Strategic Formation of Credit Networks

PRANAV DANDEKAR, ASHISH GOEL, Stanford University
MICHAEL P. WELLMAN and BRYCE WIEDENBECK, University of Michigan

Credit networks are an abstraction for modeling trust among 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. We find strong positive network formation results for the simplest theoretical model. When each agent trusts a fixed set of other agents and transacts directly only with those it trusts, all pure-strategy Nash equilibria are social optima. However, when we allow transactions over longer paths, the price of anarchy may be unbounded. On the positive side, when agents have a shared belief about the trustworthiness of each agent, simple greedy dynamics quickly converge to a star-shaped network, which is a social optimum. Similar star-like structures are found in equilibria of heuristic strategies found via simulation studies. In addition, we simulate environments 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. Overall, we find that networks formed by self-interested agents achieve a high fraction of available value, as long as this potential value is large enough to enable any network to form.

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

General Terms: Theory, Economics

Additional Key Words and Phrases: Trust, credit networks, strategic network formation, empirical game-theoretic analysis

ACM Reference Format:
Dandekar, P., Goel, A., Wellman, M. P., and Wiedenbeck, B. 2015. Strategic formation of credit networks. *ACM Trans. Internet Technol.* 15, 1, Article 3 (February 2015), 41 pages.
DOI:http://dx.doi.org/10.1145/2700058

3

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 one another based on their local information. In general, establishing a connection incurs

Preliminary versions of parts of this work were presented at the *2011 Workshop on Trading Agent Design and Analysis* (TADA11), the *21st International World Wide Web Conference* (WWW12), and the *50th Allerton Conference on Communication, Control, and Computing*.
This research was supported in part by NSF award IIS-0904325 and by the Army Research Laboratory under Cooperative Agreement W911NF-09-2-0053.
Authors' addresses: P. Dandekar (corresponding author), A. Goel, Department of Management Science and Engineering, Stanford University, Stanford, CA; email: ppd@standford.edu; M. P. Wellman, B. Wiedenbeck, Division of Computer Science and Engineering, University of Michigan, Ann Arbor, MI.

DOI:http://dx.doi.org/10.1145/2700058

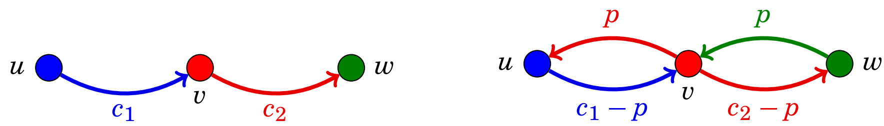

(a) before the transaction (b) after the transaction

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

a cost but also yields some benefit to connected agents. 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 [Anshelevich and Hoefer 2012; Bala and Goyal 2000; Corbo et al. 2006; Fabrikant et al. 2003; Jackson and Wolinsky 1996]. 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, how can agents select among them? This article is an investigation into some of these questions in the context of *credit networks*, an abstraction for modeling trust among agents.

1.1. Credit Networks

Credit networks model trust in terms of one agent's willingness to perform services for another without immediate compensation. We interpret services broadly as any activity, for example, completing a task or delivering a product, that is costly for the provider and beneficial for the recipient. Performing a service incurs a reciprocal obligation which may get discharged directly or indirectly in subsequent interactions. The general idea is that the agents keep track of levels of trust as *credit balances*, and update these balances as services are provided and obligations discharged.

A credit network represents these credit 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 issued by v up to value c. These IOUs can be thought of as obligations, denominated in the issuer's private *currency*. It is possible that in the future v will refuse to honor its outstanding obligations, in which case v's currency becomes worthless and u gets stuck with irredeemable IOUs. The capacity c bounds u's loss in this eventuality, and therefore can be viewed as a measure of u's trust in v.

Credit commitments between trusting nodes also enable remote transactions, as illustrated in Figure 1. Say node w wants p units of service 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 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. Observe that routing payments in credit networks is identical to routing residual flows in single-commodity 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, sellers gain it, and intermediaries exchange credit between their neighbors, with the total credit in the network unchanged.

1.1.1. Origins of the Model. The credit network model was invented independently by (at least) four distinct groups of researchers motivated by somewhat different issues and applications, but arriving at the same essential elements.

— DeFigueiredo and Barr [2005] sought a reputation system with bounded loss from coalitions of malicious users.
— Ghosh et al. [2007] aimed to support distributed payment and multi-user credit checking for multi-item auctions.
— Karlan et al. [2009] wanted to construct an economic model of informal borrowing networks.
— Mislove et al. [2008] were concerned with deterring spam.

A common thread in the objectives of these researchers was to capture a notion of pairwise trust, representable in quantified terms. In each case, the trust measure is grounded by interpreting the quantity as a capacity for transaction. That is, the degree of trust in one agent for another is measured by how much it is willing to expose itself to transactions with that counterpart. In other words, the model operationalizes trust as an extension of *credit*, in a framework where a credit balance entitles an agent to transact with the agent granting credit. The common underlying credit model of these four proposals was first noticed by Dandekar et al. [2011], who introduced the unifying term "credit network" and its formal definition. By introducing suitable definitions of transaction, credit networks can support a wide variety of applications. For example, the inventors enumerated before interpret transactions, respectively, as obtaining references guaranteeing good behavior [DeFigueiredo and Barr 2005], paying for auction winnings [Ghosh et al. 2007], borrowing an asset [Karlan et al. 2009], and communicating messages [Mislove et al. 2008]. Subsequent authors proposed using this framework to support networked asynchronous bilateral trading [Liu et al. 2010], and bartering of tutorial services [Limpens and Gillet 2011].

1.1.2. Properties. Routing payments along chains of trust ensures that agents hold IOUs issued only by other agents that they *directly* trust. As a result, if an agent defaults on its outstanding obligations, the only agents that incur a loss are those that extended credit to the defaulting agent. Thus losses from default are *localized*. Moreover, the total loss incurred is *bounded* by the total credit extended to the defaulting agent. These properties make credit networks robust against two attacks to which centralized currency systems are vulnerable: *whitewashing* [Friedman and Resnick 2001], and *Sybil attacks* [Friedman et al. 2007]. Viswanath et al. [2012] argue that, in fact, all reputation schemes designed for Sybil tolerance have essentially been versions of the credit network idea. They propose an approximation to the max-flow calculation that enables scalability to very large networks.

The effectiveness of credit networks in supporting distributed transactions was most powerfully demonstrated by Dandekar et al. [2011]. Their analysis posits that nodes repeatedly transact with each other according to a known probability distribution. In particular, they showed analytically and via simulations that, for several classes of graphs and with symmetric transaction probabilities, the long-term transaction failure probability in credit networks is comparable to that in equivalent centralized currency systems. Thus, in addition to being robust against attacks by malicious agents, credit networks also provide a high degree of *liquidity*: the ability to sustain long sequences of transactions.

1.2. Formation of Credit Networks

Extending credit to other agents increases liquidity in the network, enabling more profitable transactions to go through. However, it also entails risk, since a counterparty might default on its outstanding obligations. This raises the natural question: if agents rationally weigh these risks and benefits, what kinds of networks will they form? In order to use credit networks for practical applications, it is critical to understand the structural and economic properties of the credit networks formed by strategic agents. We address this question in this article.

1.2.1. Our Setting. 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. They derive value from successful transactions. Extending credit to other agents increases transaction success probability, thus contributing to utility. On the other hand, when an agent defaults, it imposes losses on its creditors up to the amount of the credit it received. Thus, an agent's net utility is its total value from successful transactions minus the utility loss from extending credit to untrustworthy agents. We study several variants of this game with different models of risk, both analytically and through simulations. Our simulations employ an approach known as *empirical game-theoretic analysis* (EGTA) [Wellman 2006].

We start with a model of *dichotomous risk*, where agents divide their counterparts into fully trusted and untrusted categories. We assume the trust relation is symmetric, and consider it as arising from a *social network* represented by an *undirected* graph. Agents trust their neighbors in the social network and are willing to extend them credit. They consider everyone else untrustworthy, and consequently never extend credit to non-neighbors. Though the strict dichotomy surely oversimplifies, it may approximate reality for situations where dealing with strangers is particularly risky (e.g., where identities are weak and norms of good behavior are poorly established). It may also capture heuristic behavior of agents who are highly risk averse or especially prize simplicity in social rules.

We also study a model of *global risk*, which represents an opposite extreme to the dichotomous risk model. Under global risk, each agent has a publicly known probability of default. This model captures situations involving small, densely interacting social groups, or where there are institutions such as credit-reporting agencies that systematically gather and disseminate relevant risk information.

Finally, we study a model of *graded risk* that bridges the gap between global and dichotomous risk. Under this model, agents are related on a social network, and 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.

1.2.2. Our 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 3.2). This implies that best-response dynamics always converge to a *pure-strategy Nash equilibrium* (PSNE). 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 3.3). We further show that the credit networks generated by any two PSNE are *cycle reachable* from each other (Theorem 3.5), which means that they support the same sequences of transactions [Dandekar et al. 2011].

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

Global Risk. Under global risk, we analyze several scenarios. First, we investigate PoA and the structure of equilibria when each agent is limited to extend credit to at most one other agent. We prove that, if we disallow the empty network as an outcome, the PoA of the formation game is unbounded (Theorem 4.4), even though all PSNE networks have a star-like structure (Theorem 4.3). Instead we focus on the structure of equilibria under two simple dynamics: sequential arrival and greedy dynamics. 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 to which v extends credit (Theorem 4.6). Thus the resulting network has a *comb-like* structure. Under greedy dynamics, nodes extend their entire credit budget to the node that has the lowest risk of default. If the default probabilities are all distinct, 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 PoA can be unbounded, nodes can easily find the optimal network using greedy dynamics.

We also use empirical game-theoretic simulations to investigate a richer model of global risk in which agents are not constrained to a single link or a fixed budget, and transaction probabilities and values may be asymmetric. Under global risk, we find star-like equilibrium networks under all conditions, but settings with high default rates or low transaction surpluses also have empty network equilibria. In addition, when default rates are low, transaction-dependent credit-issuing strategies can appear in equilibrium.

Graded Risk. We address graded risk exclusively through empirical game simulations. In this setting, the star-like equilibria disappear, highlighting the importance of ensuring that central nodes are unlikely to default. Empty network equilibria are present for exactly the same settings as under global risk, indicating that the conditions under which agents issue no credit may not depend as strongly on the information structure. Transaction-dependent equilibria arise under the same conditions as for global risk as well as those with high transaction surplus.

Summary and Key Insights. The theoretical analyses demonstrate generally how network formation depends on environmental conditions. We get strong positive results for dichotomous risk and symmetric bilateral transactions but, once we allow transactions on paths, the worst-case network formation performance becomes arbitrarily bad. Our simulation-based analysis therefore investigates how credit network formation by self-interested agents might play out in representative, less technically constrained environments. The results confirm that the worst case indeed cannot be avoided in that networks may simply fail to form if the benefits are not sufficiently large. On the other hand, we find that when the potential value of the network is large, self-interested agents will pursue strategies that lead to networks that achieve the lion's share of this benefit. In other words, in complex environments it is too much to expect perfect network formation or even guaranteed lower bounds on relative performance, but when demand for profitable transactions is present, self-interested agents will form viable credit networks.

1.2.3. Comparison to Other Network Formation Games. The various models of strategic credit network formation that we analyze relate to an extensive literature on network creation games. Our simplest model, with dichotomous risk and bilateral transactions, can be viewed as a special case of a *network contribution game* as defined by Anshelevich and Hoefer [2012]. In a network contribution game, agents allocate a budget of effort among their neighbors and receive a reward for each link that is increasing in the sum of their effort and their neighbor's. The authors state a result that generalizes Theorem 3.2 but otherwise focus on bilateral link creation models. Because credit edges are by nature directed, our analysis of this game focuses on the implications of equilibria under unilateral deviations.

When we lift the constraint of bilateral transactions, path lengths become a relevant strategic consideration. Fabrikant et al. [2003] studied a model where agents benefit from reducing their distance to others in the network. In their network creation games, agents incur costs for creating edges and also for the sum of their distances to all other nodes. In a credit network, agents care about the existence and the capacity of paths, as well as the likelihood of neighbors defaulting. In both the global and graded risk settings, our results depend more strongly on considerations of default risk than on distance between nodes.

Network formation in the presence of risk was studied by Blume et al. [2011] 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). The credit network model flips this: nodes derive benefit from transactions along direct as well as multihop paths, whereas only direct edges are sources of risk.

2. MODEL AND DEFINITIONS

Let V denote the set of n agents, and equivalently, the n nodes of the formed credit network.[1]

2.1. Game Model

Agents play a one-shot game where they choose credit allocations to form an *initial network* G. An edge from node u to node v of capacity $c_{uv}(G)$ represents the credit extended by agent u to agent v in the network G. Thus G is a directed graph with edge capacities. Agents choose how much credit to extend to whom based on public or private information about transactions (Section 2.2), as well as the default risk of other agents (Section 2.4). The credit agent u offers may be constrained by a budget $B_u \geq 0$, representing the total credit that u can extend to other agents.

A strategy for agent u is a mapping of u's public and private information to a *feasible* credit allocation $\{c_{uv}(G), v \in V : c_{uv}(G) \geq 0 \text{ and } \sum_{v\in V} c_{uv}(G) \leq B_u\}$. The combination of agent strategies applied to the given information induces the initial network G. The payoff to agent u is determined by the transactions and defaults that play out on the network starting from G, according to the models described next.

[1] These and other symbols employed in the article are summarized in the appendix, Table III.

2.2. Transaction Model

Once a network G is formed, agents engage in a sequence of transactions with each other. Following prior work [Dandekar et al. 2011], transactions are generated according to a stochastic process. At each time step $t = 1, 2, \ldots$, 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^t_{uv} between u and v is drawn from a *transaction size distribution* over $[0, \infty)$ with a *probability density function* (pdf) $X_{uv}(\cdot)$ and a corresponding *cumulative distribution function* (cdf) $\mathcal{X}_{uv}(\cdot)$. We assume that the pdfs $X_{uv}(\cdot)$ are public. Let $\mathbb{X} := \{X_{uv}(\cdot) : u, v \in V\}$ be the pdf matrix.

In the general credit network framework, transaction payments may be routed as a flow over multiple paths. For all the settings considered in this article (bilateral or single-unit transactions, defined shortly), it is sufficient to route payment over a single path. Given a transaction size x, a *feasible path* in the network G from node v to node u is a set of directed edges $\mathcal{P} = \{(v, u_1), (u_1, u_2), \ldots, (u_{k-1}, u_k), (u_k, u)\}$ such that for all $(w, y) \in \mathcal{P}$, $c_{wy}(G) \geq x$. We route payments along the shortest feasible path in the network. Let $\mathcal{P}^t_{vu}$ be the shortest feasible path in the credit network from v to u at time t. A successful transaction of size x^t_{uv} results in a change of credit capacities along edges in $\mathcal{P}^t_{vu}$ as follows. Let G^t denote the state of the network G at time $t = 0, 1, 2, \ldots$, where $G^0 = G$. Then, for $w, y \in V$ and a successful transaction at time $t > 0$,

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

So, in order for a payment x^t_{uv} 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 $u, v \in V$, $c_{uv}(G^t) + c_{vu}(G^t) = c_{uv}(G) + c_{vu}(G)$.

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

2.3. Utility

Agents choose credit allocations to maximize their utility. Agents derive value from successful transactions, but they risk loss of utility when they extend credit to potentially untrustworthy agents. We model this risk in several ways, but can generally denote the expected loss of utility to u associated with the prospect of default by v by $\Delta_{uv}(G)$, with the constraints that $\Delta_{uv}(G) \geq 0$ and $\Delta_{uv}(G) > 0$ only if $c_{uv}(G) > 0$.

In our analytical treatments (Sections 3 and 4), we assume that transaction values are uniform across (u, v) pairs. We can therefore define the value derived from transactions as proportional to the steady-state success probability of all transactions where the agent is a payer. Let $p_{uv}(G)$ be the steady-state success probability of the transactions from u to v (i.e., where u is the payer) when the initial network is G. Then, the total utility of an agent u when the initial network is G is given by

$$U_u(G) = \gamma \sum_{w \in V} p_{uw}(G) - \sum_{v \in V : c_{uv}(G) > 0} \Delta_{uv}(G), \tag{1}$$

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

It is difficult to characterize the steady-state transaction success probabilities for arbitrary networks and transaction regimes. However, for the settings we analyze, we are able to exploit results established by Dandekar et al. [2011].

2.4. Risk Model

In order to model variation in $\Delta_{uv}(G)$, we assume the agents are embedded in an exogenously defined *social network* represented by a simple undirected graph $H = (V, E)$. The social network H influences how $\Delta_{uv}(G)$ 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.

2.4.1. Dichotomous Risk. In this model, an agent u partitions the set of other agents according to whether they are neighbors or non-neighbors in H. For network G, agent u estimates risk exposure to be

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

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

2.4.2. 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 $c_{uv}(G)$ to v loses $c_{uv}(G)$ units. Thus $\Delta_{uv}(G) = \delta_v c_{uv}(G)$.

2.4.3. 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 $\delta_v \sim Beta(\alpha, \beta)$. Agent u then receives a signal in the form of some number of samples ∂_{uv} drawn from the binomial distribution on δ_v, where ∂_{uv} decreases exponentially with social network distance. Given this form of evidence, the posterior is also beta distributed. Specifically, if u's sample for v includes $\hat{d}$ defaults, then its posterior default probability is $(\alpha + \hat{d})/(\alpha + \beta + \partial_{uv})$.

3. NETWORK FORMATION UNDER DICHOTOMOUS RISK

Recall that under dichotomous risk, $\Delta_{uv}(G)$ is defined by (2); as a result nodes extend credit only 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), the transaction fails.

First we derive an expression for steady-state transaction success probability $p_{uv}(G)$. For edge $e = (u, v) \in E$, let $c_e(G) := c_{uv}(G) + c_{vu}(G)$.

LEMMA 3.1. *Fix a credit network G. For nodes $u, v \in V$ such that $e = (u, v) \in E$, the steady-state transaction success probability $p_{uv}(G)$ is given by*

$$p_{uv}(G) = p_{uv}(c_e(G)) = \begin{cases} \frac{\lambda_{uv}}{c_e(G)} \int_0^{c_e(G)} \mathcal{X}_{uv}(y)dy, & \text{if } c_e(G) > 0 \\ 0, & \text{if } c_e(G) = 0 \end{cases}. \quad (3)$$

See Appendix B.1 for a proof. Observe that $p_{uv}(G)$ depends on only the total credit capacity $c_e(G)$ along the edge $e = (u, v)$. How the total capacity is divided initially between the two directions does not affect the steady state, hence $p_{uv}(G) = p_{vu}(G)$. For the rest of this section, we exploit the special form of (3) and cast p_{uv} directly as a function of $c_e(G)$. That is, $p_{uv}(x)$ is the steady-state transaction success probability along edge $e = (u, v)$ when the total credit allocated along it is x. We write $p_{uv}(G)$ to mean $p_{uv}(c_e(G))$ when there is no ambiguity.

In a *potential game* [Monderer and Shapley 1996], there exists a global potential function that captures the change in payoff with respect to changes in strategy.

THEOREM 3.2. *The network formation game under a symmetric bilateral transaction regime is a potential game.*

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

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

As we are in a symmetric bilateral transaction regime, it follows from Lemma 3.1 that $p_{uv}(G) = p_{vu}(G)$ for all $(u, v) \in E$, and $p_{uv}(G) = 0$ if $(u, v) \notin E$. Therefore,

$$\sum_{u \in V} \sum_{v \in V} p_{uv}(G) = 2 \sum_{(u,v) \in E} p_{uv}(G).$$

This implies $\Phi(G) = \gamma \sum_{(u,v) \in E} p_{uv}(G)$. To show that $\Phi(G)$ is a potential function, let us fix a node $u \in V$. Consider a network G' which differs from G only in the credit allocation of u. Formally, for all $w, y \in V$,

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

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 $e' = (u', v') \notin E_u$, $c_{u'v'}(G) = c_{u'v'}(G')$. As a result, $p_{u'v'}(G) = p_{u'v'}(G')$. It follows that

$$\Phi(G) - \Phi(G') = \gamma \sum_{(u,v) \in E_u} \big(p_{uv}(G) - p_{uv}(G')\big) = U_u(G) - U_u(G').$$

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

From well-established properties of potential games [Monderer and Shapley 1996], Theorem 3.2 implies that, in this setting, a pure-strategy Nash equilibrium always exists, and best-response dynamics always converge to a PSNE. Moreover, because the potential function is given by $\Phi(G) = U(G)/2$, the price of stability is 1.

3.1.1. Nash Equilibria Maximize Social Welfare. Next we show that, under certain technical conditions on the transaction size distributions, every Nash equilibrium under a symmetric bilateral transaction regime maximizes social welfare.

THEOREM 3.3. *Assume that for every edge $(u, v) \in E$: (i) $X_{uv}(\cdot)$ is nonincreasing, (ii) $X_{uv}(\cdot)$ has support over $[0, \infty)$, and (iii) $\mathcal{X}_{uv}(\cdot)$ is twice differentiable. Let G be an arbitrary Nash equilibrium of the network formation game under a symmetric bilateral transaction regime $\langle \Lambda, \mathbb{X} \rangle$. Then G maximizes social welfare $U(G)$.*

In order to establish this theorem, we prove some properties of the functions $p_{uv}(\cdot)$.

LEMMA 3.4. *For an edge $(u, v) \in E$, the steady-state transaction success probability $p_{uv}(\cdot)$ is continuously differentiable, strictly increasing, and concave.*

See Appendix B.2 for a proof. As a corollary, if $X_{uv}(\cdot)$ is strictly decreasing instead of nonincreasing, $p_{uv}(\cdot)$ is strictly concave. Recall from Theorem 3.2 that the formation game under a symmetric bilateral transaction regime is a potential game and $\Phi(G) = U(G)/2 = \gamma \sum_{e \in E} p_{uv}(G)$ is a potential function. From Lemma 3.4, we know that $p_{uv}(\cdot), (u, v) \in E$, are concave and continuously differentiable, which implies $\Phi(\cdot)$ is concave and continuously differentiable. It was shown by Neyman [1997] that any Nash equilibrium of a potential game with a concave and continuously differentiable potential is also a potential maximizer. Therefore, G maximizes $\Phi(G)$ or, equivalently, $U(G)$. This completes the proof of Theorem 3.3.

Note that the conditions imposed upon the transaction size distributions $X_{uv}(\cdot)$ in the theorem are satisfied by many natural distributions, including exponential, mean-zero normal, and power-law distributions.

3.1.2. Nash Equilibria are Cycle Reachable. Theorem 3.3 implies an equivalence in social welfare among Nash equilibria. Next we show that if the transaction size distributions $X_{uv}(\cdot)$ are strictly decreasing instead of nonincreasing as in Theorem 3.3, the pure-strategy Nash equilibria of this game are equivalent in a much stronger sense.

Definition 3.1 [Dandekar et al. 2011]. Let G and G' be two credit networks. We say that G' is *cycle reachable* from G if G can transformed into G' by routing a sequence of payments along feasible cycles (i.e., from a node to itself along a feasible path).

The significance of this property, established by Dandekar et al. [2011], is that the sequences of transactions that succeed starting from G and starting from G' are identical.

THEOREM 3.5. *Assume that for every edge $(u, v) \in E$: (a) $X_{uv}(\cdot)$ is strictly decreasing, (b) $X_{uv}(\cdot)$ has support over $[0, \infty)$, and (c) $\mathcal{X}_{uv}(\cdot)$ is twice differentiable. Let G and G' be networks generated by two PSNE of the network formation game under the symmetric bilateral transaction regime $\langle \Lambda, \mathbb{X} \rangle$. Then G and G' are cycle reachable from each other.*

In order to prove this theorem, we first show that the total credit capacity of any edge in E is identical in any PSNE.

LEMMA 3.6. *For all edges $e = (u, v) \in E, c_e(G) = c_e(G')$.*

PROOF. We state here the propositions that prove this lemma; proofs can be found in Appendix B.3. First, observe that, for an edge $(u, v) \in E$, the steady-state transaction success probability $p_{uv}(\cdot)$ is strictly concave, strictly increasing, and continuously differentiable (by Lemma 3.4). Let us define the marginal utility of an edge.

Definition 3.2. The *marginal utility* of an edge $(u, v) \in E$ is the function $r_{uv} : \mathbb{R}_+ \to \mathbb{R}_+$ given by

$$r_{uv}(x) = p'_{uv}(x) = \frac{dp_{uv}(x)}{dx}.$$

We first show that for any edge $(u, v) \in E, r_{uv}(G) = r_{uv}(G')$.

Since $p_{uv}(\cdot)$ is strictly concave, strictly increasing, and continuously differentiable, $r_{uv}(\cdot)$ is continuous, strictly decreasing, and strictly positive. In network G, the marginal utility on an edge $e = (u, v) \in E$ is given by $r_{uv}(c_e(G))$. We denote it by $r_{uv}(G)$ when there is no ambiguity.

Definition 3.3. For a node $u \in V$ and a network G, we define $r^*_u(G) := \max_{(u,w)\in E_u} r_{uw}(G)$ and $E^*_u(G) \subseteq E_u$ as the set of edges $(u, w) \in E_u$ such that $r_{uw}(G) = r^*_u(G)$.

In words, $E^*_u(G)$ is the set of edges incident on node u that have the highest marginal utility in network G among all edges in E_u. We show that in any PSNE network G, each node u exhausts its entire budget and allocates nonzero credit only along edges in $E^*_u(G)$.

PROPOSITION 3.1. *Let G be a PSNE network. Then, for all nodes $u \in V$:*

(1) $\sum_{v:(u,v)\in E} c_{uv}(G) = B_u$; and
(2) for each $(u, v) \in E$, if $(u, v) \notin E^*_u(G)$, then $c_{uv}(G) = 0$.

Next we define a *slack edge*.

Definition 3.4. Let G be a PSNE network. We call an edge $e = (u, v) \in E$ a *slack* edge in G if $e \notin E^*_u(G)$ or $e \notin E^*_v(G)$.

Note that by Proposition 3.1, if (u, v) is a slack edge in PSNE network G, it must be that $c_{uv}(G) = 0$ or $c_{vu}(G) = 0$.

Definition 3.5. Let G be a credit network. We define:

(1) $r^{\min}_G := \min_{(u,v)\in E} r_{uv}(G)$ to be the minimum marginal utility of any edge;
(2) the set $E^{\min}_G := \{(u, v) \in E \mid r_{uv}(G) = r^{\min}_G\}$;
(3) the set $V^{\min}_G := \{u \in V \mid u$ is incident on some edge in $E^{\min}_G\}$; and
(4) the set $V^+_G \subseteq V^{\min}_G$ as

$$V^+_G := \{u \in V \mid u \text{ is incident upon some edge in } E^{\min}_G \text{ and upon some edge in} E \setminus E^{\min}_G\}.$$

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

PROPOSITION 3.2. *Let G and G' be two PSNE networks. Then $r_G^{\min} = r_{G'}^{\min}$.*

Moreover, in any two PSNE networks G and G', the set of edges with the minimum marginal utility in G is identical to that in G'.

PROPOSITION 3.3. *Let G and G' be two PSNE networks. Then $E_G^{\min} = E_{G'}^{\min}$.*

COROLLARY 3.1. *Let G and G' be two PSNE networks. Then $V_G^{\min} = V_{G'}^{\min}$ and $V_G^+ = V_{G'}^+$.*

Thus we have established that, for any two PSNE networks G and G', $r_{uv}(G) = r_{uv}(G')$ for all edges $(u,v) \in E_G^{\min}$. We show using an inductive argument that this is true of all edges in E.

Definition 3.6. Given an instance $I : H = (V,E); p_{uv}, (u,v) \in E; B_u, u \in V$ defining a network formation game under a symmetric bilateral transaction regime, a credit network G, and an arbitrary set of edges $F \subseteq E$, we define the (G,F)*-restriction* of I, denoted $I_{(G,F)}$, as $H^{(F)} := (V, E \setminus F)$, $p_{uv}^{(F)} := p_{uv}, (u,v) \in E \setminus F$, and

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

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

Definition 3.7. Given a credit network G and an arbitrary set of edges $F \subseteq E$, we define an F*-restriction* of G, denoted $G_{(F)}$, as follows: for all edges $(u,v) \in E \setminus F, c_{uv}(G_{(F)}) = c_{uv}(G)$ and $c_{vu}(G_{(F)}) = c_{vu}(G)$.

PROPOSITION 3.4. *If G is a PSNE network for instance I of the network formation game in the symmetric bilateral transaction setting, then $G_{(F)}$ is a PSNE network for $I_{(G,F)}$ for any set $F \subseteq E$.*

PROPOSITION 3.5. *Let G and G' be two PSNE networks for instance I of the network formation game under a symmetric bilateral transaction regime. Then for all edges $(u,v) \in E$, $r_{uv}(G) = r_{uv}(G')$.*

Observe that, since $p_{uv}(\cdot)$ is strictly concave, $r_{uv}(\cdot)$ is strictly decreasing. Therefore, Proposition 3.5 implies that for all $e = (u,v) \in E$, $c_e(G) = c_e(G')$, completing the proof of Lemma 3.6. □

Lemma 3.6 allows us to show that any two PSNE networks are cycle reachable. First we define the *generalized score* of a node as the total credit issued to it in network G.

Definition 3.8 [Dandekar et al. 2011]. The *generalized score* of node u in G is $d_u(G) := \sum_{v\in V} c_{vu}(G)$.

Next we show that any two PSNE networks have identical generalized scores for all nodes.

PROPOSITION 3.6. *Let G and G' be two PSNE networks. Then, for all $u \in V$, $d_u(G) = d_u(G')$.*

PROOF. Fix a node $u \in V$. Recall from Proposition 3.1 that

$$\sum_{v:(u,v)\in E} c_{uv}(G) = \sum_{v:(u,v)\in E} c_{uv}(G') = B_u. \tag{4}$$

Also, from Lemma 3.6, we know that for all edges $e \in E$,

$$c_e(G) = c_e(G'). \tag{5}$$

It follows from (4) and (5) that

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

□

PROPOSITION 3.7 [DANDEKAR ET AL. 2011]. *Two credit networks G and G' are cycle reachable if and only if, for all $u \in V$, $d_u(G) = d_u(G')$.*

From Propositions 3.6 and 3.7, it follows that G and G' are *cycle reachable*. This completes the proof of Theorem 3.5.

3.2. Long-Distance 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. The analytical results we prove in this section and the next employ instances where all transactions are for a single unit.

Definition 3.9. A *unit transaction regime* over credit network G is a transaction regime $\langle \Lambda, \mathbb{X} \rangle$ where, for all $u, v \in V$ and for all $t > 0$, the transaction size $x^t_{uv} = 1$, the transaction rate matrix Λ is symmetric, and the Markov chain $\mathcal{M}(G, \Lambda, \mathbb{X})$ is ergodic.

When the network G is acyclic (ignoring directionality), there is a single path $\mathcal{P}_{uv}$ between any pair of nodes u and v. For this case, Dandekar et al. [2011] characterize the steady-state success probabilities under a unit transaction regime.

LEMMA 3.7 [DANDEKAR ET AL. 2011]. *Consider a credit network G. Assume that G is acyclic if we ignore the directions of the edges in G. Then, in a unit transaction regime over G, the steady-state transaction success probability $p_{uv}(G)$ from node $u \in V$ to node $v \in V$ is given by*

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

Using this characterization, we show that the network formation game in this setting is not a potential game.

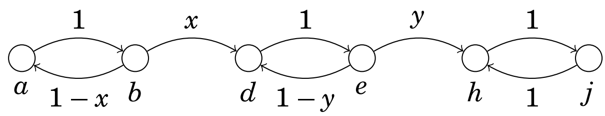

Fig. 2. Example of a formation game that does not admit a PSNE.

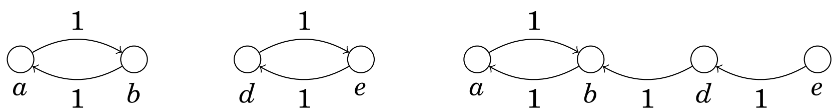

(a) an equilibrium network G (b) an optimal network G^*

Fig. 3. Example of a game with an unbounded price of anarchy.

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

PROOF. 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 nonzero 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 3.7 to compute the steady-state transaction success probabilities between nodes.

Let G be a PSNE network. Then, it must be that $c_{ab}(G) = c_{de}(G) = c_{hj}(G) = c_{jh}(G) = 1$. Let $c_{bd}(G) = x$ and $c_{ba}(G) = 1-x$. Similarly, let $c_{eh}(G) = y$ and $c_{ed}(G) = 1-y$. Observe that, since all transactions are of size one and G 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). We can easily 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 renders G a Nash equilibrium (see Figure 2). □

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.

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

PROOF. Consider a game with four agents: $V = \{a, b, d, e\}$. The graph H is a line graph over nodes in V with edges $(a, b), (b, d)$, and (d, e). For each node $u \in V$, $B_u = 1$. The nonzero transaction rates are given by $\lambda_{ab} = \lambda_{ba} = \lambda_{de} = \lambda_{ed} = \lambda_1 > 0, \lambda_{ae} = \lambda_{ea} = \lambda_2 \gg \lambda_1$. All other entries in the transaction rate matrix Λ are zero. All transactions are of size one.

Consider the network G shown in Figure 3(a). Observe that we can use Lemma 3.7 to compute the steady-state transaction success probabilities between nodes and verify that G is a Nash equilibrium. The overall social welfare $U(G)$ in network G is given by

$$U(G) = \sum_{u \in V} U_u(G) = \sum_{u \in V} \sum_{v \in V} p_{uv}(G) = 2p_{ab}(G) + 2p_{de}(G) = 2\lambda_1 \frac{2}{3} + 2\lambda_1 \frac{2}{3} = \lambda_1 \frac{8}{3}.$$

Now consider the socially optimal network G^* in Figure 3(b). The social welfare $U(G^*)$ is given by

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

In the limit as $\lambda_1 \to 0$, the ratio $U(G^*)/U(G) = \infty$. □

4. NETWORK FORMATION UNDER GLOBAL RISK: SINGLE-MINDED AGENTS

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 $c_{uv}(G)$ to v loses $c_{uv}(G)$ units. Thus $\Delta_{uv}(G) = \delta_v c_{uv}(G)$.

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 G, either $c_{uv}(G) = 0$ for all $v \in V$ or there exists a single agent $w \in V$ such that $c_{uw}(G) = B_u$.

Further, we assume that: (i) the transaction rate matrix Λ is uniform, that is, for all $u, v \in V$, $\lambda_{uv} = \lambda = 1/(n(n-1))$; (ii) all transactions are size one, that is, for all $u, v \in V$ and for all $t > 0$, $x^t_{uv} = 1$; 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 PoA is ∞.

Example 4.2. Suppose that, for all $u \in V, \gamma\lambda(h + h^2) > \delta_u c > \gamma\lambda h$, where $h = c/(c+1)$. Let G be the empty network. Observe that, by Lemma 3.7, the utility to a node u from extending credit to any node v in G is $\gamma\lambda h$, which by assumption is less than $\delta_v c$. Thus G is a Nash equilibrium. On the other hand, since $\gamma\lambda(h + h^2) > \delta_u c$ 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 PoA is ∞.

The empty-network equilibrium is only one of potentially many equilibria for a particular game instance. To focus on the more interesting equilibria and to simplify analysis, we assume for the rest of this section 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. Since an agent extends credit to exactly one other, we introduce the function $\tau_s : V \to V$ to denote this "trustee": $\tau_G(u) = v$ implies $c_{uv}(G) = c$.

The following observation is a consequence of the analysis by Dandekar et al. [2011] of the steady-state success probability in trees under a unit transaction regime.

LEMMA 4.1 [DANDEKAR ET AL. 2011]. *Consider a network G. Let $u \in V$ be a node such that no node extends credit to u in G and let $\tau_s(G) = v$. Assume the transaction rate matrix Λ is uniform and G is under a unit transaction regime. Then, for any node $w \in V \setminus \{u, v\}$, $p_{uw}(G) = hp_{vw}(G)$, where $h = c/(c+1)$.*

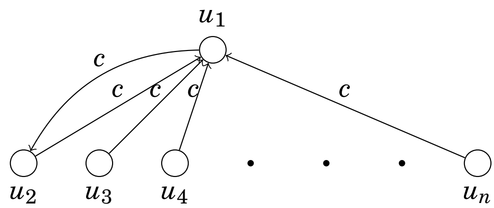

(a) socially optimal network G^*

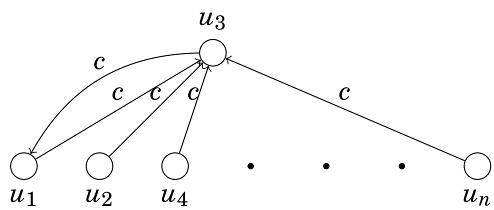

(b) Nash equilibrium G_1 where node u_3 is the root node

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

4.1. 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. Consider a network G^* such that, for all nodes $u \in V \setminus \{v^*\}$, $\tau_{G^*}(u) = v^*$, and $\tau_{G^*}(v^*) = \arg\min_{v \in V\setminus\{v^*\}} \delta_v$. Then, G^* is a Nash equilibrium and maximizes social welfare.*

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

Despite ruling out the empty network as a Nash equilibrium, PoA 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. Assume, without loss of generality, that for nodes $u_1, \ldots, u_n \in V$, $\delta_{u_1} \leq \cdots \leq \delta_{u_n}$. Let $\delta_{u_1} c = \gamma\lambda(n-3)h^2 2c/(2c+1)$, and $\delta_{u_2} = \delta_{u_3} = \gamma\lambda(n-3)h^2$, with $h = c/(c+1)$. Consider the network G^* in Figure 4(a). It follows from Lemma 4.2 that G^* is a socially optimal network. Consider the network G_1 in Figure 4(b). Observe that Lemma 3.7 can be used to compute the steady-state transaction success probabilities, and hence the utilities, of all nodes in G_1. Since $c(\delta_{u_3} - \delta_{u_1}) \leq (n-3)\gamma\lambda\frac{h^2}{2c+1}$, nodes in G_1 cannot benefit from extending credit to u_1 or u_2 instead of u_3. Thus G_1 is a Nash equilibrium. Note that, since G^* and G_1 are structurally identical

$$
\begin{aligned}
&\sum_{u,v} p_{uv}(G^*) = \sum_{u,v} p_{uv}(G_1),\\
&= \lambda(n-2)\left((n-3)h^2 + 2h\frac{2c}{2c+1} + 2h\right) + 2\lambda\frac{2c}{2c+1},\\
&= \lambda(n-2)(n-3)h^2 + \Theta(n).
\end{aligned}
$$

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

$$U(G^*) = \gamma \sum_{u,v} p_{uv}(G^*) - (n-1)\delta_{u_1}c - \delta_{u_2}c,$$
$$= \gamma\lambda(n-3)h^2\left((n-2) - (n-1)\frac{2c}{2c+1}\right) + \Theta(n) = \Theta(n^2).$$

On the other hand,

$$U(G_1) = \gamma \sum_{u,v} p_{uv}(G_1) - (n-1)\delta_{u_3}c - \delta_{u_1}c,$$
$$= \gamma \sum_{u,v} p_{uv}(G_1) - \gamma\lambda(n-1)(n-3)h^2 - \delta_{u_1}c = \Theta(n).$$

Since the price of anarchy is lower-bounded by $U(G^*)/U(G_1)$, we have that $\text{PoA} = \Omega(n)$. □

4.2. Dynamics of Network Formation

Despite the arbitrarily high PoA, we demonstrate that greedy dynamics can quickly converge to a socially optimal network.

4.2.1. Greedy Response. For network G and an agent u, we define *greedy response* by u as follows: let $v_u^* \in \arg\min_{v \in V\setminus\{u\}} \delta_v$ be a node with the lowest default probability among all nodes except u. Then, u's greedy response is to extend credit to v_u^*, that is, $\tau_{G'}(u) = v_u^*$, where G' is the network resulting from u's greedy response in G. For nodes $w, y \in V$,

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

THEOREM 4.5. *Assume that the default probabilities δ_v, $v \in V$ are all distinct. Let G^* be the network obtained after all agents have played greedy response, starting from arbitrary G. Then G^* maximizes social welfare.*

PROOF. Since the default probabilities are all distinct, there exists a unique v^* with minimum default probability, and node $v_{v^*}^*$ with the second-lowest default probability. Then, observe that for all $u \in V \setminus \{v^*\}$, $\tau_{G^*}(u) = v^*$ and $\tau_{G^*}(v^*) = v_{v^*}^*$, and G^* is exactly the credit network established as optimal in Lemma 4.2. □

4.2.2. Sequential Arrival. We consider a model where agents arrive sequentially, and decide which among the agents already in the network to which to extend credit. Let G_0 be a network of two agents, say u_0 and v_0, such that $\tau_{G_0}(u_0) = v_0$ and $\tau_{G_0}(v_0) = u_0$. At each time $t = 1, 2, \ldots$, an agent u_t arrives and extends credit to one of agents in the network G_{t-1} in order to maximize $U_{u_t}(G_t)$, where G_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 $\tau_{G_{t-1}}(u_{t-1})$.

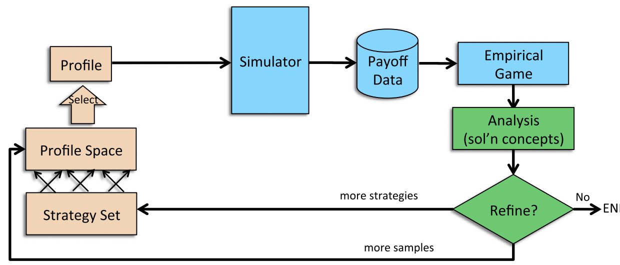

Fig. 5. Iterative procedure for empirical game-theoretic analysis.

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

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

5. SIMULATION ANALYSIS OF CREDIT NETWORK FORMATION

Covering more relaxed scenarios such as those with more flexible Λ, unconstrained budgets, multiple credit issuance, or graded risk has thus far proved elusive for analytic treatments. Several factors contribute to the difficulty of game-theoretic analysis of this problem. First, the strategy space is combinatorial and multidimensional. Strategies for this game are mappings from all the information an agent has about the environment (default probabilities for all other agents, probabilities, and values of transactions with all other agents) to all possible credit assignments to the other agents. Second, the expected value to an agent of a credit assignment is defined in terms of the outcome of a stochastic transaction sequence, intermixed with adjustments of credit balances that have important but indirect effects on the probabilities of downstream transactions.

For this reason, we employ simulation to analyze environments that relax the conditions for which we have theorems. Simulation handily deals with the complex stochastic and dynamic factors bearing on the evaluation of credit-issuing strategies. Although analysis of simulation-based games is inherently limited by the set of strategies considered, a systematic exploration of well-motivated heuristics can provide an illuminating complement to theoretical analysis.

5.1. Empirical Game-Theoretic Analysis

Our investigation of the credit network formation game employs an approach called *empirical game-theoretic analysis* (EGTA) [Wellman 2006]. In EGTA, techniques from simulation, search, and statistics combine with game-theoretic concepts to characterize strategic properties of a domain.

5.1.1. Iterative EGTA Process. A high-level view of the EGTA process is presented in Figure 5. We start with an enumerated set of strategies, typically heuristics derived from domain knowledge or experience, often parametrized by meaningful strategy features. The basic EGTA step is simulation of a strategy profile, determining a payoff observation (i.e., a sample drawn from the outcome distribution generated by the simulation environment), which gets added to the database of payoffs. Based on the accumulated data, we induce an empirical game model. On this model we may perform any of the standard computations applied to game forms (e.g., identifying dominated strategies, finding equilibria). Based on the results, we may choose to refine the model by considering more strategies or strategy profiles, or by obtaining more samples of profiles already evaluated.

The most straightforward way to define the empirical game is simply to estimate payoffs for evaluated profiles by their sample mean.[2] We employ this method for the baseline game model, but then produce an approximate *reduced game* model as well, by the technique described in Section 5.1.2.

When games (such as our version of credit network formation) exhibit significant symmetry, this can be exploited in representation and reasoning. Even for a fully symmetric game, however, the profile space grows exponentially with the lesser of number of players n and number of strategies m. There are $\binom{n+m-1}{n}$ distinct profiles, to be precise. For even moderate n and m, therefore, we generally cannot afford to evaluate every profile through simulation. We thus require analysis techniques that operate on incompletely specified games. Within a game that is incomplete due to unevaluated strategy profiles, it is often useful to work with *complete subgames*, defined by sets of strategies over which we have evaluated all profiles.

The EGTA process we followed in this study can be described in terms of two key procedures. The first, termed the *EGTA inner loop*, searches for an equilibrium within a fixed strategy set $\mathcal{S}$.[3] The second, *outer loop*, extends the strategy set through local search, implementing the selection of more strategies depicted in Figure 5.

The EGTA inner loop starts by performing an initial set of simulations covering all profiles over a small subset $\mathcal{S}_0 \subset \mathcal{S}$. It then iterates the following steps.

(1) Identify the *maximal complete subgames*, $\{\mathcal{S}^1, \dots\}$. A complete subgame is maximal if adding any strategy would render it incomplete.
(2) For each maximal complete subgame $\mathcal{S}^i$, search for *symmetric mixed-strategy Nash equilibria* (SMSNE). We employ replicator dynamics [Gintis 2000, Chapter 9] for this purpose, from a diverse set of starting points. Let σ_j^i denote the jth SMSNE found for subgame $\mathcal{S}^i$. These subgame SMSNE are *candidate* equilibria for the full game over $\mathcal{S}$.
(3) For each σ_j^i, check the strategies $s' \in \mathcal{S} \setminus \mathcal{S}^i$ such that we have evaluated all profiles where one player plays s' and the other $n-1$ play strategies in the support of σ_j^i. For any pair (σ_j^i, s') where s' is a beneficial deviation, the candidate σ_j^i is *refuted*. For instances σ_j^i such that all possible deviating strategies have been evaluated without refutation, we say that σ_j^i is *confirmed*.
(4) If there remains an SMSNE candidate σ_j^i that is neither refuted nor confirmed, simulate the profiles necessary to check another strategy $s'' \in \mathcal{S} \setminus \mathcal{S}^i$ not yet fully evaluated in context σ_j^i, and repeat from step 1.

[2]More sophisticated approaches may generalize from simulation data using regression or other machine learning techniques [Jordan and Wellman 2009; Vorobeychik et al. 2007].
[3]Our description here assumes the game is symmetric, but the procedure can be straightforwardly extended to general *role-symmetric* games [Wellman et al. 2013].

(5) If there exists a refuted SMSNE candidate σ_j^i, such that the support of σ_j^i plus its best-response refuting strategy is not subsumed by any complete subgame, simulate the profiles necessary to complete this subgame and repeat from step 1.
(6) If there exists at least one confirmed SMSNE candidate σ_j^i, return. Otherwise, choose a subgame i, extend it with some strategy $s' \in \mathcal{S} \setminus \mathcal{S}^i$, and repeat from step 1.

On termination, the empirical game is in a state where all SMSNE candidates are confirmed and all maximal subgame best responses are themselves in a completed subgame. As long as the operation of identifying subgame equilibria is complete, the procedure is guaranteed to identify at least one confirmed SMSNE candidate.

The outer loop takes as input a confirmed SMSNE from the inner loop and attempts to find a new strategy $s' \notin \mathcal{S}$ that refutes this SMSNE. It assumes a parametrically structured strategy space and performs local search in this space from a given starting point and subject to given constraints. We describe how this was implemented for the credit network formation game next; the general procedure is presented in detail elsewhere [Wellman et al. 2013].

5.1.2. Deviation-Preserving Reduction. One of the virtues of credit networks is their ability to support transactions among nodes only indirectly related by paths of credit. This property is particularly advantageous for large populations, where directly connecting all pairs that might transact would be too unwieldy. Our analysis of strategic network formation therefore requires a sufficiently large number of agents to reap the benefits of distributed credit allocation.

Increasing the number of agents, however, tends to blow up the profile space. For example, with 61 players (the number of nodes considered in this study), even a subgame of three strategies requires 1,953 profiles to complete, and four strategies requires 41,664. It would not be feasible to explore very many subgames at this population size. We therefore seek to approximate the 61-player game by a smaller game. We call this approach *player reduction*, and in prior work employed a hierarchical approach where each player in the reduced game controls a proportional number of players in the full game [Wellman et al. 2005].

In the present study, we employ a recently introduced technique called *deviation-preserving reduction* (DPR) [Wiedenbeck and Wellman 2012]. DPR is motivated by the assumption that an agent's payoff is sensitive to its own choice of strategy and to the strategies of its opponents in the aggregate, but that small numbers of opponents changing strategy can be ignored. To calculate the payoff of a player i for a profile in the reduced DPR game, we consider the full-game profile where one player plays i's designated strategy, and the remaining players are divided proportionally among the other strategies in the reduced game profile.

Specifically, we focus analysis on a six-player reduced game derived from simulations on 61-agent credit networks, depicted in Figure 6. For example, we construct the six-player DPR profile $\langle 1 \times s_1, 3 \times s_2, 2 \times s_3 \rangle$ where one player plays strategy s_1, three play s_2, and two play s_3, by simulating three 61-agent profiles. The payoff to the player playing s_1 comes from the full-game profile $\langle 1 \times s_1, 36 \times s_2, 24 \times s_3 \rangle$, the payoff for s_2 from $\langle 12 \times s_1, 25 \times s_2, 24 \times s_3 \rangle$, and for s_3 from $\langle 12 \times s_1, 36 \times s_2, 13 \times s_3 \rangle$. In effect, each reduced game player views itself as controlling one full-game agent, while its reduced game opponents represent the fraction of full-game opponents playing each strategy. By this description, we see that deviation-preserving reduction applies most straightforwardly when the reduced game size divides $n - 1$ (hence our choice

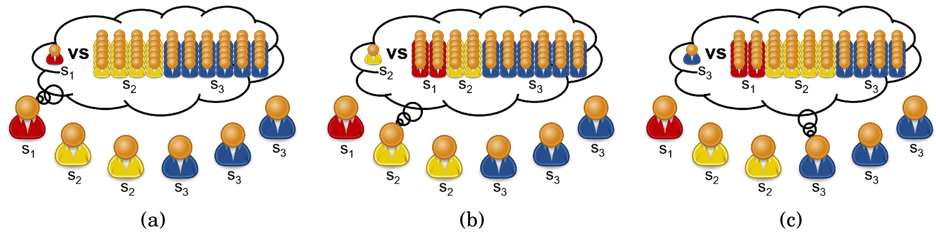

Fig. 6. Deviation-preserving reduction example. Each of the six reduced game players views itself as controlling one of the 61 full-game agents while each opponent controls an equal fraction (12) of the remaining full-game agents. The payoff to the player playing s_1 in the reduced game profile comes from the payoff to s_1 in the full-game profile depicted in (a); s_2 from (b); s_3 from (c).

of $n = 61$ for this study). The DPR technique, however, is defined more generally for nondivisible reduction factors, as well as for games that are symmetric only within roles [Wiedenbeck and Wellman 2012].

5.2. Credit Network EGTA Study: Setup

As noted previously, we consider a population of 61 agents. We performed EGTA analyses for 12 environments differing in risk model, default prevalence, and transaction value. Each simulation run evaluates a profile of heuristic strategies for issuing credit. The 61 agents apply their assigned strategies to issue credit based on their available information, forming an initial credit network. The simulation processes probabilistic defaults and a stochastic sequence of transactions to generate sample payoffs for the strategy profile. Strategies are selected from a suite of heuristics which implement a variety of criteria for issuing credit, parametrized by how many agents to issue credit and how much. We drive the choice of profiles to simulate according to the EGTA process described before (Section 5.1), producing empirical game models for each of the 12 environments.

Each run of the scenario comprises 10,000 transaction request events. The transaction rate λ_{uv} for each pair of agents $u \neq v$ is drawn uniformly and then normalized. All transaction requests from u the *buyer* to v the *seller* are for a single unit. The value to u of a successful transaction with v is drawn uniformly, $x_{uv} \sim U[1, \bar{x}]$, with $\bar{x}$ set to either 1.2 (*low value*) or 2 (*high value*). The buyer pays a constant price of one to the seller (i.e., expends one unit of credit), so the average buyer surplus per transaction is thus either 0.1 or 0.5. We assume the seller makes no profit, but note that v benefits by accruing credit it may subsequently use in transactions as a buyer.

Default probabilities δ_v for each agent are drawn from a beta distribution: $Beta(1, 9)$ (average default probability $\frac{1}{10}$) in the *low-default* setting, $Beta(1, 2)$ (average $\frac{1}{3}$) in the *medium-default* setting, and $Beta(1, 1)$ (average $\frac{1}{2}$) in the *high-default* setting. In the global risk environment these default probabilities are revealed to all, whereas in the graded risk environment each agent gets sample data from the default distribution of others, with the number of samples ∂_{uv} determined by the social network distance between u and v. The social network itself is an Erdös-Rényi graph. We take $\partial_{uv} = 100$ if u and v are neighbors, $\partial_{uv} = 10$ if they are linked through one other node, $\partial_{uv} = 1$ if they have a shortest path of length three, and $\partial_{uv} = 0$ otherwise.

We explored environments with high, medium, or low default, and high or low value, for each of global and graded risk. The 12 environments are listed in Table I, along with the number of profiles and strategies we ended up simulating in both the full and reduced games. Three-letter environment names are coded by risk

Table I. Exploration Performed by the Iterative EGTA Process under Various Environment Settings

name	Risk model	Default prob	buyer surplus	Full-game profiles (I/II)		DPR profiles (I/II)		Strategies (I/II)	
CLL	Global	low	low	4619	11695	1497	3946	17	32
CLH	Global	low	high	2179	9861	765	3359	17	32
CML	Global	med	low	3557	9425	1036	3213	8	32
CMH	Global	med	high	8619	20192	2622	6474	15	32
CHL	Global	high	low	3134	5901	1045	2090	17	32
CHH	Global	high	high	3202	6155	1101	2196	17	32
ILL	Graded	low	low	991	9322	394	3148	17	32
ILH	Graded	low	high	5377	28721	1824	8786	17	32
IML	Graded	med	low	1766	8927	565	3109	8	32
IMH	Graded	med	high	24612	39728	6818	11356	18	32
IHL	Graded	high	low	656	5080	261	1881	17	32
IHH	Graded	high	high	430	10529	201	3516	17	32

Strategies gives the number of strategies added by the outer loop. *Full-game profiles* gives the number of 61-agent profiles sampled by the inner loop. *DPR profiles* gives the number of 6-player profiles in the empirical game model.

model (C[omplete information] for global risk, I[ncomplete] for graded risk), default probability (L[ow]/M[edium]/H[igh]), and buyer value (L[ow]/H[igh]). These numbers of profiles and strategies are broken down by two stages (I and II) of search, as described next.

A heuristic strategy is defined by three parameters: (i) a criterion for ranking the other agents, (ii) the number k of agents to issue credit (the best k according to the ranking criterion), and (iii) the number of units q of credit to issue to each of these chosen agents. The criteria we included in heuristics along with the (k, q) values we considered in this study are enumerated as follows, defined from the perspective of agent u's evaluation of credit prospect v.

— *Default probability*. This is the lowest known default (δ_v) for global risk, or lowest estimated default based on samples ∂_{uv} for graded risk.
— *Buy rate*. This is the highest probability of transacting (λ_{uv}).
— *Sell rate*. It is the highest probability of serving a transaction (λ_{vu}).
— *Trade value*. This is the highest expected value of transaction per event ($\lambda_{uv}x_{uv}$).
— *Trade profit*. It is the highest difference, expected value of transaction minus expected value of served transaction ($\lambda_{uv}x_{uv} - \lambda_{vu}$).
— *Index*. This is the lowest node number (arbitrary global labeling).
— *Random*. This is uniform choice.

In addition, we included the no-credit strategy, Zero, which issues no credit to anyone.

Observe that the Default strategies behave qualitatively differently in the global and graded risk environments. Under global risk, all agents have the same information about default probabilities. Therefore, when agents issue credit to the least-likely

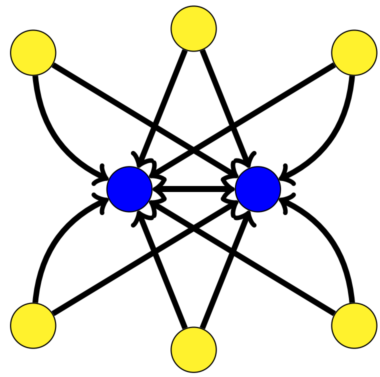

Fig. 7. A star-like credit network similar to one that would arise from the mutual application of default-based strategies in the global risk model.

Table II. Strategies Included in the EGTA Study

Criterion	Predefined (k, q)	Automatically Generated (k, q)
Default	$(1,1), (2,2), (3,2), (5,2)$	$(3,1), (4,1), (5,1), (6,1), (8,1), (9,1)$
Buy rate	$(1,1), (1,2), (2,2)$	$(2,1), (4,1), (8,1), (10,1)$
Sell rate	$(2,2)$	$(6,1)$
Trade value	$(2,2), (5,2)$	$(2,1), (3,1)$
Trade profit	$(2,2), (6,2), (8,1)$	$(3,2), (5,1)$
Index	$(1,1), (2,2)$	
Random	$(2,2)$	
Zero	$(0,0)$	

These are predefined for Stage I, and automatically generated in the outer loop.

defaulters, they are creating edges to the same target agents. This leads to a centralized or star-like credit network as illustrated in Figure 7. Such coordination on credit targets has potential advantages. If everyone including u offers credit to v, then once v transacts with u, u enjoys credit paths to everyone in the network. This coordination does not result, in contrast, from mutual application of Default in the graded risk model. Under graded risk, agents have different information based on their positions in the social network. The counterparts judged to have lowest default probably are invariably those with whom the agent has had most positive experience. Since there is little experience of any kind with social strangers, these are unlikely to be judged most trustworthy (this happens only if one is unlucky enough to have only very untrustworthy friends). Indeed, under graded risk we find that 95% of the top five least likely defaulters are one or two hops away on the social network. Finally, note that the Index strategies do coordinate on a star-like network, in either the global or graded risk model. Since both Default and Index achieve coordination in the global risk model but only the former minimizes default, by comparing the two strategies we can separate the pure benefits of coordination from the benefits of avoiding defaulters.

In Stage I of the analysis, we considered a fixed set of 17 strategies and ran the inner loop on eight of the twelve environments: those with high- or low- (not medium-) default probabilities. The 17 predefined strategies were selected in an ad hoc manner through exploration in a preliminary study and are enumerated in Table II.

For the four-medium default environments, we started with a smaller set of eight predefined strategies (the first listed for each criterion) and employed the automated strategy generation procedure (EGTA outer loop) to extend the set. On each iteration, we searched for refutations of an equilibrium σ confirmed for the existing strategy set, employing local search from a particular existing strategy. The search algorithm simply hill-climbs from the existing strategy, holding its credit criterion fixed but incrementing or decrementing its k and q parameters. The search halted upon reaching a local maximum in payoff, assuming all other nodes in the network play according to σ. If this local maximum exceeds the equilibrium payoff, the new strategy is added to the set and we proceed with another round of the EGTA inner loop. If instead we cycle through the strategy categories (in this case, defined by credit criterion) without finding a beneficial deviation, the entire process concludes.

As indicated in Table I, the CML and IML environments found no new strategies, whereas the CMH environment added seven strategies to the original eight, and IMH added ten. Together, there were 15 automatically generated strategies not included among the 17 predefined Stage-I strategies. These are listed in the final column of Table II. For Stage II, we constructed the combined set of 32 strategies and ran the EGTA inner loop for each environment with this set.

With 17 strategies, there are 1.4×10^{16} distinct strategy profiles for the full 61-player game, and 74,613 for the six-player DPR game. These numbers grow to 3.0×10^{24} (61-player DPR) and 2,324,784 (six-player DPR) for the Stage-II set of 32 strategies. As indicated in Table I, our EGTA process evaluated only a very small fraction of these profiles at each stage. Nevertheless it was able to identify equilibria in each environment.

Altogether we evaluated 165,536 full-game strategy profiles across the 12 credit network environments, from which we estimated payoffs for 53,074 DPR profiles. Each full-game profile evaluated was simulated at least 1,000 and usually upwards of 2,000 times. Our simulations were performed on a computing cluster operated by the University of Michigan, using an experiment management facility designed expressly for EGTA studies [Cassell and Wellman 2013].

5.3. Results

Through the process described in Section 5.1.1, we successfully derived equilibria for each of the 12 credit network games. Specifically, we identified between one and six SMSNEs for the reduced six-player DPR games corresponding to each environment. All candidate subgame equilibria were either confirmed or refuted by the process, and the subgames covering best responses to all candidates were completed.

The strategies Sell rate, Index, and Random are not supported in any equilibria. This confirms our intuitions: substantive criteria are better than random; coordination based on default risk makes more sense than index coordination; and buyer transaction rate is more directly reflected in value than seller rate.

To characterize the equilibria qualitatively, we partition the remaining strategies as follows. Class **D** represents Default, **Z** represents Zero, and **T** groups together strategies based on criteria related to transaction probability and value: Buy rate, Trade value, and Trade profit. The SMSNEs identified are summarized in Figure 8. In the figure, there is one cell for each environment, displaying class labels for strategies supported in some equilibrium. A class letter circled means that a strategy in this class was confirmed as a pure-strategy Nash equilibrium. Interestingly, whereas many of the equilibria found were mixed and several environments had equilibria in multiple classes, in no case did a single SMSNE mix across the class partitions defined earlier.

From the figure, we see that there is a no-credit equilibrium in eight of the twelve environments: all but those with low or medium default and high buyer value. The

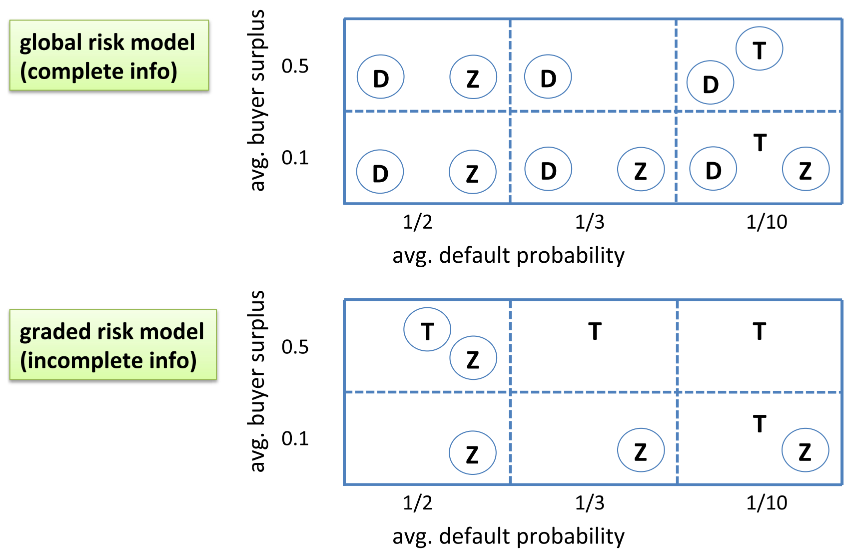

Fig. 8. Equilibria found for the 12 credit network environments. Letters denote the strategy classes represented in equilibrium, with circled letters indicating pure-strategy equilibria.

two least favorable environments—graded risk with high or medium default and low value—have only this equilibrium, whereas all the others have some equilibrium where credit is provided. All of the global risk environments have an equilibrium where everybody plays Default, but this strategy does not appear in equilibrium for any graded risk environments. Indeed, there is a one-to-one correspondence between the equilibria for the two risk classes, except that the graded risk environments omit these Default equilibria, and when buyer value is high, these are replaced by transaction-based equilibria. The weakened information about defaults plus the lack of coordinating power render this a poor credit-issuing criterion in graded risk environments.

For completeness, we list the equilibria found. Groups in brackets with probabilities represent SMSNEs, and ungrouped strategies indicate pure equilibria.

CLL. Default(1,1); Default(3,1); Default(4,1); Zero; [Buy rate(2,1), 0.899; Trade value(3,1), 0.101]; [Trade value(2,1), 0.806; Trade value(3,1), 0.194]
CLH. Default(3,2); Trade profit(5,1); [Default(6,1), 0.951; Default(8,1), 0.049]; [Default(5,1), 0.744; Default(8,1), 0.256]
CML. Default(1,1); Default(3,1); Zero
CMH. Default(2,2); [Default(2,2), 0.014; Default(6,1), 0.986]; [Default(3,2), 0.880; Default(4,1), 0.120]; [Default(5,1), 0.821; Default(6,1), 0.179]
CHL. Default(1,1); Default(3,1); Zero
CHH. Default(2,2); Zero; [Default(3,2), 0.081; Default(4,1), 0.919]
ILL. Zero; [Trade value(2,1), 0.637; Trade value(3,1), 0.363]
ILH. [Buy rate(4,1), 0.229; Trade profit(5,1), 0.771]
IML. Zero
IMH. [Buy rate(4,1),0.172; Trade profit(5,1),0.785; Trade value(3,1),0.042]
IHL. Zero
IHH. Buy rate(4,1); Zero

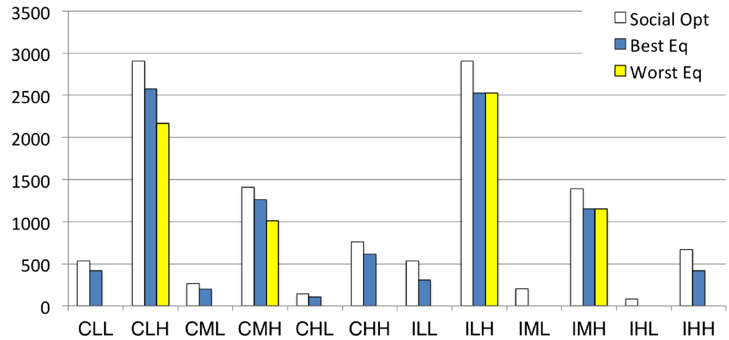

Fig. 9. Welfare at social optimum compared to welfare at equilibrium.

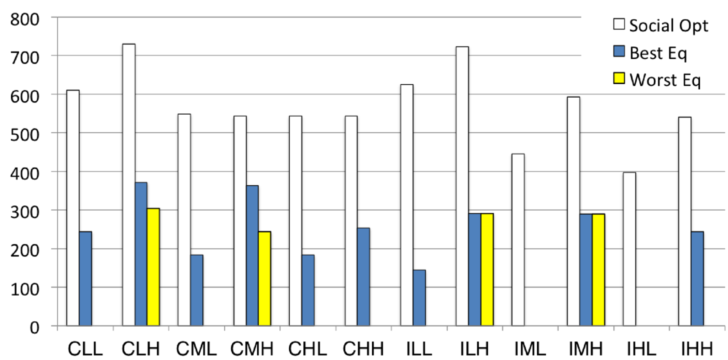

Fig. 10. Total units of credit issued at social optimum compared to equilibrium.

Whereas the set of equilibria evolved as the empirical game was refined from Stage I to Stage II, the qualitative categories of strategy profiles represented in equilibrium (as depicted in Figure 8, ignoring the circle designations) remained constant.

We next turn to the question: How well do the credit networks generated in equilibrium perform? Figure 9 compares the welfare (sum of agent utility) of equilibrium outcomes to that of an estimated social optimum. Our estimate is actually a lower bound, equal to the greatest social welfare seen in any full-game profile simulated. Equilibrium welfare varies across equilibria, hence we present the best and worst of those identified. In eight of twelve environments, the worst is the Zero equilibrium, which supports no transactions and thus yields zero welfare. What we find overall is that, when there is a substantial amount of welfare possible (i.e., the most favorable environments), equilibrium network formation does a good job of obtaining most of it. For less favorable environments, a network (if it forms at all) tends to produce little utility.

We can also observe directly the amount of credit issued in equilibrium networks as compared to the social optimum—which is not necessarily the credit-maximizing network. As seen in Figure 10, the comparison mirrors that for welfare, but with lower ratios of equilibrium to social optimum across the board. This is due to the diminishing returns to credit once the network has ample credit capacity. In other words, we can achieve a substantial fraction of available social welfare without issuing this same fraction of the credit that a social planner would.

As acknowledged at the outset, all of these results are relative to the particular strategy space included in the empirical game analysis. Our choice was driven by an effort to span a diverse space and to include strategies successful in preliminary studies or otherwise representing plausible prospects for refuting initial equilibrium candidates. The fact that adding strategies in Stage II based on automated exploration of parametric variations on the original strategies did not change the qualitative character of equilibria lends support to the robustness of these results.

6. 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.

In the most restrictive case of dichotomous risk with only bilateral transactions permitted, we show that the formation game is a potential game and, under many transaction size distributions, 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* 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 PoA 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 PoA is unbounded, simple greedy dynamics quickly converge to a social optimum. Even when agents are allowed to extend credit to multiple agents, we show using empirical game simulation that nonempty equilibria tend to be star like.

Our empirical game simulations confirm the finding of star-like equilibrium networks under global risk, even in less restrictive scenarios. In addition, we study a graded risk model where agents have partial information about default risks. We find that star-like equilibria disappear because agents are unable to coordinate on highly trustworthy central nodes. We also find that whether empty networks can occur in equilibrium depends primarily on the relative profitability of transactions, and not on the structure of information about default probabilities.

APPENDIXES

A. TABLE OF SYMBOLS

Table III summarizes the notation we employ in this article.

Table III. Notational Symbols and their Interpretation

Symbol	Meaning
$u, v, \ldots$	Agents; nodes in a credit network
V	Set of agents/nodes
n	Number of agents/nodes
c_{uv}	Credit on edge from u to v
B_u	Credit budget for agent u
G	Credit network
t	Current time step[4]
λ_{uv}	Probability that u is chosen to transact with v
Λ	Transaction rate matrix
x^t_{uv}	Transaction size between u and v
$X_{uv}(\cdot)$	Probability distribution over transaction sizes
$\mathcal{X}_{uv}(\cdot)$	Cumulative version of $X_{uv}(\cdot)$
$\mathbb{X}$	Matrix of transaction-size distributions
$\langle \Lambda, \mathbb{X} \rangle$	Transaction regime
$\mathcal{P}$	Path in the credit network
$\mathcal{P}_{uv}$	Shortest feasible path between u and v
$\mathcal{M}(G, \Lambda, \mathbb{X})$	Markov chain over the states of the network
$U_u(G)$	Utility (payoff) of u for initial network G
$U(G)$	Social welfare for initial network G
$\Delta_{uv}(G)$	Expected loss of utility to u from prospect of default by v
$p_{uv}(G)$	Steady-state success probability of transactions from u to v
γ	Factor converting p to utility units
$H = (V, E)$	Social network, with edges E
E_u	Set of edges in E incident on node u
δ_v	Probability that agent v defaults
δ_{uv}	Signal agent u receives about probability v defaults
$Beta(\alpha, \beta)$	Beta distribution with parameters α and β
∂_{uv}	Number of samples generating signal δ_{uv}
$\Phi(G)$	Potential function
r_{uv}	Marginal utility of edge (u, v)
r^*_u	Maximum marginal utility among edges incident on u
E^*_u	Edges incident on u maximizing marginal utility
$r^{\min}_G$	Minimum marginal utility among edges of G
$I_{(G,F)}$	(G, F)-restriction of setup instance I
$d_u(G)$	*Generalized score*: Total credit issued to u in G
c	Constant integer credit issued in unit transaction regime
$\tau_G(u)$	The agent issued credit by single-minded agent u
v^*	Node with minimum default probability
v^*_u	Node other than u with minimum default probability
x_{uv}	Value to u of a successful transaction with v

[4]Symbols denoting time-varying objects may be superscripted with time; omitted superscripts indicate $t = 0$ (initial value).

B. PROOFS OF THE BILATERAL TRANSACTION REGIME

B.1. Proof of Lemma 3.1

Recall from Section 2 that the total credit capacity along edge $e = (u, v)$ remains constant over time, that is, for all $t > 0$, $c_{uv}(G^t) + c_{vu}(G^t) = c_{uv}(G) + c_{vu}(G) = c_e(G)$. The repeated probabilistic transactions along edge e induce a Markov chain over $[0, c_e(G)]$ which is governed entirely by λ_{uv} and $X_{uv}(\cdot)$. We denote the Markov chain along edge e by $\mathcal{M}_e(\lambda_{uv}, X_{uv}) = \{Y_e(t) \mid t \geq 0\}$. A state $Y_e(t) = y$ of the Markov chain represents the current division of total credit capacity across the two directions, so can be encoded by $c_{uv}(G^t) = y, c_{vu}(G^t) = c_e(G) - y$. Let P_e be the transition kernel of the Markov chain and let ρ_e be the corresponding density. Note that since $X_{uv}(\cdot)$ has support on $[0, \infty)$, for all $0 \leq y, z \leq c_e(G)$, $\rho_e(y, z) > 0$. This implies that $\mathcal{M}_e$ is irreducible.

PROPOSITION B.1. *The Markov chain $\mathcal{M}_e(\lambda_{uv}, X_{uv})$ is irreducible.*

PROOF. If $c_e(G) = 0$, $\mathcal{M}_e(\lambda_{uv}, X_{uv})$ is trivially irreducible. If $c_e(G) > 0$, observe that, for any $y \in [0, c_e(G)]$ and any closed interval $A \subseteq [0, c_e(G)]$ of nonzero length, $P_e(y, A) > 0$. □

Next we show that $\mathcal{M}_e(\lambda_{uv}, X_{uv})$ has a uniform steady-state distribution over $[0, c_e(G)]$.

PROPOSITION B.2. *The Markov chain $\mathcal{M}_e(\lambda_{uv}, X_{uv})$ has a uniform steady-state distribution over $[0, c_e(G)]$.*

PROOF. The statement is trivially true if $c_e(G) = 0$. Assume $c_e(G) > 0$. Consider a uniform distribution given by π_e such that $\pi_e(y) = 1/c_e(G)$ for all $y \in [0, c_e(G)]$. Then, π_e satisfies the detailed balance condition, $\pi_e(y)\rho_e(y, z) = \pi_e(z)\rho_e(z, y)$ for all $0 \leq y, z \leq c_e(G)$. Therefore, π_e is a steady-state distribution of $\mathcal{M}_e$. Further, since $\mathcal{M}_e$ is irreducible (by Proposition B.1), π_e is the unique steady-state distribution of $\mathcal{M}_e$. □

Now we are ready to characterize the steady-state transaction probability $p_{uv}(G)$ between nodes u and v. Fix a time step t. Assume that $Y_e(t-1) = z$, where $0 \leq z \leq c_e(G)$. Then the transaction at time step t where u pays v succeeds if $y^t \leq c_e(G) - z$. So the success probability at time step t is given by $\lambda_{uv}\mathbb{P}[y^t \leq c_e(G) - z] = \lambda_{uv}\mathcal{X}_{uv}(c_e(G) - z)$. Therefore, the steady-state success probability is given by

$$p_{uv}(G) = \int_0^{c_e(G)} \pi_e(z)\lambda_{uv}\mathcal{X}_{uv}(c_e(G) - z)dz.$$

When $c_e(G) = 0$, the integral evaluates to zero and therefore $p_{uv}(G) = 0$. On the other hand, when $c_e(G) > 0$,

$$p_{uv}(G) = \frac{\lambda_{uv}}{c_e(G)} \int_0^{c_e(G)} \mathcal{X}_{uv}(c_e(G) - z)dz.$$

since $\pi_e(z) = 1/c_e(G)$ for all $0 \leq z \leq c_e(G)$. Substituting $y = c_e(G) - z$, we get

$$p_{uv}(G) = -\frac{\lambda_{uv}}{c_e(G)} \int_{c_e(G)}^0 \mathcal{X}_{uv}(y)dy = \frac{\lambda_{uv}}{c_e(G)} \int_0^{c_e(G)} \mathcal{X}_{uv}(y)dy.$$

B.2. Proof of Lemma 3.4

B.2.1. Proof that p_{uv} is Continuously Differentiable and Strictly Increasing. Recall from (3) that for an edge $e \in E$, p_{uv} is given by

$$p_{uv}(x) = \begin{cases} \frac{\lambda_{uv}}{x} \int_0^x \mathcal{X}_{uv}(y)dy, & \text{if } x > 0 \\ 0, & \text{if } x = 0 \end{cases}.$$

Recall that $\mathcal{X}_{uv}(\cdot)$ is twice differentiable, which implies $p_{uv}(x)$ is continuously differentiable in $(0, \infty)$. We show that both p_{uv} and its derivative p'_{uv} are continuous at 0. For $x > 0$,

$$\lim_{x\to 0} p_{uv}(x) = \lim_{x\to 0} \frac{\lambda_{uv}}{x} \int_0^x \mathcal{X}_{uv}(y)dy.$$

Since both the numerator and denominator in the prior expression are continuously differentiable functions over $[0, \infty)$, we can use L'Hôpital's rule to evaluate the limit. Therefore

$$\lim_{x\to 0} \frac{\lambda_{uv}}{x} \int_0^x \mathcal{X}_{uv}(y)dy = \lim_{x\to 0} \frac{\lambda_{uv}\mathcal{X}_{uv}(x)}{1} = 0 = p_{uv}(0),$$

where the first equality results from L'Hôpital's rule. Therefore p_{uv} is continuous at 0. Next we show that p'_{uv} is also continuous at 0. For $x > 0$, $p'_{uv}(x)$ is given by

$$\begin{aligned} p'_{uv}(x) = \frac{dp_{uv}(x)}{dx} &= \lambda_{uv}\left(-\frac{1}{x^2}\int_0^x \mathcal{X}_{uv}(y)dy + \frac{1}{x}\mathcal{X}_{uv}(x)\right), \\ &= \frac{\lambda_{uv}}{x}\left(\mathcal{X}_{uv}(x) - \frac{1}{x}\int_0^x \mathcal{X}_{uv}(y)dy\right). \end{aligned} \tag{6}$$

Note that

$$\begin{aligned} \int_0^x (x-y)X_{uv}(y)dy &= x\int_0^x X_{uv}(y)dy - \int_0^x yX_{uv}(y)dy, \\ &= x\mathcal{X}_{uv}(x) - \int_0^x yX_{uv}(y)dy, \\ &= x\mathcal{X}_{uv}(x) - \left(y\mathcal{X}_{uv}(y)\big|_0^x - \int_0^x \mathcal{X}_{uv}(y)dy\right) = \int_0^x \mathcal{X}_{uv}(y)dy. \end{aligned}$$

Substituting $\int_0^x (x-y)X_{uv}(y)dy$ for $\int_0^x \mathcal{X}_{uv}(y)dy$ in (6), we get

$$\begin{aligned} p'_{uv}(x) &= \frac{\lambda_{uv}}{x}\left(\mathcal{X}_{uv}(x) - \frac{1}{x}\int_0^x (x-y)X_{uv}(y)dy\right), \\ &= \frac{\lambda_{uv}}{x}\left(\mathcal{X}_{uv}(x) - \frac{1}{x}\int_0^x xX_{uv}(y)dy + \frac{1}{x}\int_0^x yX_{uv}(y)dy\right), \\ &= \frac{\lambda_{uv}}{x}\left(\mathcal{X}_{uv}(x) - \mathcal{X}_{uv}(x) + \frac{1}{x}\int_0^x yX_{uv}(y)dy\right). \end{aligned}$$

Therefore

$$p'_{uv}(x) = \frac{\lambda_{uv}}{x^2}\int_0^x yX_{uv}(y)dy. \tag{7}$$

Taking the limit as $x \to 0$, we get

$$\lim_{x\to 0} p'_{uv}(x) = \lim_{x\to 0} \frac{\lambda_{uv}}{x^2}\int_0^x yX_{uv}(y)dy.$$

Recall that, by assumption, $\mathcal{X}_{uv}$ is twice differentiable, which implies X_{uv} is differentiable. Therefore, the numerator and denominator in the expression for $p'_{uv}(x)$ are continuously differentiable functions over $[0, \infty)$. So we can again use L'Hôpital's rule to evaluate the limit. This gives

$$\lim_{x\to 0} \frac{\lambda_{uv}}{x^2}\int_0^x yX_{uv}(y)dy = \lim_{x\to 0} \frac{\lambda_{uv} x X_{uv}(x)}{2x} = \frac{\lambda_{uv}X_{uv}(0)}{2}. \tag{8}$$

Now, let us evaluate the derivative, p'_{uv}, at 0.

$$p'_{uv}(0) = \lim_{\delta\to 0} \frac{p_{uv}(0+\delta) - p_{uv}(0)}{\delta} = \lim_{\delta\to 0} \frac{p_{uv}(\delta)}{\delta}$$
$$= \lim_{\delta\to 0} \frac{\lambda_{uv}}{\delta^2} \int_0^{\delta} yX_{uv}(y)dy = \frac{\lambda_{uv}X_{uv}(0)}{2} \text{ (from (8))}. \quad (9)$$

From (8) and (9), it follows that $\lim_{x\to 0} p'_{uv}(x) = p'_{uv}(0)$, therefore p'_{uv} is continuous over $[0,\infty)$. Next we show that p_{uv} is strictly increasing over $[0,\infty)$. Recall that X_{uv} has support over $[0,\infty)$. Therefore $\mathcal{X}_{uv}$ is strictly increasing. As a result, p_{uv} is strictly increasing over $(0,\infty)$. Furthermore, $p_{uv}(0) = 0$ and p_{uv} is strictly positive over $(0,\infty)$. Therefore p_{uv} is strictly increasing over $[0,\infty)$.

B.2.2. Proof that p_{uv} is Concave. First, let us assume $x > 0$. Recall from (7) that

$$p'_{uv}(x) = \frac{\lambda_{uv}}{x^2} \int_0^x yX_{uv}(y)dy.$$

Since $X_{uv}(\cdot)$ is differentiable, we can differentiate p'_{uv} again to get

$$p''_{uv}(x) = \frac{d^2p_{uv}}{dx^2} = \lambda_{uv}\left(\frac{-2}{x^3}\int_0^x yX_{uv}(y)dy + \frac{1}{x}X_{uv}(x)\right),$$
$$= \lambda_{uv}\left(\frac{-2}{x^3}\int_0^x yX_{uv}(y)dy + \frac{2}{x^3}X_{uv}(x)\int_0^x ydy\right),$$
$$= \frac{2\lambda_{uv}}{x^3}\int_0^x y\left(X_{uv}(x) - X_{uv}(y)\right)dy.$$

Since $X_{uv}(\cdot)$ is nonincreasing (by assumption), we have that, for any $x \geq 0$, $X_{uv}(y) \geq X_{uv}(x)$ for all $y \leq x$. This in turn implies that $p''_{uv}(x)$ is nonpositive for $x > 0$, which implies $p_{uv}(\cdot)$ is concave over $(0,\infty)$. We would like to show that it is concave over its entire domain $[0,\infty)$. Consider any $x \in (0,\infty)$ and a small value $\delta > 0$. Then, since $p_{uv}(\cdot)$ is concave over $(0,\infty)$, we have that, for all $\alpha \in [0,1]$,

$$p_{uv}(\alpha\delta + (1-\alpha)x) \geq \alpha p_{uv}(\delta) + (1-\alpha)p_{uv}(x).$$

Taking the limit as $\delta \to 0$, we get

$$\lim_{\delta\to 0} p_{uv}(\alpha\delta + (1-\alpha)x) \geq \alpha \lim_{\delta\to 0} p_{uv}(\delta) + (1-\alpha)p_{uv}(x).$$

Now,

$$\lim_{\delta\to 0} p_{uv}(\delta) = \lim_{\delta\to 0} \frac{\lambda_{uv}\mathcal{X}_{uv}(\delta)}{1} \text{ (by L'Hôpital's rule)} = 0 = p_{uv}(0).$$

Therefore we have that, for any $x > 0$ and for all $\alpha \in [0,1]$,

$$p_{uv}(0 + (1-\alpha)x) \geq \alpha p_{uv}(0) + (1-\alpha)p_{uv}(x).$$

This implies that $p_{uv}(\cdot)$ is concave over $[0,\infty)$.

B.3. Proofs of Propositions in Lemma 3.6

PROOF OF PROPOSITION 3.1. By contradiction. Suppose there exists a node u such that $\sum_{v:(u,v)\in E} c_{uv}(G) < B_u$. Recall that $r_{uv}(G) > 0$ for all $(u,v) \in E$. Therefore, there must exist an edge (u,v) incident on u such that $r_{uv}(G) > 0$. So, u can allocate the remaining credit along e and improve its payoff, contradicting the assumption that G is a Nash equilibrium.

Next, suppose there exists an edge $e = (u,v) \notin E^*_u(G)$. This implies $r^*_u(G) > r_{uv}(G) > 0$. Assume that $c_{uv}(G) > 0$. If u lowers its credit allocation along e by some $\varepsilon > 0$, its utility on e changes by $\gamma(p_{uv}(c_e(G) - \varepsilon) - p_{uv}(c_G(e)))$. If u puts this ε credit on an edge $e' \in E^*_u(G)$, its utility on e' changes by $\gamma(p_{e'}(c_{e'}(G) + \varepsilon) - p_{e'}(c_{e'}(G)))$. In the limit as $\varepsilon \to 0$, $p_{uv}(c_e(G) - \varepsilon) - p_{uv}(c_e(G)) = -\varepsilon r_{uv}(G)$ and $p_{e'}(c_{e'}(e') + \varepsilon) - p_{e'}(c_{e'}(G)) = \varepsilon r^*_u(G)$. The net change in u's utility is $\gamma(\varepsilon(r^*_u(G) - r_{uv}(G))) > 0$. This again implies that G cannot be a Nash equilibrium, resulting in a contradiction. □

PROOF OF PROPOSITION 3.2. By contradiction. Assume $r^{\min}_{G'} < r^{\min}_G$.

Let $V^{\text{int}}_{G'} := V^{\min}_{G'} - V^+_{G'}$. For a node $u \in V^{\text{int}}_{G'}$, all edges incident upon u are in $E^{\min}_{G'}$. We first prove a few properties of edges in $E^{\min}_{G'}$ and nodes in $V^{\text{int}}_{G'}$.

P1. For all edges $e \in E^{\min}_{G'}$, $c_e(G') > 0$.
Proof. Consider an edge $e = (u,v) \in E^{\min}_{G'}$. We know that in state G, $r_{uv}(G) \geq r^{\min}_G$. Assume $c_e(G') = 0$. Since $r_{uv}(\cdot)$ is a strictly decreasing function, this implies $r_{uv}(t) \leq r^{\min}_{G'}$ for any network t. This contradicts the statement that $r_e(G) \geq r^{\min}_G > r^{\min}_{G'}$. Therefore $c_e(G') > 0$.

P2. For each edge $e = (u,v) \in E^{\min}_{G'}$, at least one of u and v must be in $V^{\text{int}}_{G'}$.
Proof. Consider edge $e = (u,v) \in E^{\min}_{G'}$. Assume $u \notin V^{\text{int}}_{G'}$ and $v \notin V^{\text{int}}_{G'}$. This implies that u and v each have at least one incident edge that is not in $E^{\min}_{G'}$ (by definition of $V^{\text{int}}_{G'}$). This implies $e \notin E^*_u(G')$ and $e \notin E^*_v(G')$ (by definition of $E^{\min}_{G'}$). Therefore, by Proposition 3.1, $c_{uv}(G') = c_{vu}(G') = 0$. But this contradicts P1. Therefore at least one of the endpoints of e must be in $V^{\text{int}}_{G'}$.

P3. An edge $e \in E^{\min}_{G'}$ is a slack edge if and only if exactly one of its endpoints is in $V^{\text{int}}_{G'}$.
Proof. Fix edge $e = (u,v) \in E^{\min}_{G'}$.
First, assume e is slack. It is obvious that at least one of u and v must not be in $V^{\text{int}}_{G'}$ (by definition of a slack edge). But, from P2, we know that at least one of u and v must be in $V^{\text{int}}_{G'}$. Therefore exactly one of u and v must be in $V^{\text{int}}_{G'}$.
Now assume that exactly one endpoint, say node u, is in $V^{\text{int}}_{G'}$ and $v \notin V^{\text{int}}_{G'}$. Then, there exists an edge $e' \in E - E^{\min}_{G'}$ incident on v. This implies $e \notin E^*_v(G')$ (by definition of $E^{\min}_{G'}$). Therefore e must be slack (by definition of slackness).

P4. An edge $e \in E^{\min}_{G'}$ is not slack if and only if both its endpoints are in $V^{\text{int}}_{G'}$ (this follows from P2 and P3).

P5. For all edges $e \in E^{\min}_{G'}, c_e(G') > c_e(G)$ (since $r^{\min}_{G'} < r^{\min}_G$).

P6. Let $Y \subseteq E^{\min}_{G'}$ be the set of slack edges in $E^{\min}_{G'}$. Then, for all edges $e = (u,v) \in Y$ where $u \in V^{\text{int}}_{G'}, c_e(G') = c_{uv}(G')$ (since $v \in V^+_{G'}, c_{vu}(G') = 0$).

Using these properties, we show that at least one node in $V_{G'}^{\text{int}}$ is not expending its entire credit budget in network G. Therefore G cannot be an equilibrium. The total credit allocated by nodes in $V_{G'}^{\text{int}}$ in equilibrium G' is given by

$$\begin{aligned}
\sum_{u \in V_{G'}^{\text{int}}} \sum_{v:(u,v)\in E} c_{uv}(G') &= \sum_{u \in V_{G'}^{\text{int}}} \sum_{e=(u,v)\in E_{G'}^{\min}} c_{uv}(G') \text{ (by definition of } V_{G'}^{\text{int}}),\\
&= \sum_{e \in E_{G'}^{\min} - Y} c_e(G') + \sum_{\substack{e=(u,v)\in Y\\ u \in V_{G'}^{\text{int}}}} c_{uv}(G') \text{ (by P3 and P4)},\\
&= \sum_{e \in E_{G'}^{\min}} c_e(G') \text{ (from P6)},\\
&> \sum_{e \in E_{G'}^{\min}} c_e(G) \text{ (from P5)},\\
&= \sum_{e \in E_{G'}^{\min} - Y} c_e(G) + \sum_{e' \in Y} c_{e'}(G),\\
&\geq \sum_{e \in E_{G'}^{\min} - Y} c_e(G) + \sum_{\substack{e'=(u,v)\in Y\\ u \in V_{G'}^{\text{int}}}} c_{uv}(G) \text{ (because } 1 \forall e = (u,v) \in E, c_e(G) \geq c_{uv}(G)),\\
&= \sum_{u \in V_{G'}^{\text{int}}} \sum_{v:(u,v)\in E} c_{uv}(G) \text{ (by definition of } V_{G'}^{\text{int}} \text{ and by P3 and P4).}
\end{aligned}$$

Thus we have shown that the total credit allocated by nodes in $V_{G'}^{\text{int}}$ in equilibrium G' is greater than that in equilibrium G. This implies that there exists at least one node in $V_{G'}^{\text{int}}$, say u^*, such that

$$\sum_{v:(u^*,v)\in E} c_{u^*v}(G') > \sum_{v:(u^*,v)\in E} c_{u^*v}(G).$$

So, u^* is violating the conditions of Proposition 3.1 in network G and therefore G cannot be a Nash equilibrium. □

PROOF OF PROPOSITION 3.3. Consider the set $E_G^{\min}$. We know that for all edges $e \in E_G^{\min}$, $r_e(G') \geq r_{uv}(G)$ (by Proposition 3.2). We partition the set $E_G^{\min}$ into sets T_1 and T_2 (informally corresponding to "type I" and "type II" edges) as follows.

$$T_1 = \{e \in E_G^{\min} \mid r_{uv}(G') = r_{uv}(G)\}; \quad T_2 = \{e \in E_G^{\min} \mid r_{uv}(G') > r_{uv}(G)\}$$

Note that $E_G^{\min} = E_{G'}^{\min}$ if and only if T_2 is empty. We prove that T_2 is empty by contradiction. Assume $T_2 \neq \emptyset$.

P1. For all edges $e \in T_2,\ c_e(G') < c_e(G)$ (since $r_{uv}(\cdot)$ is a nonincreasing function). The following derives from P1 and the definition of $V_G^{\min}$.

P2.

$$\sum_{u \in V_G^{\min}} \sum_{e=(u,v)\in E \cap T_2} c_{uv}(G') < \sum_{u \in V_G^{\min}} \sum_{e=(u,v)\in E \cap T_2} c_{uv}(G)$$

The following comes from P2.

P3. $\exists u \in V_G^{\min}$ (say u^*) such that

$$\sum_{e=(u^*,v)\in E\cap T_2} c_{u^*v}(G') < \sum_{e=(u^*,v)\in E\cap T_2} c_{u^*v}(G).$$

P4. $u^* \notin V_G^+$.
Proof. Assume $u^* \in V_G^+$. Then, there exists an edge $e' \in E - E_G^{\min}$ incident upon u^*, which means edges in $E_G^{\min}$ incident upon u^* are not in $E_{u^*}^*(G)$ (by definition of $E_G^{\min}$). This implies that, for all edges $e = (u^*, v) \in E_G^{\min}$, $c_{u^*v}(G) = 0$. But this implies P3 cannot be true; we have a contradiction.
P5. For all edges $e = (u^*, v) \notin T_2$, $c_{u^*v}(G') = 0$.
Proof. Since $u^* \notin V_G^+$ (P4), all edges incident upon u^* are in $E_G^{\min}$. Consider an edge $e = (u^*, v) \notin T_2$. Since T_1 and T_2 partition $E_G^{\min}$, this implies $e \in T_1$. We know that $r_{uv}(G') = r_G^{\min} < r_{e'}(G')$ for all edges $e' \in T_2$ incident upon u^* (by definition of T_1 and T_2). Therefore $c_{u^*v}(G') = 0$ (by Proposition 3.1).

From P3 and P5, it follows that

$$\sum_{e=(u^*,v)\in E} c_{u^*v}(G') < \sum_{e=(u^*,v)\in E} c_{u^*v}(G).$$

That is, u^* is not allocating its entire budget in G'. Hence G' is not a Nash equilibrium, resulting in a contradiction. □

PROOF OF PROPOSITION 3.4. Observe that any improving unilateral deviation for node u in $I_{(G,F)}$ starting from $G_{(F)}$ is also a valid improving unilateral deviation for u in I starting from G. □

PROOF OF PROPOSITION 3.5. Proof by induction on the number of edges m in the network.

Inductive Hypothesis. Assume the theorem holds for instances with at most m edges.

Base Case. $m = 1$. If $r_{uv}(G) > 0$ in Nash equilibrium G, then G is the only Nash equilibrium. Otherwise, $r_{uv}(G) = 0$ in all Nash equilibria.
We now show that the theorem is also true for instances with $m + 1$ edges. We are given the instance $I : G = (V, E); p_{uv}, e \in E; B_u, u \in V$, with $|E| = m + 1$. Let G and G' be two Nash equilibria of I. By Proposition 3.3, $E_G^{\min} = E_{G'}^{\min}$. Let $F := E_G^{\min}$.
By Proposition 3.4, $G_{(F)}$ is a Nash equilibrium for the instance $I_{(G,F)}$ and $G'_{(F)}$ is a Nash equilibrium for the instance $I_{(G',F)}$. Instances $I_{(G,F)}$ and $I_{(G',F)}$ have the same graph $G^{(F)}$ and the same utilities $p_{uv}^{(F)}$, since these depend only on F. We show next that they also have the same budgets: for all nodes $u \in V$, $B_u^{(G,F)} = B_u^{(G',F)}$. From Corollary 3.1, we know that $V_G^{\min} = V_{G'}^{\min}$ and $V_G^+ = V_{G'}^+$. We divide the proof into three cases.

Case (i). $u \in V \setminus V_G^{\min}$. This implies u is not incident on any edge in F. Therefore $B_u^{(G,F)} = B_u^{(G',F)} = B_u$.
Case (ii). $u \in V_G^+$. Then, by slackness and as explained in the proof of Proposition 3.3, for all edges $e \in F$, $c_G(u, e) = c_{G'}(u, e) = 0$. This implies $B_u^{(G,F)} = B_u^{(G',F)} = B_u$.
Case (iii). $u \in V_G^{\min} \setminus V_G^+$. Then $B_u^{(G,F)} = B_u^{(G',F)} = 0$ (by construction, since $E_u \subseteq F$).

Hence, $I_{(G,F)}$ and $I_{(G',F)}$ are the same problem instance, and $G_{(F)}$ and $G'_{(F)}$ are two Nash equilibria of this instance. Further $G^{(F)}$ has at most m edges. Therefore, by the inductive hypothesis, for all edges $e \in E \setminus F,\ r_{uv}(G_{(F)}) = r_{uv}(G'_{(F)})$. But

$$\forall e \in E \setminus F,\ c_e(G_{(F)}) = c_e(G) \text{ and } c_e(G'_{(F)}) = c_e(G')$$

by definition of an F-restriction. This implies

$$\forall (u,v) \in E \setminus F,\ r_{uv}(G) = r_{uv}(G').$$

Also, $\forall (u,v) \in F,\ r_{uv}(G) = r_{uv}(G')$ (from Proposition 3.3). This completes the inductive proof. □

C. PROOFS OF THEOREMS ABOUT THE GLOBAL RISK MODEL

C.1. Proof of Theorem 4.3

The proof consists of two parts. First, we show that G is weakly connected. Second, we show that when n is large, G is a star network.

Throughout this proof, we call a node u a *leaf node* in G if no nodes extend credit to u in G. We define some notation that we need in the proof. Observe that, since agents are single minded, in any network G', for any node u, there is exactly one edge leaving u. For any network G' and for nodes $u, v \in V$, we define the network $\Gamma_{G'}(u,v) := G' \setminus \{(u, \tau_{G'}(u))\} \cup \{(u,v)\}$, that is, $\Gamma_{G'}(u,v)$ is the network obtained by deleting the edge leaving u in G' and adding the edge (u,v). Let $h := c/(c+1)$. For an integer $k \geq 2$ and $1 \leq l \leq k-1$, let

$$q(l,k) := \frac{h^l + h^{k-l} - 2h^k}{1-h^k}. \tag{10}$$

For an integer $k \geq 2$, let

$$\rho_C(k) := \sum_{l=1}^{k-1} q(l,k) = 2\sum_{l=1}^{k-1} \frac{h^l - h^k}{1-h^k}. \tag{11}$$

Observe that:

(a) $\lim_{k\to\infty} \rho_C(k) = \frac{2h}{(1-h)}$, and
(b) for all finite $k \geq 2$, $\frac{2h}{(1-h)} > \rho_C(k)$ (since $h < 1$).

For an integer $k \geq 2$, let

$$\rho_L(k) := \sum_{l=1}^{k-1} h^l = h\frac{1-h^{l-1}}{1-h}. \tag{12}$$

In a unit transaction regime $q(l,k)$ is the steady-state transaction success probability between two nodes separated by l edges in a circular graph with k edges of capacity c each, whereas $\rho_L(k)$ is the sum of the steady-state transaction success probabilities for a node at the end of a k-node line graph [Dandekar et al. 2011].

We use these quantities to argue about improving unilateral deviations of nodes in G.

PROPOSITION C.1. *Let G be a Nash equilibrium. Then, G is weakly connected.*

PROOF. By contradiction. Assume that G is not weakly connected. Consider two weakly connected components G_1 and G_2 of G. Let V_1 and V_2 be the set of nodes in G_1 and G_2, respectively. Let $|V_1| = n_1$ and $|V_2| = n_2$. Observe that: (i) both $n_1, n_2 \geq 2$, and

(ii) G_1 and G_2 have n_1 and n_2 edges, respectively, and consequently, each contains a directed cycle. We divide the proof into three cases.

— *Case I*. Both G_1 and G_2 each have at least one leaf node. Let $u \in V_1$ and $v \in V_2$ be leaf nodes. Further assume there exist nodes $y \in V_1$ and $z \in V_2$ such that $\tau_G(u) = y$ and $\tau_G(v) = z$. For any network t, define $F(u,t) := \sum_{w \in V \setminus \{u,v,y,z\}} p_{uw}(t)$ and let $F(v,t) := \sum_{w \in V \setminus \{u,v,y,z\}} p_{vw}(t)$.
Then we have

$$U_u(G) = \gamma \left(p_{uy}(G) + F(u,G)\right) - \delta_y c,$$
$$U_v(G) = \gamma \left(p_{vz}(G) + F(v,G)\right) - \delta_z c.$$

Consider another network $G' = \Gamma_G(u,z)$, that is, in G' node u extends credit to z instead of y. Then,

$$U_u(G') = \gamma \left(p_{uz}(G') + p_{uv}(G') + F(u,G')\right) - \delta_z c.$$

Similarly, consider a network $G'' = \Gamma_G(v,y)$, that is, in G'' node v extends credit to y instead of z. Then,

$$U_v(G'') = \gamma \left(p_{vy}(G'') + p_{vu}(G'') + F(v,G'')\right) - \delta_y c.$$

However, note that $U_u(G') > U_v(G)$, since $p_{uv}(G')$ is an extra term and all other terms are identical. For similar reasons, $U_v(G'') > U_u(G)$. Thus we have

$$U_u(G') > U_v(G) \geq U_v(G'') > U_u(G),$$

where the second inequality holds because G is a Nash equilibrium, which implies v does not have an improving unilateral deviation. However, $U_u(G') > U_u(G)$ implies that u has an improving unilateral deviation, contradicting the assumption that G is a Nash equilibrium.
Case II. Neither G_1 nor G_2 has a leaf node (i.e., G_1 and G_2 are both directed cycles). Observe that, for any node $u \in V_1$ such that $\tau_G(u) = v$, $U_u(G) = \gamma\lambda\rho_C(n_1) - \delta_v c$, where $\rho_C(n_1)$ is defined in (11). Similarly, for any node $u \in V_2$ such that $\tau_G(u) = v$, $U_u(G) = \gamma\lambda\rho_C(n_2) - \delta_v c$. Let $v_1 \in \arg\min_{v \in V_1} \delta_v$ be a node in V_1 that has the smallest value of δ_v in V_1. Similarly, let $v_2 \in \arg\min_{v \in V_2} \delta_v$.
Since G is a Nash equilibrium, no node in V_2 has an incentive to switch and extend credit to v_1 instead. This implies that, for all nodes $u \in V_2$,

$$\gamma\lambda \left(h + h\rho_C(n_1) + \rho_L(n_2)\right) - \delta_{v_1} c \leq \gamma\lambda\rho_C(n_2) - \delta_u c. \tag{13}$$

Let u_1 be the node in V_1 that extends credit to v_1, that is, $\tau_G(u_1) = v_1$. We show that u_1 now has an incentive to extend credit to v_2. Instantiating (13) for v_2, we get

$$\gamma\lambda \left(h + h\rho_C(n_1) + \rho_L(n_2)\right) - \delta_{v_1} c \leq \gamma\lambda\rho_C(n_2) - \delta_{v_2} c,$$

or equivalently,

$$(\delta_{v_1} - \delta_{v_2})c \geq \gamma\lambda \left(h + h\rho_C(n_1) - \rho_C(n_2) + \rho_L(n_2)\right). \tag{14}$$

Next we show that $h + \rho_L(n_2) > \rho_C(n_1)(1-h)$. Recall that for all finite $k \geq 2$, $\rho_C(k) < 2h/(1-h)$. This implies $\rho_C(n_1)(1-h) < 2h$. Also recall that for all $k \geq 2$, $\rho_L(k) \geq h$. This implies $h + \rho_L(n_2) \geq 2h > \rho_C(n_1)(1-h)$.
Replacing $h + \rho_L(n_2) \geq 2h$ with $\rho_C(n_1)(1-h)$ in (14) , we get

$$(\delta_{v_1} - \delta_{v_2})c > \gamma\lambda \left(\rho_C(n_1) - \rho_C(n_2)\right). \tag{15}$$

Using a similar argument as earlier, it is easy to see that $h + \rho_L(n_1) > \rho_C(n_2)(1-h)$, or equivalently,

$$\rho_C(n_2) < h\rho_C(n_2) + h + \rho_L(n_1).$$

Substituting the preceding upper bound on $\rho_C(n_2)$ in (15), we get

$$(\delta_{v_1} - \delta_{v_2})c > \gamma\lambda\left(\rho_C(n_1) - h\rho_C(n_2) - h - \rho_L(n_1)\right).$$

Rearranging, we get

$$\gamma\lambda\left(h + h\rho_C(n_2) + \rho_L(n_1)\right) - \delta_{v_2}c > \gamma\lambda\rho_C(n_1) - \delta_{v_1}c.$$

Observe that the left-hand side of the prior inequality is the utility of u_1 when it extends credit to v_2 instead of v_1, whereas the right-hand side is $U_{u_1}(G)$. Thus, u_1 can improve its utility in G by extending credit to v_2, contradicting the assumption that G is a Nash equilibrium.

Case III involves exactly one of G_1 and G_2 having at least one leaf node. The argument in this case is similar to that in Case II. □

PROPOSITION C.2. *Let G be a Nash equilibrium. For a sufficiently large n, there must exist a node, say u^*, such that for all nodes $v \in V \setminus \{u^*\}$, $\tau_G(v) = u^*$.*

PROOF. By contradiction. From Proposition C.1, we know that G is weakly connected. We divide the proof into cases based on the number of leaf nodes in G.

— *Case I*. G has zero leaf nodes (i.e., G is a circular graph). Observe that the total utility $U_y(G)$ of a node y in G is given by

$$U_y(G) = \gamma\lambda\rho_C(n) - \delta_z c,$$

where $z = \tau_G(y)$ is the node to which y extends credit, and $\rho_C(n)$ is defined in (11). Since $\rho_C(n)$ is independent of y, the node that maximizes $U_y(G)$ must be extending credit to the node with the smallest value of δ_v. Let $v_1 \in \arg\min_{v\in V} \delta_v$ be a node with the smallest value of δ_v. Let v_2 be the node that extends credit to v_1, that is, $v_1 = \tau_G(v_2)$. Further let u be the node extending credit to v_2, that is, $v_2 = \tau_t(u)$. We will show that u can improve its utility by extending credit to v_1 instead of v_2.
Let $G' = \Gamma_G(u, v_1)$, that is, G' is the network obtained when u extends credit to v_1 instead of u_2. The utility $U_u(G')$ of u in G' is given by

$$\begin{aligned}
U_u(G') &= \gamma\left(\lambda\rho_C(n-1) + p_{uv_2}(G')\right) - \delta_{v_1}c,\\
&= \gamma\left(\lambda p(n-1) + hp_{uv_1}(G')\right) - \delta_{v_1}c \ (p_{uv_2}(G') = p_{v_2u}(G') = hp_{v_1u}(G'),\\
&\quad \text{by Lemma 4.1}),\\
&= \gamma\lambda\left(p(n-1) + hq(1, n-1)\right) - \delta_{v_1}c \ (\text{since } u \text{ and } v_1 \text{ are adjacent}\\
&\quad \text{in a cycle of length } n-1),\\
&> \gamma\lambda p(n) - \delta_{v_1}c \ (\text{for a sufficiently large } n),\\
&= U_{v_2}(G) \geq U_u(G) \ (\text{since } \delta_{v_1} \leq \delta_{v_2}).
\end{aligned}$$

Since u has a unilateral improving deviation in G, G cannot be a Nash equilibrium.

— *Case II*. G has exactly one leaf node (that is, G is a balloon-like graph).
Observe that, since G has n edges, it must contain a directed cycle. Let $C \subseteq V$ be the set of nodes that are part of the cycle in G. Let $T = V \setminus C$ be the set of "tail" nodes. Let $|C| = n_1$ and $|T| = n_2$, where $n_1 + n_2 = n$. Since G has exactly one leaf node, observe that: (a) T is nonempty, and (b) nodes in T form a line graph.
Let $v_1 \in C$ be the node in the cycle that connects to T, that is, there exists a node $w \in T$ such that $\tau_G(w) = v_1$. Then, since G is a Nash equlibrium, it must be that $v_1 \in \arg\min_{v\in C} \delta_v$, that is, v_1 must have the lowest default probability among nodes in C.

As n grows, at least one of n_1 and n_2 must grow. We divide the proof into two cases.

— *Case II(a)*. We will prove that when n_1 is sufficiently large, G cannot be a Nash equilibrium. Let v_2 be the node that extends credit to v_1, that is, $v_1 = \tau_G(v_2)$. Let u be the node extending credit to v_2, that is, $v_2 = \tau_G(u)$. We will show that u can improve its utility by extending credit to v_1 instead of v_2.

Let $G' = \Gamma_G(u, v_1)$, that is, G' is the network obtained when u extends credit to v_1 instead of u_2. Observe that, by Lemma 4.1, $p_{uv_2}(G') = p_{uv_1}(G')p_{v_1v_2}(G') = hp_{uv_1}$. Similarly, for a node $w \in T$, $p_{uw}(G') = p_{uv_1}(G')p_{v_1w}(G')$. Also, observe that $p_{uv_1}(G') = \lambda q(1, n_1 - 1) > \lambda q(1, n_1)$, where $q(1, n_1)$ is defined in (10). Finally, $\sum_{w\in T} p_{v_1w}(G') = \sum_{w\in T} p_{v_1w}(G) = \lambda\rho_L(n_2+1)$, where $\rho_L(n_2+1)$ is defined in (12). We will use these observations to show that, for a sufficiently large n_1, $U_u(G') > U_u(G)$.

$$\begin{aligned} U_u(G') &= \gamma \left(\lambda\rho_C(n_1-1) + p_{uv_2}(G') + \sum_{w\in T} p_{uw}(G') \right) - \delta_{v_1} c \\ &= \gamma\lambda \left(\rho_C(n_1-1) + q(1, n_1-1)(h + \rho_L(n_2+1)) \right) - \delta_{v_1} c \end{aligned}$$

Recall that as $n \to \infty$, $\rho_C(n) = 2h/1-h)$. Also, $q(1, n_1 - 1) > q(1, n_1)$. Therefore for a sufficiently large n_1,

$$\rho_C(n_1-1) + q(1, n_1-1)(h + \rho_L(n_2+1)) > \rho_C(n_1) + q(1, n_1)\rho_L(n_2+1).$$

This implies, for a sufficiently large n_1, that

$$\begin{aligned} U_u(G') &> \gamma\lambda \left(\rho_C(n_1) + q(1, n_1)\rho_L(n_2+1)\right) - \delta_{v_1} c, \\ &= U_{v_1}(G) \geq U_u(G) \text{ (since } \delta_{v_1} \leq \delta_{v_2}). \end{aligned}$$

Thus, u can improve its utility by deviating, contradicting the assumption that G is a Nash equilibrium.

— *Case II(b)*. We will prove that when n_2 is sufficiently large, G cannot be a Nash equilibrium. Let y be the leaf node in G and let w be the node in G to which y extends credit. We will show that, for a sufficiently large n_2, v_1 can improve its utility by extending credit to w. Let z be the node in C, to which v_1 extends credit. Since G is a Nash equilibrium, it must be that $\delta_w \leq \delta_{v_1} \leq \delta_z$ (otherwise y would have an improving deviation).

Observe that $U_{v_1}(G) = \lambda(\rho_C(n_1) + \rho_L(n_2+1)) - \delta_z c$. Let $G' = \Gamma_G(v_1, w)$, that is, G' is the network obtained when v_1 extends credit to w instead of z. Then, $U_{v_1}(G') = \lambda(\rho_C(n_2) + \rho_L(n_1)) + p_{v_1y}(G') - \delta_w c$, where $p_{v_1y}(G') = p_{v_1w}(G')p_{wy}(G') = hq(1, n_2)$. Observe that, for a sufficiently large n_2, $\rho_L(n_2+1) - \rho_L(n_2) = h^{n_2} < hq(1, n_2) = p_{v_1y}(G')$. Therefore

$$\begin{aligned} U_{v_1}(G') - U_{v_1}(G) &> \lambda(\rho_C(n_2) + \rho_L(n_1)) - \delta_w c - (\lambda(\rho_C(n_1) + \rho_L(n_2)) - \delta_z c), \\ &\geq \lambda(\rho_C(n_2) + \rho_L(n_1) - \rho_C(n_1) - \rho_L(n_2)) \text{ (since } \delta_w \leq \delta_z), \\ &= \lambda(B(n_2) - B(n_1)), \end{aligned}$$

where $B(k) := \rho_C(k) - \rho_L(k)$. Substituting the values of $\rho_C(k)$ and $\rho_L(k)$, we get

$$B(k) = \frac{h}{1-h} + \frac{h^k}{1-h^k}\left(\frac{1-h^k}{1-h} - 2k\right).$$

Observe that: (i) $\lim_{k\to\infty} B(k) = h/(1-h)$, and (ii) for all $k \geq 2$, $B(k) \leq h/(1-h)$, since $(1-h^k)/(1-h) \leq 2k$. Therefore for a fixed n_1,

$$\lim_{n_2\to\infty} B(n_2) - B(n_1) = \frac{h}{1-h} - B(n_1) \geq 0.$$

We have shown that $U_{v_1}(G') - U_{v_1}(G) > 0$, that is, v_1 has an improving deviation in G, contradicting the assumption that G is a Nash equilibrium.

—*Case III*. G has two or more leaf nodes. The proof of this case is similar to that in Case I of Proposition C.1. Let $u, v \in V$ be leaf nodes in G. Further assume for the purpose of contradiction that there exist nodes $y, z \in V$ such that $\tau_G(u) = y$ and $\tau_G(v) = z$. For any network t, let $F(u,t) := \sum_{w\in V\setminus\{u,v,y,z\}} p_{uw}(t)$ and let $F(v,t) := \sum_{w\in V\setminus\{u,v,y,z\}} p_{vw}(t)$.
Then we have

$$U_u(G) = \gamma\left(p_{uy}(G) + p_{uz}(G) + p_{uv}(G) + F(u,G)\right) - \delta_y c,$$
$$U_v(G) = \gamma\left(p_{vy}(G) + p_{vz}(G) + p_{vu}(G) + F(v,G)\right) - \delta_z c.$$

Consider another network $G' = \Gamma_G(u,z)$, that is, in G' node u extends credit to z instead of x. Then,

$$U_u(G') = \gamma\left(p_{uy}(G') + p_{uz}(G') + p_{uv}(G') + F(u,G')\right) - \delta_z c.$$

Similarly, consider a network $G'' = \Gamma_G(v,y)$, that is, in G'' node v extends credit to y instead of z. Then,

$$U_v(G'') = \gamma\left(p_{vy}(G'') + p_{vz}(G'') + p_{vu}(G'') + F(v,G'')\right) - \delta_y c.$$

However, note that $U_u(G') > U_v(G)$, since $p_{uv}(G') > p_{vu}(G)$ and all other terms are identical. For similar reasons, $U_v(G'') > U_u(G)$. Thus we have

$$U_u(G') > U_v(G) \geq U_v(G'') > U_u(G).$$

where the second inequality holds because G is a Nash equilibrium, which implies v does not have an improving unilateral deviation. However, $U_u(G') > U_u(G)$ implies that u has an improving unilateral deviation, contradicting the assumption that G is a Nash equilibrium.

□

This concludes the proof of Theorem 4.3.

PROOF OF THEOREM 4.6. The statement clearly holds for u_1. We prove the general statement by contradiction. Assume that node u_i, for some $i > 1$, is the first node such that $\tau_{G_i}(u_i) \notin \{u_{i-1}, \tau_{G_{i-1}}(u_{i-1})\}$. Let $\tau_{G_i}(u_i) = u'$. Let $\tau_{G_{i-1}}(u_{i-1}) = v$. Observe that both $u', v \in V_{i-2}$. Since u_i extended credit to u' instead of v, it must be that the utility of u_i from extending credit to u' is at least the utility that u_i would have derived from extending credit to v.

$$U_{u_i}(G_{i-1} \cup \{(u_i, u')\}) \geq U_{u_i}(G_{i-1} \cup \{(u, v)\}) \quad (16)$$

Let $G(u') := G_{i-1} \cup \{(u_i, u')\}$, and $G(v) := G_{i-1} \cup \{(u, v)\}$. Then, (16) implies

$$\gamma \sum_{w\in V_i} p_{u_i w}(G(u')) - \delta_{u'} c \geq \gamma \sum_{w\in V_i} p_{u_i w}(G(v)) - \delta_v c. \quad (17)$$

For a network G_i and a node $u \in V_i$, define $F(u, G_i) := \sum_{w \in V_i} p_{uw}(G_i)$. Using Lemma 4.1, $\sum_{w \in V_i} p_{u_i w}(G(u'))$ can be written as

$$\sum_{w \in V_i} p_{u_i w}(G(u')) = p_{u_i u'}(G_i) + hF(u', G_{i-1}) = \lambda h + h\left(F(u', G_{i-2}) + p_{u' u_{i-1}}(G_{i-1})\right).$$

Here $h = c/(c+1)$. Similarly,

$$\sum_{w \in V_i} p_{u_i w}(G(v)) = \lambda h + h\left(F(v, G_{i-2}) + p_{v u_{i-1}}(G_{i-1})\right).$$

Substituting the expressions for $\sum_{w \in V_i} p_{u_i w}(G(u'))$ and $\sum_{w \in V_i} p_{u_i w}(G(v))$ in (17), we get

$$\gamma h\left(\lambda + F(u', G_{i-2}) + p_{u' u_{i-1}}(G_{i-1})\right) - \delta_{u'} c \geq \gamma h\left(\lambda + F(v, G_{i-2}) + p_{v u_{i-1}}(G_{i-1})\right) - \delta_v c. \tag{18}$$

Note that, since u_{i-1} extended credit to v instead of to u', u_{i-1} is at least at distance two from u' in G_{i-1}. Therefore

$$p_{u' u_{i-1}}(G_{i-1}) < \lambda h = p_{v u_{i-1}}(G_{i-1}). \tag{19}$$

(18) and (19) imply that

$$\gamma h\left(\lambda + F(u', G_{i-2})\right) - \delta_{u'} c > \gamma h\left(\lambda + F(v, G_{i-2})\right) - \delta_v c.$$

However,

$$\gamma h(\lambda + F(u', G_{i-2}) - \delta_{u'} c = U_{u_{i-1}}(G_{i-2} \cup \{(u_{i-1}, u')\})$$

and

$$\gamma h(\lambda + F(v, G_{i-2}) - \delta_v c = U_{u_{i-1}}(G_{i-2} \cup \{(u_{i-1}, v\}).$$

Thus our analysis implies that $U_{u_{i-1}}(G_{i-2} \cup \{(u_{i-1}, u')\}) > U_{u_{i-1}}(G_{i-2} \cup \{(u_{i-1}, v\})$, that is, u_{i-1} would have derived a higher utility by extending credit to u', contradicting the assumption that u_{i-1} maximized its utility by extending credit to v. □

ACKNOWLEDGMENTS

Tae Hyung Kim assisted in the implementation of the automated EGTA process, and Ben-Alexander Cassell extended our EGTA infrastructure to facilitate this work.

REFERENCES

Elliot Anshelevich and Martin Hoefer. 2012. Contribution games in networks. *Algorithmica 63*, 51–90.

Venkatesh Bala and Sanjeev Goyal. 2000. A noncooperative model of network formation. *Econometrica, 68*, 1181–1229.

Larry Blume, David Easley, Jon Kleinberg, Robert Kleinberg, and Eva Tardos. 2011. Network formation in the presence of contagious risk. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC'11).* 1–10.

Ben-Alexander Cassell and Michael P. Wellman. 2013. EGTAOnline: An experiment manager for simulation-based game studies. In *Proceedings of the International Workshop on Multi-Agent Based Simulation (MABS'13).* Lecture Notes in Artificial Intelligence, vol. 7838, Springer, 85–100.

Jacomo Corbo, Antoni Calvo-Armengol, and David Parkes. 2006. A study of Nash equilibrium in contribution games for peer-to-peer networks. *SIGOPS Oper. Syst. Rev. 40*, 61–66.

Pranav Dandekar, Ashish Goel, Ramesh Govindan, and Ian Post. 2011. Liquidity in credit networks: A little trust goes a long way. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC'11).* 147–156.

Dimitri B. Defigueiredo and Earl T. Barr. 2005. TrustDavis: A non-exploitable online reputation system. In *Proceedings of the 7th IEEE International Conference on E-Commerce Technology (CEC'05).* 274–283.

Alex Fabrikant, Ankur Luthra, Elitza N. Maneva, Christos H. Papadimitriou, and Scott Shenker. 2003. On a network creation game. In *Proceedings of the 22nd Annual ACM Symposium on Principles of Distributed Computing (PODC'03).* 347–351.

Eric Friedman and Paul Resnick. 2001. The social cost of cheap pseudonyms. *J. Econ. Manag. Strategy 10*, 2, 173–199.

Eric Friedman, Paul Resnick, and Rahul Sami. 2007. Manipulation-resistant reputation systems. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds., Cambridge University Press, 677–697.

Arpita Ghosh, Mohammad Mahdian, Daniel M. Reeves, David M. Pennock, and Ryan Fugger. 2007. Mechanism design on trust networks. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE'07).* 257–268.

Herbert Gintis. 2000. *Game Theory Evolving*. Princeton University Press.

Matthew O. Jackson and Asher Wolinsky. 1996. A strategic model of social and economic networks. *J. Econ. Theory 71*, 1, 44–74.

Patrick R. Jordan and Michael P. Wellman. 2009. Generalization risk minimization in empirical game models. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'09).* 553–560.

Dean Karlan, Markus Mobius, Tanya Rosenblat, and Adam Szeidl. 2009. Trust and social collateral. *Quart. J. Econ. 124*, 3, 1307–1361. http://ideas.repec.org/a/tpr/qjecon/v124y2009i3p1307-1361.html.

Freddy Limpens and Denis Gillet. 2011. A competence bartering platform for learners. In *Proceedings of the 10th International Conference on Advances in Web-Based Learning (ICWL'11).* 148–153.

Zhengye Liu, Hao Hu, Yong Liu, Keith W. Ross, Yao Wang, and Markus Mobius. 2010. P2P trading in social networks: The value of staying connected. In *Proceedings of the 29th IEEE International Conference on Computer Communications (INFOCOM'10).* 2489–2497.

Alan Mislove, Ansley Post, Peter Druschel, and Krishna P. Gummadi. 2008. Ostra: Leveraging trust to thwart unwanted communication. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI'08).* 15–30. http://www.usenix.org/events/nsdi08/tech/fullpapers/mislove/mislove.pdf.

Dov Monderer and Lloyd S. Shapley. 1996. Potential games. *Games Econ. Behav.14*, 124–143.

Abraham Neyman. 1997. Correlated equilibrium and potential games. *Int. J. Game Theory 26*, 223–227.

Bimal Viswanath, Mainack Mondal, Krishna P. Gummadi, Alan Mislove, and Ansley Post. 2012. Canal: Scaling social network-based Sybil tolerance schemes. In *Proceedings of the 7th European Conference on Computer Systems (EuroSys'12).* 309–322.

Yevgeniy Vorobeychik, Michael P. Wellman, and Satinder Singh. 2007. Learning payoff functions in infinite games. *Mach. Learn. 67*, 145–168.

Michael P. Wellman. 2006. Methods for empirical game-theoretic analysis (extended abstract). In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI'06).* 1552–1555.

Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, Shih-Fen Cheng, and Rahul Suri. 2005. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI'05).* 502–508.

Michael P. Wellman, Tae Hyung Kim, and Quang Duong. 2013. Analyzing incentives for protocol compliance in complex domains: A case study of introduction-based routing. In *Proceedings of the 12th Workshop on the Economics of Information Security (WEIS'13).*

Bryce Wiedenbeck and Michael P. Wellman. 2012. Scaling simulation-based game analysis through deviation-preserving reduction. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'12).* 931–938.

Received October 2013; revised July 2014; accepted October 2014