\ m_1'(1-1) + m_2'\lambda_{im_2}' &= aA_i \\ \lambda_{im_2}' &= \frac{aA_i}{m_2'} \end{aligned}$$

where $m_2' \geq 1$ because i must be solvent if all banks agree to reduce its debts. Therefore, creditor banks' incoming payments are

$$\lambda_{im_1}' = 0, \qquad \lambda_{im_2}' = \frac{aA_i}{m_2'}. \tag{7}$$

Again, $\lambda_{im_1}' < \lambda_{im_2}'$, independent of m_2', which indicates that the cancellation strategy is still dominated by the non-cancellation strategy. This will naturally lead to a state where, in case i is not rescued, no bank uses the cancellation strategy.

However, at the point of critical value, the cancellation strategy can bring a higher payoff than the non-cancellation strategy. Specifically, recall that in (6), we have $\lambda_{im_1} \to l_{im}$ when $A_i \to L_i$. On the other hand, according to (6) and (7), for the minimum $m_2' = m_2^* + 1$,

$$\lambda_{im_2}' = \frac{aA_i}{m_2^* + 1} = \frac{aA_i}{\left\lfloor \frac{A_i}{l_{im}} \right\rfloor + 1} < al_{im}.$$

Therefore, for any given recovery parameter $a \in (0,1)$, there exists an A_i such that $\lambda_{im_1} > \lambda_{im_2}'$. However, note that the latter strict inequality does not always hold for given a and A_i. In other words, the ultimate result is either all banks do not cancel any debts, or m_1^* banks remove part of i's debts at equilibrium, depending on the values of a and A_i.

Polarised banking network. Secondly, we consider a banking network where i is the only insolvent bank, and there is a single bank (denoted by j) claiming a large enough inter-bank liability l_{ij} with a further k banks claiming identical small liabilities l_{ik}, satisfying

$$l_{ij} + kl_{ik} = L_i, \quad l_{ij} \to L_i.$$

In this situation, $A_i < l_{ij}$; therefore, i cannot survive even though the k banks give up all the debts they claim unless j participates in the debt cancellation.

If j agrees to reduce i's debts, the following equation holds

$$\lambda_{ij} + k_1\lambda_{ik_1} + k_2\lambda_{ik_2} = A_i$$

where k_1 represents the number of banks removing part of liabilities and k_2 the number of banks which do not cancel any debts, satisfying $k_1, k_2 \geq 0, k_1 + k_2 = k$. Let d denote the cancellation proportion, we have $\lambda_{ij} = l_{ij}(1-d)$, $\lambda_{ik_1} = l_{ik_1}(1-d)$ and $\lambda_{ik_2} = l_{ik_2}$, so

$$\begin{aligned} l_{ij}(1-d) + k_1 l_{ik}(1-d) + k_2 l_{ik} &= A_i \\ d &= \frac{L_i - A_i}{l_{ij} + k_1 l_{ik}} \end{aligned}$$

Therefore, creditor banks' incoming payments are

$$\lambda_{ij} = l_{ij} \cdot \frac{A_i - k_2 l_{ik}}{L_i - k_2 l_{ik}}, \quad \lambda_{ik_1} = l_{ik} \cdot \frac{A_i - k_2 l_{ik}}{L_i - k_2 l_{ik}}, \quad \lambda_{ik_2} = l_{ik}.$$

Clearly, $\lambda_{ik_1} < \lambda_{ik_2}$, so for the k banks, the cancellation strategy is always dominated by the non-cancellation strategy. Therefore, at equilibrium, $k_1 = 0$ and only j reduces i's debts. In this case, banks' incoming payments become

$$\lambda_{ij} = A_i - kl_{ik}, \qquad \lambda_{ik} = l_{ik}.$$

On the other hand, if j does not reduce i's debts, i is not rescued, and a solution (k_1', k_2') is attained that satisfies

$$\begin{aligned} \lambda_{ij}' + k_1'\lambda_{ik_1}' + k_2'\lambda_{ik_2}' &= aA_i \\ \lambda_{ij}' + k_1'(1-1) + k_2'\lambda_{ik_2}' &= aA_i \end{aligned}$$

Creditor banks' incoming payments are in this case given by

$$\lambda_{ij}' = aA_i \cdot \frac{l_{ij}}{l_{ij} + k_2' l_{ik}}, \quad \lambda_{ik_1}' = 0, \quad \lambda_{ik_2}' = aA_i \cdot \frac{l_{ik}}{l_{ij} + k_2' l_{ik}}.$$

Again, $\lambda_{ik_1}' < \lambda_{ik_2}'$, so a state will emerge where the k banks do not cancel i's debts. In this case, banks' incoming payments become

$$\lambda_{ij}' = aA_i \cdot \frac{l_{ij}}{L_i}, \qquad \lambda_{ik}' = aA_i \cdot \frac{l_{ik}}{L_i}.$$

Since $l_{ij} \to L_i$, we have $\lambda_{ij} \to A_i$ and $\lambda_{ij}' \to aA_i$. Therefore, for a given $a \in (0,1)$, we have $\lambda_{ij} > \lambda_{ij}'$, indicating that j will always reduce i's debts to rescue it. Nevertheless, note that in this situation, only one bank reduces debts.

To summarise the analysis in this subsection: in a homogeneous banking network, either no bank cancels debts, or some cooperate in rescuing the insolvent bank. In a polarised banking network, only the bank claiming the majority of the debts of the insolvent bank will reduce its liabilities. We may consider these as two extremes, in the sense that the emerging strategic behaviour in other financial networks falls somewhere in between these cases. We further observe that under both extremes, either no debts are cancelled, or an as-small-as-possible number of banks claiming large enough debts remove some at equilibrium. The incentive and game dynamics analyses in this subsection explain why it is difficult for creditor banks to cooperate in rescuing insolvent banks, consistent with the experimental results shown in Section 4.1 at a higher level.

4.3 Effects Analysis

We proceed to quantitatively analyse the impact of debt cancellations at equilibrium. Figure 4 shows the number of insolvent banks before debt cancellations and the reduction in defaults at equilibrium and in optimal cases. From Figure 4(a), we can see that when external asset parameter $c = 0$, implying a state in which there are no surplus external assets in the banking network, a financial shock caused by removing a bank's necessary external assets will trigger a large-scale default cascade and cause a number of banks to default. In this case, the size of the default cascade significantly depends on the value of

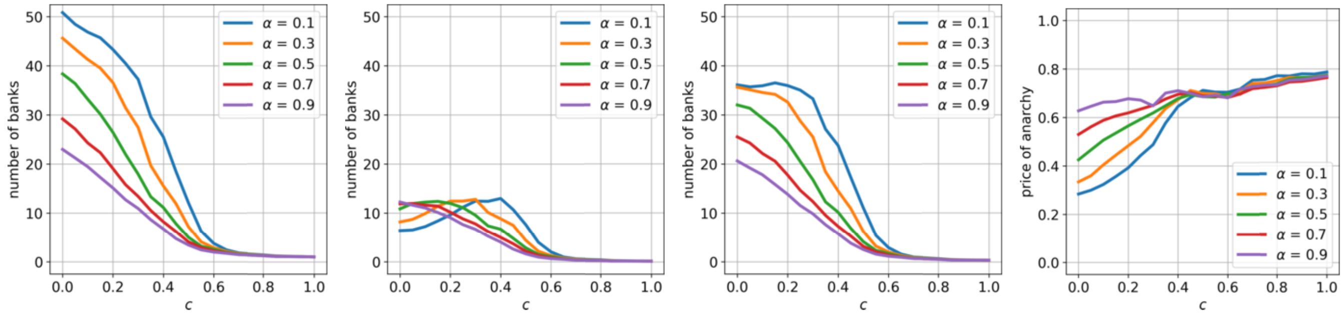

(a) Before debt cancellations (b) Equilibrium rescues (c) Optimum rescues (d) Price of Anarchy

Figure 4. Average number of banks in various profiles and PoA with different external asset parameters c and recovery parameters a.

the recovery parameter a. When $a = 0.1$, over half of the banks are in default, and when $a = 0.9$, about a quarter of the banks cannot repay their debts. On the other hand, when $c \geq 0.5$, implying a state in which the banks hold some surplus external assets, the size of the default cascade is usually small and the recovery parameter will have a minor impact on that. In general, with the increase in external asset parameter c, the number of insolvent banks decreases dramatically as the surplus external assets can be used to repay liabilities in case banks cannot receive sufficient incoming payments. In addition, for any given c, with the increase in recovery parameter a, the number of insolvent banks declines, because the losses caused by a default cascade to the creditor banks become smaller.

Figure 4(b) and Figure 4(c) show the reductions in the number of insolvent banks at equilibrium and at the optimum, respectively. From Figure 4(b), we recognise that debt cancellations effectively rescue insolvent banks and weaken the damage of default cascades, even though few banks reduce debts at equilibrium. For each recovery parameter a, with the increase in the external asset parameter c, the number of rescued banks first increases and then decreases, so there exists a peak. In general, an increase in a appears to move the peak towards smaller c. The main reason for these phenomena is that according to the results shown in Section 4.1, a hardly influences strategy choice, so there will be a similar number of banks joining in the debt cancellation at the same level of c. As stated above, the higher the a, the smaller the losses caused by a default cascade, so the lower the debt cancellation proportion. Therefore, it can be more likely for the participant banks to save the insolvent banks successfully, which raises the number of rescued insolvent banks up. In addition, Figure 4(c) indicates that most insolvent banks can be rescued under the optimum, for many of the parameter choices. Generally, the larger the size of the default cascade, the more insolvent banks can receive rescues at the optimum. However, with the increase in a, the ratio of the rescued banks to the total insolvent banks can be higher for the same reason mentioned above.

Comparing Figures 4(b) and 4(c), we observe a big difference between the equilibrium versus optimum performances. To measure the performance at equilibrium, we calculate its *price of anarchy (PoA)*, given by

$$\mathrm{PoA} = \begin{cases} \frac{n_r^*}{n_r^+} & \text{if } n_r^+ > 0, \\ 1 & \text{otherwise,} \end{cases}$$

where n_r^* and n_r^+ are defined in Algorithm 1 as number of banks rescued at equilibrium and optimum, respectively. Figure 4(d) shows the average PoA of repetitive experiments for 100 times. From Figure 4(d), we see that the recovery parameter a has an important impact on the PoA, especially when the banking network is short of surplus external assets. Specifically, when the external asset parameter c is small, the PoA under larger a is significantly higher than the PoA under a smaller a. With the increase in c, the PoA keeps increasing and the gaps between the PoAs at different levels of a become smaller, converging to the same value at around $c = 0.5$. Nevertheless, the average PoA never exceeds 0.8 in any of our experiments, and this quantifies the empirical loss of performances due to selfishness to at least 20%, indicating that the performances of selfish debt cancellations are not perfect.

5 Conclusion

Our research studied debt cancellation as an inter-bank rescue strategy for the banking system to "self-heal" in an incentive-compatible way, and to what extent selfish debt cancellations can reduce systemic risk caused by default cascades.

The experimental results show that the average number of banks participating in debt cancellations with low probability is always greater than the sum of those reducing debts with medium and high probability in the equilibrium situation. On the other hand, in the optimal situation, banks are more likely to reduce debts to rescue more insolvent banks. To explain this difference, we analysed the incentives and game dynamics and concluded that for any independent debt cancellation games, either no debts are cancelled, or an as-small-as-possible number of banks claiming large enough debts are willing to remove some at equilibrium.

Our results indicate that debt cancellation can indeed save insolvent debtor banks and reduce systemic risk when the banks act according to an equilibrium strategy profile which is compatible with the incentives of the individual banks, although there are usually only a few banks joining in the debt cancellations. In further analysis, we also find that although the recovery parameter hardly influences the strategy choice, it can have an impact on the effects of the debt cancellations by affecting the rescue difficulty. Finally, we calculated the price of anarchy to measure the overall performance of the equilibria, finding that when the banks in the system are short of money, the performance of the equilibria increases with the recovery parameter. In addition, although the PoA keeps increasing with the growth of external assets, it never exceeds 0.8 in any of our experiments.

An important limitation of our work is that only non-cooperative games are considered and the communication among banks is ignored. Therefore, one direction for future work is to extend the current model to a cooperative game model, where banks can form a binding coalition and distribute payoffs. This attempt can simulate the real-world financial market from another view, leading to a deeper insight into how cancelling debts can reduce systemic risk.

Acknowledgements

This work was partially supported by the UKRI Trustworthy Autonomous Systems Hub (EP/V00784X/1). Bart de Keijzer was partially supported by EPSRC grant (EP/X021696/1). This work was also supported by the King's Computational Research, Engineering and Technology Environment (CREATE) at King's College London [10].

References

[1] M. Akhtaruzzaman, S. Boubaker, and J. W. Goodell. Did the collapse of silicon valley bank catalyze financial contagion? *Finance Research Letters*, page 104082, 2023.

[2] N. Allouch, M. Jalloul, and A. Duncan. Strategic default in financial networks. *Games and Economic Behavior*, 142:941–954, 2023.

[3] J. Andria, G. di Tollo, and J. Kalda. Propagation of bankruptcy risk over scale-free economic networks. *Entropy*, 24(12):1713, 2022.

[4] Bank of England. 2023 list of UK firms designated as other systemically important institutions (O-SIIs). https://www.bankofengland.co.uk/prudential-regulation/publication/2023/november/2023-list-of-uk-firms-designated-as-osiis, 2023.

[5] B. Bollobás, C. Borgs, J. T. Chayes, and O. Riordan. Directed scale-free graphs. In *SODA*, volume 3, pages 132–139, 2003.

[6] L. Eisenberg and T. H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.

[7] J. C. González-Avella, V. H. de Quadros, and J. R. Iglesias. Network topology and interbank credit risk. *Chaos, Solitons & Fractals*, 88:235–243, 2016.

[8] P. Kanellopoulos, M. Kyropoulou, and H. Zhou. Forgiving debt in financial network games. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 335–341. ijcai.org, 2022.

[9] P. Kanellopoulos, M. Kyropoulou, and H. Zhou. Debt transfers in financial networks: Complexity and equilibria. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 260–268, 2023.

[10] King's College London. King's computational research, engineering and technology environment (create). https://doi.org/10.18742/rnvf-m076, 2024.

[11] M. Lanctot, E. Lockhart, J.-B. Lespiau, V. Zambaldi, S. Upadhyay, J. Pérolat, S. Srinivasan, F. Timbers, K. Tuyls, S. Omidshafiei, et al. Openspiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.

[12] S. Langfield and K. Soramäki. Interbank exposure networks. *Computational Economics*, 47:3–17, 2016.

[13] K. Mayo and M. P. Wellman. A strategic analysis of portfolio compression. In *Proceedings of the Second ACM International Conference on AI in Finance*, pages 1–8, 2021.

[14] M. Montagna and T. Lux. Contagion risk in the interbank market: A probabilistic approach to cope with incomplete structural information. *Quantitative Finance*, 17(1):101–120, 2017.

[15] S. Omidshafiei, C. Papadimitriou, G. Piliouras, K. Tuyls, M. Rowland, J.-B. Lespiau, W. M. Czarnecki, M. Lanctot, J. Perolat, and R. Munos. α-rank: Multi-agent evaluation by evolution. *Scientific reports*, 9(1):9937, 2019.

[16] P. A. Papp and R. Wattenhofer. Network-aware strategies in financial systems. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 91:1–91:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[17] P. A. Papp and R. Wattenhofer. Debt swapping for risk mitigation in financial networks. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 765–784, 2021.

[18] M. Ponsen, K. Tuyls, M. Kaisers, and J. Ramon. An evolutionary game-theoretic analysis of poker strategies. *Entertainment Computing*, 1(1):39–45, 2009.

[19] L. C. Rogers and L. A. Veraart. Failure and rescue in an interbank network. *Management Science*, 59(4):882–898, 2013.

[20] S. Schuldenzucker and S. Seuken. Portfolio compression in financial networks: Incentives and systemic risk. In *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, page 79. ACM, 2020.

[21] Statista Research Department. Total number of Monetary Financial Institutions (MFI) operating in the United Kingdom (UK) as of March 2023, by country of origin. https://www.statista.com/statistics/870166/number-of-banks-operating-in-the-uk-by-country-of-residence/, 2023.

[22] K. Tuyls, J. Pérolat, M. Lanctot, J. Z. Leibo, and T. Graepel. A generalised method for empirical game theoretic analysis. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, pages 77–85. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2018.

[23] K. Tuyls, J. Perolat, M. Lanctot, E. Hughes, R. Everett, J. Z. Leibo, C. Szepesvári, and T. Graepel. Bounds and dynamics for empirical game theoretic analysis. *Autonomous Agents and Multi-Agent Systems*, 34:1–30, 2020.

[24] W. E. Walsh, D. C. Parkes, and R. Das. Choosing samples to compute heuristic-strategy nash equilibrium. In *Agent-Mediated Electronic Commerce V. Designing Mechanisms and Systems: AAMAS 2003 Workshop, AMEC 2003, Melbourne, Australia, July 15, 2003, Revised Selected Papers 5*, pages 109–123. Springer, 2004.

[25] M. P. Wellman. Methods for empirical game-theoretic analysis. In *AAAI*, volume 980, pages 1552–1556, 2006.

[26] M. P. Wellman. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30:1175–1189, 2016.