

Strategic Formation of Credit Networks

Pranav Dandekar
Stanford University
ppd@stanford.edu

Ashish Goel
Stanford University
ashishg@stanford.edu

Michael P. Wellman
University of Michigan
wellman@umich.edu

Bryce Wiedenbeck
University of Michigan
btwied@umich.edu

ABSTRACT

Credit networks are an abstraction for modeling trust between agents in a network. Agents who do not directly trust each other can transact through exchange of IOUs (obligations) along a chain of trust in the network. Credit networks are robust to intrusion, can enable transactions between strangers in exchange economies, and have the liquidity to support a high rate of transactions. We study the formation of such networks when agents strategically decide how much credit to extend each other. When each agent trusts a fixed set of other agents, and transacts directly only with those it trusts, the formation game is a potential game and all Nash equilibria are social optima. Moreover, the Nash equilibria of this game are equivalent in a very strong sense: the sequences of transactions that can be supported from each equilibrium credit network are identical. When we allow transactions over longer paths, the game may not admit a Nash equilibrium, and even when it does, the price of anarchy may be unbounded. Hence, we study two special cases. First, when agents have a shared belief about the trustworthiness of each agent, the networks formed in equilibrium have a star-like structure. Though the price of anarchy is unbounded, myopic best response quickly converges to a social optimum. Similar star-like structures are found in equilibria of heuristic strategies found via simulation. In addition, we simulate a second case where agents may have varying information about each others' trustworthiness based on their distance in a social network. Empirical game analysis of these scenarios suggests that star structures arise only when defaults are relatively rare, and otherwise, credit tends to be issued over short social distances conforming to the locality of information.

Categories and Subject Descriptors

K.4.4 [Computing Milieux]: Electronic Commerce; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Theory, Economics

Keywords

trust, credit networks, strategic network formation, empirical game-theoretic analysis

1. INTRODUCTION

The study of *strategic network formation* seeks to understand the emergent behavior and properties of a network when self-interested agents establish connections to other agents based on their local information. In general, establishing a connection incurs a cost but also yields some benefit to agents connected through that edge. The agents are deemed to be utility-maximizing, that is, they make decisions in order to maximize the difference between their total benefit and their total cost. This problem has been studied in many different settings [11, 2, 8, 5, 1]. One can ask interesting questions about the emergent properties of the networks formed in each setting: What network topologies are feasible in equilibrium? How do equilibrium networks differ from socially optimal ones? How does this depend upon the cost of forming an edge and the benefit derived from having a connection? If there are multiple equilibria, can agents select among them through some kind of iterated better-response dynamics?

This paper is an investigation into some of these questions in the context of *credit networks*, an abstraction for modeling trust among autonomous agents. A credit network represents trust relationships through a directed graph with edge capacities. Nodes in this graph correspond to agents, and edges correspond to credit relationships between them. An edge of capacity c from node u to node v indicates that agent u extends c units of credit to agent v , or equivalently, u is committed to accept IOUs (obligations) issued by v up to value c . The capacity of this edge can be viewed as a measure of u 's trust in v . Nodes pay for goods and services by issuing their own IOUs, instead of using a common currency. Credit commitments between trusting nodes also enable remote transactions, as illustrated in Fig. 1. Say node w wants to buy a good worth p units from node u . Nodes u and w can transact—even though u does not directly trust w —via the trusted intermediary v . Assuming $p \leq \min\{c_1, c_2\}$, the payment proceeds by w issuing an IOU to v worth p units, and v issuing an IOU to u worth p units. If, however, $p > \min\{c_1, c_2\}$, the transaction fails. As a result of a successful transaction, the credit capacities c_{uv} and c_{vw} decrease by p , representing the remaining credit commitments. In addition, the capacities c_{vu} and c_{wv} both increase to p from zero, since v and w will both accept the return of their own IOUs as payment. Thus arbitrary payments can be routed through a credit network by passing IOUs along a chain of trusting agents, obviating the need for a common currency. Observe that routing payments in credit networks is identical to routing residual flows in general flow networks. Also note that payment flows in the opposite direction of credit, so a payment merely results in a redistribution of credit: buyers expend credit and sellers gain it while intermediaries exchange credit between their neighbors, but the total credit in the network remains unchanged.




Figure 1: Updating credit to process a transaction between u and w worth p units.

The credit network model was introduced independently by De-Figueiredo and Barr [7], Ghosh et al. [9], and Karlan et al. [13] as a mechanism for enabling transactions among untrusting agents in a network. This model is particularly well-suited for transactions in exchange economies such as P2P networks where it has been shown to improve inefficiencies resulting from asynchronous demand and bilateral trading [14]. It has been used to thwart spam in e-mail and content-sharing systems such as YouTube [16]. It can also be used in settings such as packet routing in mobile ad-hoc networks and combating spam in viral marketing over social networks. There is a large body of work in economics and sociology on social capital and favor exchanges in networks [10]. This model not only provides a rigorous way of keeping track of favors owed to and by each individual in a network, but also facilitates exchange of favors, via trusted intermediaries, between individuals who do not know each other directly.

Prior research shows that liquidity (ability to route payments) in many credit network topologies is comparable to that in equivalent centralized currency systems [6]. Whereas that work takes the credit network to be exogenously defined, here we address the question of how credit networks may be formed in the first place. We endogenize network formation by explicitly modeling the decision by each agent to issue credit to others. Issuing credit entails risk (a counterparty may violate the trust extended), as well as benefits (it increases the probability that profitable transactions may be completed). A natural question is whether agents who rationally weigh these risks and benefits will actually form viable credit networks. Network formation in the presence of risk was recently studied by Blume et al. [4] in a model motivated by financial contagion and epidemic diseases. In their setting, nodes derive utility only from direct edges, whereas risk is contagious (*i.e.*, failure of distant nodes is also a source of risk). Our model flips this: nodes derive benefit from transactions along direct as well as multi-hop paths, whereas only direct edges are sources of risk.

1.1 Our Setting

In our model, each agent has a *credit budget* representing the total credit that agent can extend others. Agents play a one-shot game where they determine how much credit to extend other agents, and then engage in repeated probabilistic transactions over the formed credit network. Agents derive utility from successful transactions. Extending credit to other agents increases transaction success probability, thus contributing to utility. On the other hand, when agent u extends credit to agent v , u risks a potential loss of utility resulting from violation of the trust it placed in v . Thus, an agent's net utility is its total utility from successful transactions minus the utility loss from extending credit to untrustworthy agents.

We analyze the formation of credit networks under various models of risk. We start with a model of *dichotomous risk*: agents are embedded in a *social network* represented by an *undirected* graph.

Agents trust their neighbors in the social network and may extend credit to them. However, they associate a very high loss of utility with extending credit to non-neighbors, and consequently, never extend credit to them. This setting captures situations illustrated by the following examples where directly transacting with a stranger may have grave consequences.

- During a disease epidemic within a human population, high-risk groups will limit their interactions to those who belong to similar social circles. Evidence of this has been found, for example, in the setting of HIV/AIDS [12, 3].
- Users trying to circumvent Internet censorship and evade network surveillance in repressive regimes make use of Internet proxies [15]. If caught, penalties may be severe. Thus, users rely on their friends and acquaintances to distribute proxy addresses.
- Members of covert organizations face the prospect of severe harm at the hands of the enemy if their identity is compromised. As a result, they may rely on longstanding relationships and assets built over time to conduct their business.

We also study a model of *global risk*, which represents the other extreme with respect to the dichotomous risk model. In this model, each node has a publicly known risk of default. This corresponds to situations involving small, densely interacting social groups, or where there are organizations such as credit-reporting agents that systematically gather and disseminate relevant risk information.

Finally, we study a model of *graded risk* that helps bridge the gap between global and dichotomous risk. Under this model, each agent has a private default probability. Agents receive noisy signals about each other's probability of defaulting, and these signals are more informative for neighbors in the social network.

1.2 Results

Dichotomous Risk Under dichotomous risk, when we allow only bilateral transactions (*i.e.*, transactions only between adjacent nodes in the social network, and payments routed only along the direct edge between nodes), we show that the formation game is a potential game (Theorem 3.1). This implies that best-response dynamics always converge to a Nash equilibrium¹. Moreover, for a large, natural class of transaction size distributions, we show that agents' utilities are concave in their credit allocations. This allows us to prove that every Nash equilibrium of the game maximizes social welfare (Theorem 3.4). More interestingly, we show that the Nash equilibria are equivalent in a much stronger sense: any two Nash equilibria are *cycle-reachable* from each other (Theorem 3.6), which means that it is possible to transform one equilibrium into another by routing a sequence of payments from a node to itself along a feasible path. The significance of this structural property follows from [6]: for any two Nash equilibria s and s' of the game, if an arbitrary sequence of transactions is feasible starting from s , that sequence is also feasible starting from s' .

With non-bilateral transactions, the game becomes significantly less well-behaved: the game may not admit a Nash equilibrium (Theorem 3.8), and even when it does, the price of anarchy in this setting can be unbounded (Theorem 3.9).

Global Risk Under global risk, we analyze the price of anarchy and the structure of equilibria when each agent is limited to extend credit to at most one other agent. We prove if we disallow the empty network as an outcome, the price of anarchy of the formation game

¹In this paper, the term *Nash Equilibrium* always refers to a pure Nash equilibrium, except when we explicitly consider mixed strategy equilibria of simulated games in Section 4.2.

is unbounded (Theorem 4.4), even though all Nash equilibria have a star-like structure (Theorem 4.3). Instead we focus on the structure of equilibria under two simple dynamics: sequential arrival and myopic best response. When nodes arrive sequentially and create a single link, we show that a node u always extends credit to either the node v that arrived immediately before u or to the node that v extends credit to (Theorem 4.6). Thus the resulting network has a *comb-like* structure. Under myopic best response, nodes extend their entire credit budget to the node that has the lowest risk of default. If the default risks are unique, this results in a *star-like* network structure which is also the optimal structure in terms of social welfare (Theorem 4.5). Thus, even though the price of anarchy can be unbounded, nodes can easily find the optimal network using myopic best response.

Simulations We use empirical game simulation to study a more general formulation of the global risk model, finding that non-empty equilibrium networks tend to have a centralized, star-like structure due to use of default probability as a primary credit-issuing criterion. We also analyze several graded risk settings, and find that centralized networks only arise when defaults are relatively rare, and otherwise, credit links tend to be issued over short social distances conforming to the locality of information.

2. MODEL AND DEFINITIONS

Let V denote the set of n agents. Each agent $u \in V$ has a budget $B_u \geq 0$ representing the total credit that u can extend to other agents in V . Agents play a one-shot game where they choose credit allocations to form an initial network s . Agents represent nodes of the formed network. An edge from node u to node v of capacity $c_{uv}(s)$ represents the credit extended by agent u to agent v in the network s . A strategy for agent u is a set of *feasible* credit allocations $\{c_{uv}(s), v \in V : c_{uv}(s) \geq 0 \text{ and } \sum_{v \in V} c_{uv}(s) \leq B_u\}$.

2.1 Transaction Model

Once a network s is formed, agents engage in repeated probabilistic transactions with each other. At each time step $t = 1, 2, \dots$, a pair of transacting agents $\langle u, v \rangle$, with u being the payer (buyer) and v the payee (seller), is chosen with probability λ_{uv} . The transaction rate matrix $\Lambda = \{\lambda_{uv} : u, v \in V\}$ is public, and satisfies the following properties: (i) $\lambda_{uu} = 0$, (ii) $\lambda_{uv} \geq 0$, and (iii) $\sum_{u,v} \lambda_{uv} = 1$.

Suppose agents $\langle u, v \rangle$ are chosen to transact at time t . Then the transaction size, x_{uv}^t , between u and v is drawn from a *transaction size distribution* over $[0, \infty)$ with a probability density function (pdf) $g_{uv}(\cdot)$ and a corresponding cumulative distribution function (cdf) $G_{uv}(\cdot)$. We assume that the pdfs $g_{uv}(\cdot)$ are public. Let $\mathcal{G} := \{g_{uv}(\cdot) : u, v \in V\}$ be the pdf matrix.

Given a transaction size x , a *feasible path* in the network s from node v to node u is a set of directed edges $\mathcal{P} = \{(v, u_1), (u_1, u_2), \dots, (u_{k-1}, u_k), (u_k, u)\}$ such that for all $(w, y) \in \mathcal{P}$, $c_{wy}(s) \geq x$. We route payments along the shortest feasible path in the network. Let \mathcal{P}_{vu}^t be the shortest feasible path in the credit network from v to u at time t . A successful transaction of size x_{uv}^t results in a change of credit capacities along edges in \mathcal{P}_{vu}^t as follows. Let $s^t := \{c_{uv}(s^t) : u, v \in V\}$ denote the state of the network s at time $t = 0, 1, 2, \dots$, where $s^0 = s$. Then, for $w, y \in V$ and for $t > 0$,

$$c_{wy}(s^t) = \begin{cases} c_{wy}(s^{t-1}) - x_{uv}^t, & \text{if } (w, y) \in \mathcal{P}_{vu}^{t-1} \\ c_{wy}(s^{t-1}) + x_{uv}^t, & \text{if } (y, w) \in \mathcal{P}_{vu}^{t-1} \\ c_{wy}(s^{t-1}), & \text{otherwise} \end{cases}$$

So, in order for a payment x_{uv}^t from u to v to succeed, there must exist a feasible path in the credit network *from the payee v to the payer u* . If no such path exists, the transaction fails, in which case all credit capacities remain unchanged. Thus, for all $t > 0$, and for all $u, v \in V$, $c_{uv}(s^t) + c_{vu}(s^t) = c_{uv}(s) + c_{vu}(s)$.

The repeated probabilistic transactions induce a Markov chain over the states of the network, which we denote by $\mathcal{M}(s, \Lambda, \mathcal{G})$. A *transaction regime* is defined as the tuple $\langle \Lambda, \mathcal{G} \rangle$. We say a transaction regime $\langle \Lambda, \mathcal{G} \rangle$ is *symmetric* if the transaction rate matrix Λ is symmetric: for all nodes $u, v \in V$, $\lambda_{uv} = \lambda_{vu}$, and the transaction size pdfs are symmetric: for all $u, v \in V$, $g_{uv}(\cdot) = g_{vu}(\cdot)$.

We are interested in the success probabilities of transactions in the steady-state of this Markov chain, which are difficult to characterize for arbitrary networks and transaction regimes. However, we can do so in some simple cases, including the *unit transaction regime*.

DEFINITION 2.1. A unit transaction regime over credit network s is a transaction regime $\langle \Lambda, \mathcal{G} \rangle$ where, for all $u, v \in V$ and for all $t > 0$, the transaction size $x_{uv}^t = 1$, the transaction rate matrix Λ is symmetric and the Markov chain $\mathcal{M}(s, \Lambda, \mathcal{G})$ is ergodic.

When the network s is acyclic (ignoring directionality), Dandekar et al. [6] characterize the steady-state success probabilities under a unit transaction regime.

LEMMA 2.1 ([6]). Consider a credit network s . Assume that s is acyclic if we ignore the directions of the edges in s . Let \mathcal{P}_{uv} be the set of (undirected) edges along the path between nodes u and v . Then, in a unit transaction regime over s , the steady-state transaction success probability, $f_{uv}(s)$, between two nodes $u, v \in V$ is given by

$$f_{uv}(s) = \lambda_{uv} \prod_{e=(w,y) \in \mathcal{P}_{uv}} \frac{\lfloor c_{wy}(s) \rfloor + \lfloor c_{yw}(s) \rfloor}{\lfloor c_{wy}(s) \rfloor + \lfloor c_{yw}(s) \rfloor + 1}$$

2.2 Utility

Agents choose credit allocations to maximize their utility. Successful transactions contribute to agents' utility, but agents risk loss of utility when they extend credit to potentially untrustworthy agents. We model this risk in several ways, but denote the expected loss of utility to u associated with the prospect of default by v by $\Delta_{uv}(s)$, with the constraints that $\Delta_{uv}(s) \geq 0$ and $\Delta_{uv}(s) > 0$ only if $c_{uv}(s) > 0$. Let $f_{uv}(s)$ be the steady-state success probability of the transactions from u to v when the initial network is s . Then, the total utility of an agent u when the initial network is s is given by

$$U_u(s) = \gamma \sum_{w \in V} f_{uw}(s) - \sum_{v \in V : c_{uv}(s) > 0} \Delta_{uv}(s) \quad (1)$$

where γ is a constant that converts transaction success probability into equivalent utility units. The overall *social welfare* in network s is simply the sum of utilities of all nodes in s : $U(s) = \sum_{u \in V} U_u(s)$.

2.3 Risk Model

In order to model variation in $\Delta_{uv}(s)$, we assume that the agents are embedded in an exogenously-defined *social network* represented by a simple undirected graph $H = (V, E)$. The social network H influences the how $\Delta_{uv}(s)$ for an agent u varies across agents $v \in V$. We consider three specific models of how risk changes as a function of distance between u and v in H .

Dichotomous Risk. In this model, an agent u partitions the set of agents V into two sets using H : neighbors in H and non-neighbors in H . For any network s , agent u estimates risk exposure

to be:

$$\Delta_{uv}(s) = \begin{cases} 0, & \text{if } (u, v) \in E \\ \infty, & \text{otherwise} \end{cases} \quad (2)$$

This model assumes agents are willing to interact only with their neighbors in H . For any credit network s formed under this model, $c_{uv}(s) = 0$ if $(u, v) \notin E$.

Global Risk. In this model, we assume that each agent v has a default probability $\delta_v \in (0, 1]$ which is public. If v defaults, a node u that extended credit $c_{uv}(s)$ to v loses $c_{uv}(s)$ units. Thus, $\Delta_{uv}(s) = \delta_v c_{uv}(s)$.

Graded Risk. Here, as in the Global Risk model, each agent v has default probability δ_v , but this information is not publicly known. Instead, each agent u receives a signal δ_{uv} about the default probability of each other agent v . These signals are decreasingly informative with distance in H , so agents know much more about the default probabilities of their neighbors in the social network than about distant nodes. In our simulations, we implement this by drawing agents' default probabilities from a beta distribution: $\delta_v \sim \text{Beta}(\alpha, \beta)$. Agent u then receives a signal in the form of some number of samples S_{uv} drawn from the binomial distribution on δ_v , where S_{uv} decreases exponentially with social network distance.

3. NETWORK FORMATION UNDER DICHOTOMOUS RISK

Recall that under dichotomous risk, $\Delta_{uv}(s)$ is defined by (2), as a result nodes only extend credit to their neighbors in H .

3.1 Symmetric Bilateral Transactions

We call a transaction between nodes u and v *bilateral* if $(u, v) \in E$ and the payment is routed along the edge (u, v) . Here we allow only bilateral transactions: if a payment between adjacent nodes u and v cannot be routed along the direct edge (u, v) , we fail the transaction. As a result, if $(u, v) \notin E$, the steady-state success probability $f_{uv}(s) = f_{vu}(s) = 0$. Moreover, the steady-state transaction success probabilities along an edge $e = (u, v)$ in a network s are governed only by the credit allocations, $c_{uv}(s)$, $c_{vu}(s)$, along e in s . We also assume that the transaction regime $\langle \Lambda, \mathcal{G} \rangle$ is symmetric and that $\lambda_{uv} > 0$ if $(u, v) \in E$. As a result, for all nodes u and v , $f_{uv}(s) = f_{vu}(s)$.

In our analysis of the symmetric bilateral transaction regime, for an edge $e = (u, v) \in E$, we will use $\lambda_e, g_e(\cdot), G_e(\cdot)$ and $f_e(\cdot)$ to denote $\lambda_{uv}, g_{uv}(\cdot), G_{uv}(\cdot)$, and $f_{uv}(\cdot)$, respectively.

We first show that in this setting, the network formation game is a potential game.

THEOREM 3.1. *The network formation game under a symmetric bilateral transaction regime is a potential game.*

PROOF. Consider the function $\Phi(s)$ defined as

$$\Phi(s) := \frac{U(s)}{2} = \frac{1}{2} \sum_{u \in V} U_u(s) = \frac{\gamma}{2} \sum_{u \in V} \sum_{v \in V} f_{uv}(s)$$

Since we are in a symmetric bilateral transaction regime, $f_{uv}(s) = f_{vu}(s)$ for all $(u, v) \in E$, and $f_{uv}(s) = 0$ if $(u, v) \notin E$. Therefore,

$$\sum_{u \in V} \sum_{v \in V} f_{uv}(s) = 2 \sum_{e \in E} f_e(s)$$

This implies $\Phi(s) = \gamma \sum_{e \in E} f_e(s)$. We will show that $\Phi(s)$ is a potential function. Fix a node $u \in V$. Consider a network s' which differs from s only in the credit allocation of u . Formally, for all

$w, y \in V$,

$$c_{wy}(s') = \begin{cases} c_{wy}(s), & \text{if } w \neq u \\ c'_{wy}, & \text{if } w = u \text{ and } (u, y) \in E \end{cases}$$

where $\{c'_{uy} : (u, y) \in E\}$ is any feasible allocation of u 's credit. Let $E_u \subseteq E$ be the set of edges incident upon u in E . Note that for all $e' = (u', v') \notin E_u$, $c_{u'v'}(s) = c_{u'v'}(s')$. As a result, $f_{e'}(s) = f_{e'}(s')$. It follows that

$$\Phi(s) - \Phi(s') = \gamma \sum_{e \in E_u} (f_e(s) - f_e(s')) = U_u(s) - U_u(s')$$

Thus the network formation game is a potential game with $\Phi(s)$ as the potential function. \square

Theorem 3.1 implies that in this setting, a Nash equilibrium always exists, best-response dynamics always converge to a Nash equilibrium, and finally, because the potential function is given by $\Phi(s) = U(s)/2$, the price of stability is 1. Next we will show that for a large, natural class of transaction size distributions, agents' utilities are concave, and consequently, the price of anarchy is also 1, i.e., every Nash equilibrium of the formation game maximizes social welfare.

3.1.1 Nash Equilibria Maximize Social Welfare

Consider an edge $e = (u, v) \in E$. Assume that $g_e(\cdot)$ has support over $[0, \infty)$. Also, let $G_e(\cdot)$ be twice differentiable. First we derive an expression for $f_e(s)$ in terms of the credit allocations $c_{uv}(s)$ and $c_{vu}(s)$ along edge e .

LEMMA 3.2. *Consider a credit network s . For nodes $u, v \in V$ such that $e = (u, v) \in E$, the steady-state transaction success probability, $f_e(s)$, under a symmetric bilateral transaction regime is given by*

$$f_e(s) = f_e(c_e(s)) = \begin{cases} \frac{\lambda_e}{c_e(s)} \int_0^{c_e(s)} G_e(y) dy, & \text{if } c_e(s) > 0 \\ 0, & \text{if } c_e(s) = 0 \end{cases} \quad (3)$$

where $c_e(s) = c_{uv}(s) + c_{vu}(s)$ is the total credit allocated along edge e in s .

The proof is omitted due to space constraints.² Observe from (3) that $f_e(s)$ depends only on the total credit capacity $c_e(s)$ along the edge $e = (u, v)$. Therefore, for the rest of this section, instead of thinking of f_e as a function of $c_{uv}(s)$ and $c_{vu}(s)$, we will think of f_e as the function $f_e : \mathbb{R}_+ \rightarrow [0, 1]$. That is, $f_e(x)$ is the steady-state transaction success probability along edge e when the total credit allocated along it is x . We will write $f_e(s)$ to mean $f_e(c_e(s))$ when there is no ambiguity. Next we prove some properties of the functions $f_e(\cdot)$ that enable us to establish that every Nash equilibrium maximizes social welfare.

LEMMA 3.3. *Consider a credit network s under a symmetric bilateral transaction regime. For an edge $e \in E$,*

1. *The transaction success probability, $f_e(\cdot)$, is continuously differentiable and strictly increasing.*
2. *If $g_e(\cdot)$ is non-increasing, $f_e(\cdot)$ is concave.*

As a corollary, if $g_e(\cdot)$ is strictly decreasing, $f_e(\cdot)$ is strictly concave. Many natural distributions have strictly decreasing density

²Proofs of all results are included in the full version of this paper, available at <http://www.stanford.edu/~ppd/papers/cn-formation.pdf>.

functions over $[0, \infty)$. Examples include the exponential distribution, the normal distribution $\mathcal{N}(0, \sigma^2)$, and the power-law distribution. Next we show that if the transaction success probabilities, $f_e(\cdot)$, are concave, every Nash equilibrium maximizes social welfare.

THEOREM 3.4. *Let s be a Nash equilibrium of the network formation game under a symmetric bilateral transaction regime. If the transaction success probabilities, $f_e(\cdot), e \in E$, are concave, then s maximizes social welfare $U(s)$.*

PROOF. Recall from Theorem 3.1 that the formation game under a symmetric bilateral transaction regime is a potential game and $\Phi(s) = U(s)/2 = \gamma \sum_{e \in E} f_e(s)$ is a potential function. Recall from Lemma 3.3, that $f_e(\cdot), e \in E$, are continuously differentiable, which implies $\Phi(\cdot)$ is continuously differentiable. Since $f_e(\cdot)$ are concave (by assumption), $\Phi(\cdot)$ is also concave. It was shown by Neyman [17] that any Nash equilibrium of a potential game with a concave and continuously differentiable potential is also a potential maximizer. Therefore, s maximizes $\Phi(s)$, or equivalently, $U(s)$. \square

3.1.2 Nash Equilibria are Cycle-Reachable

Theorem 3.4 implies an equivalence between the Nash equilibria of the game; any two Nash equilibria s and s' have the same social welfare, $U(s) = U(s')$. Next we show that if $f_e(\cdot), e \in E$, are strictly concave, the Nash equilibria of this game are equivalent in a much stronger sense: any two Nash equilibria s and s' are *cycle-reachable*, which, as shown by Dandekar et al. [6], implies that the sequences of transactions that succeed starting from s and starting from s' are identical.

We first show that the total credit capacity of any edge in E is identical in any Nash equilibrium.

LEMMA 3.5. *Let $f_e(\cdot), e \in E$, be strictly concave. Let s and s' be two Nash equilibria of the network formation game. Then for all edges $e \in E, c_e(s) = c_e(s')$.*

PROOF. First, let us define the marginal utility of an edge $e \in E$.

DEFINITION 3.1. *The marginal utility of an edge $e \in E$ is the function $r_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ given by*

$$r_e(x) = f'_e(x) = \frac{df_e(x)}{dx}$$

We show that for any edge $e \in E, r_e(s) = r_e(s')$. The lemma follows as a direct consequence.

Since $f_e(\cdot)$ is strictly concave (by assumption), strictly increasing and continuously differentiable (by Lemma 3.3), $r_e(\cdot)$ is continuous, strictly decreasing and strictly positive. In network s , the marginal utility on an edge $e \in E$ is given by $r_e(c_e(s))$. We denote it by $r_e(s)$ when there is no ambiguity. Let E_u be the set of edges in E incident upon node u .

DEFINITION 3.2. *For a node $u \in V$ and a network s , we define $\rho_u(s) := \max_{e \in E_u} r_e(s)$ and $E_u^*(s) \subseteq E_u$ as the set of edges $e \in E_u$ such that $r_e(s) = \rho_u(s)$.*

In words, $E_u^*(s)$ is the set of edges incident on node u that have the highest marginal utility in network s among all edges in E_u . We show that in any Nash equilibrium s , each node u exhausts its entire budget and allocates non-zero credit only along edges in $E_u^*(s)$.

PROPOSITION 1. *Let s be a Nash equilibrium. Then, for all nodes $u \in V$, both (1) and (2) are true:*

1. $\sum_{v:(u,v) \in E} c_{uv}(s) = B_u$.

2. For each $e = (u, v) \in E$, if $e \notin E_u^*(s)$ then $c_{uv}(s) = 0$.

Next we define a *slack edge*.

DEFINITION 3.3. *Let s be a Nash equilibrium. We call an edge $e = (u, v) \in E$ a slack edge in s if $e \notin E_u^*(s)$ or $e \notin E_v^*(s)$ or both.*

Note that by Proposition 1, if edge $e = (u, v)$ is a slack edge in Nash equilibrium s , either $c_{uv}(s) = 0$ or $c_{vu}(s) = 0$ or both $c_{uv}(s) = c_{vu}(s) = 0$.

DEFINITION 3.4. *Let s be a credit network. We define*

1. $r_s^{\min} := \min_{e \in E} r_e(s)$ to be the minimum marginal utility of any edge $e \in E$ in s ,
2. the set $E_s^{\min} := \{e \in E \mid r_e(s) = r_s^{\min}\}$,
3. the set $V_s^{\min} := \{u \in V \mid u \text{ is incident on some edge in } E_s^{\min}\}$, and
4. the set $V_s^X \subseteq V_s^{\min}$ as

$$V_s^X := \{u \in V \mid u \text{ is incident upon some edge in } E_s^{\min} \text{ and upon some edge in } E - E_s^{\min}\}$$

The minimum marginal utility in any two Nash equilibria is identical.

PROPOSITION 2. *Let s and s' be two Nash equilibria. Then $r_s^{\min} = r_{s'}^{\min}$.*

Moreover, in any two Nash equilibria s and s' , the set of edges with the minimum marginal utility in s is identical to that in s' .

PROPOSITION 3. *Let s and s' be two Nash equilibria. Then $E_s^{\min} = E_{s'}^{\min}$.*

COROLLARY 3.1. *Let s and s' be two Nash equilibria. Then $V_s^{\min} = V_{s'}^{\min}$ and $V_s^X = V_{s'}^X$.*

Thus, we have established that for any two Nash equilibria s and s' , $r_e(s) = r_e(s')$ for all edges $e \in E_s^{\min}$. We show using an inductive argument that this is true of all edges in E .

DEFINITION 3.5. *Given an instance $I : G = (V, E); f_e, e \in E; B_u, u \in V$ of the network formation game under a symmetric bilateral transaction regime, a network s , and an arbitrary set of edges $F \subseteq E$, we define the (s, F) -restriction of I , denoted $I_{(s, F)}$, as follows: $G^{(F)} := (V, E \setminus F)$, $f_e^{(F)} := f_e, e \in E \setminus F$, and*

$$B_u^{(s, F)} := \begin{cases} 0 & \text{if } E_u \subseteq F \\ B_u - \sum_{(u, w) \in F} c_{uw}(s) & \text{otherwise} \end{cases}$$

Note that for a node u , if $E_u \subseteq F$, then the value of $B_u^{(s, F)}$ is immaterial since u has no incident edges in $I_{(s, F)}$ along which to allocate its budget.

DEFINITION 3.6. *Given a network s and an arbitrary set of edges $F \subseteq E$, we define an F -restriction of s , denoted, $s_{(F)}$, as follows: for all edges $e = (u, v) \in E \setminus F$, $c_{uv}(s_{(F)}) = c_{uv}(s)$ and $c_{vu}(s_{(F)}) = c_{vu}(s)$.*

PROPOSITION 4. *If s is a Nash equilibrium for instance I of the network formation game in the bilateral transaction setting, then $s_{(F)}$ is a Nash equilibrium for $I_{(s, F)}$ for any set $F \subseteq E$.*

PROPOSITION 5. Let s and s' be two Nash equilibria for instance 1 of the network formation game under a symmetric bilateral transaction regime. Then for all edges $e \in E$, $r_e(s) = r_e(s')$.

Observe that since $f_e(\cdot)$ is strictly concave, $r_e(\cdot)$ is strictly decreasing. Therefore, Proposition 5 implies that for all $e \in E$, $c_e(s) = c_e(s')$. \square

Lemma 3.5 allows us to show that any two Nash equilibria are cycle-reachable.

DEFINITION 3.7 ([6]). Let s and s' be two credit networks. We say that s' is cycle-reachable from s if s can be transformed into s' by routing a sequence of payments along feasible cycles (i.e., from a node to itself along a feasible path).

THEOREM 3.6. Let $f_e(\cdot)$, $e \in E$, be strictly concave. Let s and s' be two Nash equilibria of the network formation game under the symmetric bilateral transaction regime. Then s and s' are cycle-reachable from each other.

PROOF. First we define the generalized score vector of a credit network s .

DEFINITION 3.8 ([6]). Given a credit network s of n nodes, the generalized score vector of s is the vector $D(s) = \langle d_u(s) : u \in V \rangle \in \mathbb{R}_+^n$ where for all $u \in V$, $d_u(s) := \sum_{v \in V} c_{vu}(s)$.

Next we show that any two Nash equilibria have the same generalized score vector.

PROPOSITION 6. Let s and s' be two Nash equilibria. Then, $D(s) = D(s')$.

PROOF. Fix a node $u \in V$. Recall from Proposition 1 that

$$\sum_{v:(u,v) \in E} c_{uv}(s) = \sum_{v:(u,v) \in E} c_{uv}(s') = B_u \quad (4)$$

Also, from Lemma 3.5, we know that for all edges $e \in E$,

$$c_e(s) = c_e(s') \quad (5)$$

Let E_u be the set of edges in E incident upon u . It follows from (4) and (5) that

$$\begin{aligned} d_u(s) &= \sum_{v \in V} c_{vu}(s) = \sum_{v:(u,v) \in E} c_{vu}(s) = \sum_{e \in E_u} (c_e(s) - c_{uv}(s)) \\ &= \sum_{e \in E_u} c_e(s) - B_u = \sum_{e \in E_u} c_e(s') - \sum_{e \in E_u} c_{uv}(s') = d_u(s') \end{aligned}$$

\square

PROPOSITION 7 ([6]). Two credit networks s and s' are cycle-reachable if and only if $D(s) = D(s')$.

Proposition 6 along with Proposition 7 complete the proof. \square

The significance of this result is that if s and s' are cycle-reachable, they support the same set of feasible transactions.

THEOREM 3.7 ([6]). Let s_1 and s_2 be two cycle-reachable networks. If a transaction $\theta = \langle u, v, p \rangle$ (i.e., routing p units from node u to node v) is feasible in s_1 , it is also feasible in s_2 . Further, if transaction θ in network s_1 results in a network s'_1 , and the same transaction in network s_2 results in a network s'_2 , then s'_1 and s'_2 are cycle-reachable.

Thus, for two Nash equilibria s and s' of the game, if a sequence of transactions succeeds starting from s it also succeeds starting from s' . Observe that this equivalence between Nash equilibria implied by Theorem 3.6 is stronger than that implied by Theorem 3.4.

3.2 Symmetric Transactions

Here we lift the restriction that transactions be bilateral, allowing transactions between nodes that are not neighbors in H . We also allow payments between neighboring nodes to be routed along paths other than the direct edge between them.




Figure 2: Example of a formation game that does not admit a Nash equilibrium.

THEOREM 3.8. There exists an instance of the network formation game under a symmetric transaction regime that does not admit a Nash equilibrium.

PROOF. We will construct an instance of network formation game and show that it does not admit a Nash equilibrium. Consider a game with six agents: $V = \{a, b, d, e, h, j\}$. The graph H is a line graph over nodes in V with edges $(a, b), (b, d), (d, e)$ and so on. For each node $u \in V$, $B_u = 1$. The non-zero transaction rates are given by: $\lambda_{ab} = \lambda_{ba} = \lambda_{de} = \lambda_{ed} = \lambda_{hj} = \lambda_{jh} = 0.001$, $\lambda_{ae} = \lambda_{ea} = \lambda_{bj} = \lambda_{jb} = 0.2435$, $\lambda_{ej} = \lambda_{je} = 0.01$. All other entries in the transaction rate matrix Λ are zero. All transactions are of size one. Observe that this is a unit transaction regime, so we can use Lemma 2.1 to compute the steady-state transaction success probabilities between nodes.

Let s be a Nash equilibrium. Then, it must be that $c_{ab}(s) = c_{de}(s) = c_{hj}(s) = c_{jh}(s) = 1$. Let $c_{bd}(s) = x$ and $c_{ba}(s) = 1 - x$. Similarly, let $c_{eh}(s) = y$ and $c_{ed}(s) = 1 - y$. Observe that since all transactions are of size one, and s is a Nash equilibrium, it must be that $x, y \in \{0, 1\}$ (i.e., x and y cannot be strictly between 0 and 1). Verify that for each of the four combinations of (x, y) , namely, $(0, 0), (0, 1), (1, 0)$ and $(1, 1)$, either b or e has an improving unilateral deviation. In fact, the four combinations form a best-response cycle. Hence, there is no assignment of $x, y \in [0, 1]$ that will ensure that s is a Nash equilibrium. \square

Next we show that even if agents reach a Nash equilibrium, it may be arbitrarily bad in terms of social welfare compared to a social optimum.




Figure 3: Example of a game with an unbounded price of anarchy.

THEOREM 3.9. The price of anarchy of the network formation game under a symmetric transaction regime is unbounded.

PROOF. We will construct an instance of the game and show that it has an unbounded price of anarchy. Consider a game with four agents: $V = \{a, b, c, d\}$. The graph H is a line graph over

nodes in V with edges (a, b) , (b, c) and (c, d) . For each node $u \in V$, $B_u = 1$. The non-zero transaction rates are given by: $\lambda_{ab} = \lambda_{ba} = \lambda_{cd} = \lambda_{dc} = \lambda_1 > 0$, $\lambda_{ad} = \lambda_{da} = \lambda_2 \gg \lambda_1$. All other entries in the transaction rate matrix Λ are zero. All transactions are of size one.

Consider the network s shown in Fig. 3a. Observe that we can use Lemma 2.1 to compute the steady-state transaction success probabilities between nodes in s . Verify that s is a Nash equilibrium. The overall social welfare, $U(s)$, in network s is given by

$$\begin{aligned} U(s) &= \sum_{u \in V} U_u(s) = \sum_{u \in V} \sum_{v \in V} f_{uv}(s) = 2f_{ab}(s) + 2f_{cd}(s) \\ &= 2\lambda_1 \frac{2}{3} + 2\lambda_1 \frac{2}{3} = \lambda_1 \frac{8}{3} \end{aligned}$$

Now consider the network s^* in Fig. 3b. Verify that s^* is a social optimum. The overall social welfare $U(s^*)$ is given by

$$U(s^*) = \sum_{u \in V} U_u(s^*) = 2 \left(\lambda_1 \frac{2}{3} + \lambda_1 \frac{1}{2} + \lambda_2 \frac{1}{6} \right)$$

As $\lambda_1 \rightarrow 0$, the ratio $U(s^*)/U(s) \rightarrow \infty$. \square

4. NETWORK FORMATION UNDER GLOBAL RISK

Recall that in the global risk model, each agent v has a public default probability $\delta_v \in (0, 1]$. If v defaults, a node u that extended credit $c_{uv}(s)$ to v loses $c_{uv}(s)$ units. Thus, $\Delta_{uv}(s) = \delta_v c_{uv}(s)$.

4.1 Single-Minded Agents

We analyze the setting where agents may issue credit to at most one counterpart.

DEFINITION 4.1. We say that agent $u \in V$ is single-minded if in any credit network s , either $c_{uv}(s) = 0$ for all $v \in V$, or there exists a single agent $w \in V$ such that $c_{uw}(s) = B_u$.

Further, we assume that (i) the transaction rate matrix Λ is uniform: for all $u, v \in V$, $\lambda_{uv} = \lambda = 1/(n(n-1))$, (ii) all transactions are size one: for all $u, v \in V$, and for all $t > 0$, $x_{uv}^t = 1$, and (iii) for all agents $u \in V$, the credit budget $B_u = c > 0$, where c is an integer.

First we illustrate using a simple example that if the default probabilities are in a certain range, the empty network is a Nash equilibrium, and the price of anarchy is ∞ .

Example 1: Consider a set of n agents. Further suppose that for all $u, v \in V$, $\gamma\lambda(h+h^2) > \delta_{uc} > \gamma\lambda h$, where $h = c/(c+1)$. Let s be the empty network. Observe that, Lemma 2.1, the utility to a node u from extending to any node v in s is $\gamma\lambda h$, which by assumption is less than $\delta_v c$. Thus s is a Nash equilibrium. On the other hand, since $\gamma\lambda(h+h^2) > \delta_{uc}$ for all $u \in V$, the social optimum is a star network where every node extends credit to the root, while the root extends no credit. As a result, the price of anarchy is ∞ .

For the rest of this section, we assume that extending zero credit is not part of the agents' strategy set. This assumption, coupled with the fact that agents are single-minded, implies that any credit network formed in this setting will have exactly n directed edges each of capacity c , where n is the number of agents playing the game. Since an agent extends credit to exactly one agent in any network, we define the following notation to denote the agent that has been extended credit by an agent u in network s : for a network s , we define $\tau_s : V \rightarrow V$ to be the "trustee function": $\tau_s(u) = v$ implies $c_{uv}(s) = c$.

We use the following observation to prove our results; the observation follows from the analysis by Dandekar et al. [6] of the steady-state success probability in trees under a unit transaction regime.

LEMMA 4.1 ([6]). Consider a network s . Let $u \in V$ be a node such that no node extends credit to u in s and let $\tau_s(u) = v$. Assume the transaction rate matrix Λ is uniform and s is under a unit transaction regime. Then, for any node $w \in V \setminus \{u, v\}$, $f_{uw}(s) = hf_{vw}(s)$, where $h = c/(c+1)$.

4.1.1 Price of Anarchy and Structure of Equilibria

It is easy to see that any socially optimal network will have a star-like structure where the root is a node with the minimum default probability.

LEMMA 4.2. Let $v^* \in \arg \min_{v \in V} \delta_v$ be a node with the minimum default probability. Let $u^* \in \arg \min_{v \in V \setminus \{v^*\}} \delta_v$ be a node with the minimum default probability among nodes other than v^* . Consider a network s^* such that for all nodes $u \in V \setminus \{v^*\}$, $\tau_{s^*}(u) = v^*$, and $\tau_{s^*}(v^*) = u^*$. Then, s^* maximizes social welfare. Moreover, s^* is also a Nash equilibrium.

Next we show that all Nash equilibria have a star-like structure.

THEOREM 4.3. For a sufficiently large n , in any Nash equilibrium s there exists a node u^* such that for all nodes $v \in V \setminus \{u^*\}$, $\tau_s(v) = u^*$.

Next we show that despite ruling out the empty network as a Nash equilibrium, the price of anarchy in this setting can be unbounded.

THEOREM 4.4. The price of anarchy of the network formation game with single-minded agents is unbounded.

PROOF. Consider a set of n agents. Assume, without loss of generality, that for nodes $u_1, \dots, u_n \in V$, $\delta_{u_1} \leq \dots \leq \delta_{u_n}$. Let $\delta_{u_1} c = \gamma\lambda(n-3)h^2 2c/(2c+1)$, and $\delta_{u_2} = \delta_{u_3} = \gamma\lambda(n-3)h^2$, where recall that $h = c/(c+1)$. Consider the network s^* in Fig. 4a. It follows from Lemma 4.2 that s^* is a socially optimal network. Consider the network s_1 in Fig. 4b. Observe that Lemma 2.1 can be used to compute the steady-state transaction success probabilities and, hence, the utilities, of all nodes in s_1 . Since $c(\delta_{u_3} - \delta_{u_1}) \leq (n-3)\gamma\lambda \frac{h^2}{2c+1}$, nodes in s_1 cannot benefit from extending credit to u_1 or u_2 instead of u_3 . Thus, s_1 is a Nash equilibrium. Note that since s^* and s_1 are structurally identical

$$\begin{aligned} \sum_{u,v} f_{uv}(s^*) &= \sum_{u,v} f_{uv}(s_1) \\ &= \lambda(n-2) \left((n-3)h^2 + 2h \frac{2c}{2c+1} + 2h \right) + 2\lambda \frac{2c}{2c+1} \\ &= \lambda(n-2)(n-3)h^2 + \Theta(n) \end{aligned}$$

Thus, the total social welfare in s^* is given by

$$\begin{aligned} U(s^*) &= \gamma \sum_{u,v} f_{uv}(s^*) - (n-1)\delta_{u_1} c - \delta_{u_2} c \\ &= \gamma\lambda(n-3)h^2 \left((n-2) - (n-1) \frac{2c}{2c+1} \right) + \Theta(n) = \Theta(n^2) \end{aligned}$$

On the other hand,

$$\begin{aligned} U(s_1) &= \gamma \sum_{u,v} f_{uv}(s_1) - (n-1)\delta_{u_3} c - \delta_{u_1} c \\ &= \gamma \sum_{u,v} f_{uv}(s_1) - \gamma\lambda(n-1)(n-3)h^2 - \delta_{u_1} c = \Theta(n) \end{aligned}$$

Since the price of anarchy is lower-bounded by $U(s^*)/U(s_1)$, we have that $\text{PoA} = \Omega(n)$. \square




Figure 4: Example of a game under the global risk model with an unbounded price of anarchy.

4.1.2 Dynamics of Network Formation

Despite the fact that the price of anarchy in this setting can be arbitrarily high, we demonstrate that *myopic best-response* dynamics can quickly converge to a socially optimal network.

Myopic Best Response For network s , and an agent u , we define *myopic best-response* by u as follows: let $v^* \in \arg \min_{v \in V \setminus \{u\}} \delta_v$ be a node with the lowest default probability among all nodes except u . Then, u 's myopic best response is to extend credit to v^* , i.e., $\tau_{s'}(u) = v^*$, where $s' = \{c_{uv}(s') : u, v \in V\}$ defined below is the network resulting from u 's myopic best-response in s . For nodes $w, y \in V$,

$$c_{wy}(s') := \begin{cases} c_{wy}(s), & \text{if } w \neq u \\ 0, & \text{if } w = u \text{ and } y \neq v^* \\ c, & \text{if } w = u \text{ and } y = v^* \end{cases}$$

THEOREM 4.5. *Assume that the default probabilities, $\delta_v, v \in V$, are all distinct. Consider a network s . Let s^* be the network obtained after all agents have played myopic best response, starting from s . Then s^* maximizes social welfare.*

PROOF. Since the default probabilities are all distinct, there exists a unique node, say v^* , with the lowest default probability, and another node u^* with the second lowest default probability. Then, observe that for all $u \in V \setminus \{v^*\}$, $\tau_{s^*}(u) = v^*$ and $\tau_{s^*}(v^*) = u^*$. The optimality of s^* follows from Lemma 4.2. \square

Sequential Arrival We consider a model where agents arrive sequentially, and strategically decide which one of the agents in the network to extend credit to. Let s_0 be a network of two agents, say u_0 and v_0 , such that $\tau_{s_0}(u_0) = v_0$ and $\tau_{s_0}(v_0) = u_0$. At each time $t = 1, 2, \dots$, an agent u_t arrives and extends credit to one of agents in the network s_{t-1} in order to maximize $U_{u_t}(s_t)$ where s_t is the resulting network. We denote by V_t the set of agents that have arrived up to and including time t . We show that the agent u_t arriving at time t always extends credit either to u_{t-1} or to $\tau_{s_{t-1}}(u_{t-1})$.

THEOREM 4.6. *For all $t \geq 1$, $\tau_{s_t}(u_t) \in \{u_{t-1}, \tau_{s_{t-1}}(u_{t-1})\}$.*

Since the node u_t arriving at time t always extends credit to either u_{t-1} or $\tau_{s_{t-1}}(u_{t-1})$, the resulting network has a *comb-like* structure, i.e., there is a chain of nodes forming the spine of the network, and each node in that chain is trusted by a number of leaf nodes.

4.2 Simulations on Global and Graded Risk

To address a more general case, we turn to empirical game analysis methods. In this approach, we choose a small set of heuristic strategies for agents to follow, and apply hierarchical reduction [18] to limit the number of players. We repeatedly simulate strategy profiles in this restricted game to estimate their payoffs. Evaluating the resulting empirical game yields insight on general strategic issues as exhibited by the heuristic strategies. This methodology allows us to generalize the setting in several ways: non-uniform transaction

rates and values, issuing credit to multiple counterparts, and graded risk based on incomplete information.

In the experiments reported here, we simulate 60-agent credit networks and construct 6-player hierarchically reduced games in which a multiple of 10 agents plays each strategy. In each simulation run, agents are first assigned strategies, after which the random parameters H (social network), Δ (default probabilities), Λ (transaction rates), and \mathcal{G} (transaction sizes) are realized. Then agents issue credit according to their strategies, defaults occur according to Δ , and 10,000 transactions are attempted according to Λ . Each successful transaction in which agent u buys from agent v adds x_{uv} to u 's payoff and subtracts 1 from v 's, while transferring 1 unit of credit through the network. Each agent also loses c_{uv} for each defaulter v to which it had issued credit. We calculate the payoff to a strategy s as the average payoff to agents playing s . Strategy payoffs are averaged over 250 to 3500 simulation runs as necessary to statistically distinguish empirical game equilibria.

In all simulation environments, the transaction rate λ_{uv} for each pair of agents is drawn uniformly and then normalized. The transaction size distribution $g_{uv}(\cdot) = x_{uv}$ is a singleton for each pair of agents, but the value is drawn from one of two distributions: $x_{uv} \sim U[1, 1.2]$ or $x_{uv} \sim U[1, 2]$. Note that we are using $x_{uv} \sim \mathcal{G}$ here to indicate the value to the buyer u , whereas the seller's cost, and the amount of credit transferred are fixed at one. Default probabilities δ_v for each agent are drawn from one of three Beta distributions: $Beta(1, 1)$, $Beta(1, 2)$, or $Beta(1, 9)$.

Our experiments consider two risk models: *global risk*, with no social network, and *graded risk*, where H is an Erdős-Rényi graph. Under global risk, all agents are fully informed about transaction rates (Λ), transaction values (\mathcal{G}), and default probabilities (Δ). Under graded risk, agents still know Λ and \mathcal{G} , but information about Δ comes in the form of signals whose informativeness decreases exponentially with social network distance. If we call the length of the shortest path between u and v in the social network $|\mathcal{SP}_{uv}|$, then the number of samples u receives from $Binom(\delta_v)$ is $S_{uv} = \lfloor 10^{3-|\mathcal{SP}_{uv}|} \rfloor$, meaning that agents receive 100 samples for their neighbors, 10 for nodes at distance 2, 1 at distance 3, and none at greater distances. If agent u receives a signal with S_{uv} samples including S_{uv}^d defaults, its posterior belief about v 's default probability is $\Delta_{uv} = Beta(\alpha + S_{uv}^d, \beta + S_{uv} - S_{uv}^d)$.

4.2.1 Strategies

We are particularly interested in what criteria agents might use to allocate credit. We therefore focus on heuristic strategies that create a fixed number of credit links (either 0 or 5), and allocate the same amount of credit (5 units) on all links. An agent's strategic decision is then whether to allocate any credit, and if so, what criteria to employ in picking the five nodes to which they connect.

We test eight heuristic strategies under which agent u could issue the following sets of credit links, where the pair (v, c_{uv}) indicates that u issues c_{uv} units of credit to agent v :

- ZE (zero credit): \emptyset




Figure 5: Distributions of social network distances: (a) between pairs of agents; (b) over credit links produced by strategy DP with; (c) over credit links produced by strategy TD. Parameter settings: graded risk; $\delta_v \sim Beta(1, 2)$; $x_{uv} \sim U[1, 2]$.

- IX (index): $\{(v, 5) : v \in \{v_1, \dots, v_5\}\}$
- $\{(v, 5) : v \text{ is among the 5 best agents according to } \dots\}$
 - DP (estimated default probability): Δ_{uv}
 - TV (myopic trade value): $\lambda_{uv}x_{uv}$
 - TP (net trade profit): $\lambda_{uv}x_{uv} - \lambda_{vu}$
 - EU (expected utility): $10^4(1 - \Delta_{uv})(\lambda_{uv}x_{uv} - \lambda_{vu}) - 5\Delta_{uv}$
 - TD (trade then default): $10^4(1 - \Delta_{uv})\lambda_{uv}x_{uv} - 5\Delta_{uv}$
 - DT (default then trade): $\lambda_{uv}x_{uv} - 5\Delta_{uv}$

Some of these strategies warrant further explanation: DP, which chooses the agents least likely to default, has very different results under global and graded risk. In the former case, δ_v is common knowledge for all agents, so all agents playing DP coordinate their credit issuance. Under graded risk, agents' beliefs Δ_{uv} depend on their position in the social network, hence different DP agents make varying choices. By always issuing credit to agents 1 to 5, IX provides a way for agents to coordinate in either the global or graded risk settings.

EU estimates the expected utility attributable to each agent, assuming that all attempted transactions succeed. TD does the same, but excludes the cost from selling to other agents. These strategies both tend to weight $\lambda_{uv}x_{uv}$ much more heavily than Δ_{uv} , so the strategy DT also considers both transactions and defaults, but switches the relative weights.

Fig. 5a illustrates the distribution of social network distances for all agents in the Erdős-Rényi graph under graded risk. The distribution over distances in the social network for credit links produced by strategy DP is shown in Fig. 5b. Comparing these two histograms, we can see that preferring low-default counterparts results in issuing credit to agents nearby in the social network. Although neighbors in the network have the same *prior* probabilities of default as distant agents, the superiority of information about them means that nearby agents are much more prevalent among those with lowest *posterior* probabilities. The strategy DT produces a similar histogram to DP, as it also relies heavily on beliefs about default probability, adding just a small factor for trade value that acts as a tie-breaker. The remaining six strategies are influenced only slightly or not at all by the social network, and therefore exhibit histograms (Fig. 5c) much like the underlying distance distribution shown in Fig. 5a.

4.2.2 Global Risk Model

The results of equilibrium analysis under global risk for each combination of default probability distribution, and buyer surplus

are shown in the top half of Fig. 6. Strategies appearing in a cell are supported in some symmetric mixed-strategy Nash equilibrium of the corresponding game. Each circled strategy is a symmetric pure-strategy Nash equilibrium.

Default probability is clearly the most relevant criterion in the global risk setting. At least one of DP and DT, and often both, is supported in an equilibrium of all settings except the bottom left, which has the lowest transaction values and highest default probabilities (where the empty network is the unique equilibrium). That the empty network is among the equilibria in all three environments with low transaction values is an indication of the importance of network effects: it is much more profitable to participate in a credit network if many other agents do so as well. We also observe the importance of coordinating on a centralized, star-like network, in that DP and IX both appear as symmetric pure strategy equilibria. This point is reinforced by the poor performance of the strategies relying primarily on transaction value: TV, TP, and TD.

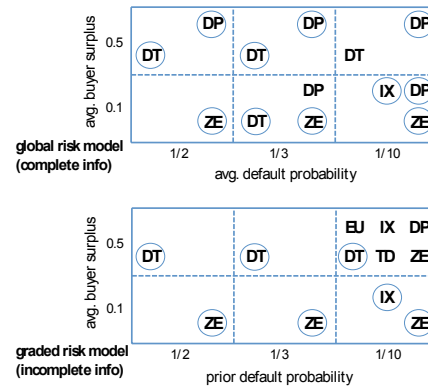


Figure 6: Strategies appearing in symmetric Nash equilibria in six global risk environments (top), and six graded risk environments (bottom). Circled strategies in a cell constitute pure symmetric equilibria of the associated game.

4.2.3 Graded Risk Model

The bottom half of Fig. 6 shows equilibrium analysis under graded risk. DP, which appears in equilibria of nearly all global risk settings is less prevalent in graded risk equilibria, indicating that it owes much of its success to its function as a coordination device. TD, on the other hand, continues to perform well, making clear

that default probability is still a relevant criterion to consider. In the low-value settings, it is unsurprising that the empty network is more likely to arise, because agents have less information available. When defaults are sufficiently rare, a centralized network from all agents playing IX is again an equilibrium. The case of high value and low default probability shows a multiplicity of equilibria, including the only appearance of one of the strategies focused on trades: TD.

The clear messages from our empirical game simulations are twofold. First, coordinating on a centralized, star-like network can be very beneficial, when such coordination is feasible. Second, network effects are very strong: none of the strategies that rely heavily on transaction-related criteria perform well; instead agents tend to play strategies like TD, IX, and DP that result in a high likelihood of remaining connected to most of the network in equilibrium.

5. CONCLUSION

Our investigation of strategic issues in the formation of credit networks characterizes, in various settings, the nature and efficiency of credit networks that are formed by self-interested agents autonomously choosing how to issue credit among available counterparts. The analysis employs game-theoretic solution concepts, employed in theoretical examination of analytic models, as well as simulation-based exploration of extended environments.

Under dichotomous risk, if only bilateral transactions are allowed, we show that the formation game is a potential game. Moreover, for many transaction size distributions, we show that agents' utilities are concave, and consequently, every Nash equilibrium of the game maximizes social welfare. More interestingly, we showed that the Nash equilibria are equivalent in a much stronger sense: all Nash equilibria are *cycle-reachable* [6] from each other, which implies that the sequences of transactions that can be supported from each equilibrium network are identical. However, when we allow transactions over longer paths, best-response dynamics may not converge, and the price of anarchy is unbounded.

Under a model of global risk, if agents are limited to extend credit to at most one other agent, we prove that the networks formed in equilibrium have a star-like structure. Although the price of anarchy is unbounded, myopic best response quickly converges to a social optimum. Even when agents are allowed to extend credit to multiple agents, we show using empirical game simulation that non-empty equilibria tend to be star-like. We also analyze several graded risk settings, and find that agents coordinate in a star-like structure only when defaults are relatively unlikely, and otherwise, credit links tend to be issued over short social distances conforming to the locality of information.

6. ACKNOWLEDGMENTS

This research was supported in part by NSF award IIS-0904325. Part of the research was also sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

7. REFERENCES

- [1] E. Anshelevich and M. Hoefer. Contribution games in networks. *Algorithmica*, pages 1–40, 2011. 10.1007/s00453-011-9520-7.
- [2] V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68:1181–1229, 2000.
- [3] M. A. Barnard. Needle sharing in context: Patterns of sharing among men and women injectors and HIV risks. *Addiction*, 88:805–812, 1993.
- [4] L. Blume, D. Easley, J. Kleinberg, R. Kleinberg, and E. Tardos. Network formation in the presence of contagious risk. In *12th ACM Conference on Electronic Commerce*, pages 1–10, San Jose, CA, 2011.
- [5] J. Corbo, A. Calvó-Armengol, and D. Parkes. A study of Nash equilibrium in contribution games for peer-to-peer networks. *SIGOPS Operating Systems Review*, 40:61–66, 2006.
- [6] P. Dandekar, A. Goel, R. Govindan, and I. Post. Liquidity in credit networks: A little trust goes a long way. In *12th ACM Conference on Electronic Commerce*, pages 147–156, San Jose, CA, 2011.
- [7] D. B. DeFigueiredo and E. T. Barr. TrustDavis: A non-exploitable online reputation system. In *7th IEEE International Conference on E-Commerce Technology*, pages 274–283, Washington, DC, 2005.
- [8] A. Fabrikant, A. Luthra, E. N. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *22nd ACM Symposium on Principles of Distributed Computing*, pages 347–351, 2003.
- [9] A. Ghosh, M. Mahdian, D. M. Reeves, D. M. Pennock, and R. Fugger. Mechanism design on trust networks. In *3rd International Workshop on Internet and Network Economics*, pages 257–268, 2007.
- [10] M. O. Jackson, T. Rodríguez Barraquer, and X. Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, pages 1–52, 2011.
- [11] M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
- [12] J. A. Jacquez, C. P. Simon, J. Koopman, L. Sattenspiel, and T. Perry. Modeling and analyzing HIV transmission: The effect of contact patterns. *Mathematical Biosciences*, 92(2):119 – 199, 1988.
- [13] D. Karlan, M. Mobius, T. Rosenblat, and A. Szeidl. Trust and social collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, 2009.
- [14] Z. Liu, H. Hu, Y. Liu, K. W. Ross, Y. Wang, and M. Mobius. P2P trading in social networks: The value of staying connected. In *29th IEEE International Conference on Computer Communications*, pages 1–9, 2010.
- [15] M. Mahdian. Fighting censorship with algorithms. In *5th International Conference on Fun with Algorithms*, pages 296–306, Ischia, Italy, 2010.
- [16] A. Mislove, A. Post, P. Druschel, and K. P. Gummadi. Ostra: Leveraging trust to thwart unwanted communication. In *5th Usenix Symposium on Networked Systems Design and Implementation*, pages 15–30, San Francisco, 2008.
- [17] A. Neyman. Correlated equilibrium and potential games. *International Journal of Game Theory*, 26:223–227, 1997.
- [18] M. P. Wellman, D. M. Reeves, K. M. Lochner, S.-F. Cheng, and R. Suri. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *20th National Conference on Artificial Intelligence*, pages 502–508, Pittsburgh, 2005.