

An Empirical Game-Theoretic Analysis of Credit Network Formation

Michael P. Wellman and Bryce Wiedenbeck

Abstract—The framework of *credit networks* provides a flexible and robust model of distributed trust, based on pairwise credit allocations representing commitments to allow transactions. Since issuing credit entails risks as well as benefits, it is unclear whether self-interested and autonomous agents will form viable credit networks. We tackle this question through an extensive simulation-based game-theoretic analysis of a 61-node credit network formation scenario, covering eight environments varying on information and cost-benefit parameters. We find that viable credit networks form in equilibrium given sufficiently high transaction value, or sufficiently low default risk. Although the amount of credit issued in equilibrium is significantly lower than in the social optimum, as is the social welfare achieved, this difference diminishes proportionally as the environment becomes more favorable overall.

I. INTRODUCTION

Most formal models of network formation posit simple reflexive or myopic behavior on the part of nodes forming connections or selecting local states. In contrast, *strategic* network formation models account for the self-interest and autonomy of the nodes (which given this treatment merit the label *agent*), typically employing the conceptual and technical framework of game theory. Predicting the likely evolution of networks under strategic behavior is a notoriously difficult problem, due to path dependence, tipping phenomena, and general sensitivity to initial conditions and expectations. Nevertheless, there is a growing game-theoretic literature on network formation, in which agents make deliberate decisions about forming connections, then incur costs and receive benefits based on the resulting network. Much of this literature adopts abstract models of the costs and benefits of establishing links. There are a wide range of models, varying for example on whether edge creation is unilateral [1], [2] or bilateral [3], or an indirect consequence of other decisions (e.g., according to an affiliation model [4]). The models also reflect many possibilities for relating costs and benefits to network structure. Regardless of formulation, analysis of network formation games tends to be extremely challenging, due to the size of game (number of players, combinatorics of available actions), and the complexity of dependence between network structure and outcomes. This naturally leads analysts to focus on highly stylized abstract models, as a tractable means to produce general insights.

Given a *particular* scenario in network formation, analyses from prior literature are unlikely to be directly applicable, as the details of the scenario will inevitably not be accounted for in the stylized models. Every domain will have its own, typically complicated way that network structure translates

to cost and benefits for individual agents. Thus, there is no alternative to performing customized analyses for particular scenarios of interest, dealing with whatever scale and complexity challenges they present.

The investigation reported here is an example of just such a particular network formation scenario. We seek to understand the prospects for autonomous self-interested agents to form viable networks of distributed trust, called *credit networks*. As sufficiently rich models of these scenarios appear to be beyond the scope of analytic tractability, we rely on simulation and empirical game-theoretic analysis to derive our strategic conclusions. By finding interesting relationships between environments and equilibrium strategies, we also demonstrate how empirical methods can provide insights complementary to purely analytic methods.

II. CREDIT NETWORKS

Our analysis focuses on an abstract model of *credit networks*: weighted directed graphs that represent the capacity of agents (each agent represented as a node in the graph) to transact with each other.

A. 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 & Barr [5] sought a reputation system with bounded loss from coalitions of malicious users.
- Ghosh et al. [6] aimed to support distributed payment and multi-user credit checking for multi-item auctions.
- Karlan et al. [7] wanted to construct an economic model of informal borrowing networks.
- Mislove et al. [8] 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. By introducing suitable definitions of transaction, credit networks can support a wide variety of applications. For example, the inventors enumerated above interpret transactions respectively as obtaining references

guaranteeing good behavior [5], paying for auction winnings [6], borrowing an asset [7], and communicating messages [8]. Subsequent authors proposed using this framework to support networked asynchronous bilateral trading [9], and bartering of tutorial services [10].

Much of the power of credit networks accrues due to the *transitivity* of credit. If A issues credit to B who issues credit to C , then we have directly that B can transact with A and C with B . As long as the underlying transaction composes (as it does for the above cases: reference, payment, borrowing, and communication), then these credit relationships also entail that C can transact with A . For example, a payment of C to B combined with an equal payment of B to A is tantamount to a net payment from C to A . In other words, we can *route* transactions along paths of credit relationships. More generally, with credit issued to quantitative degrees, one agent can transact with another at a given quantity if it is possible to route a total *flow* of that quantity across credit relationships in the network.

B. Formal Definition

The common underlying credit model of these four proposals was first noticed by Dandekar et al. [11], who introduced the unifying term “credit network” and its formal definition. We employ their terminology and notation here. A credit network $G = (V, E)$ is a directed graph with n nodes representing agents, and m weighted edges representing credit relationships. Let c_{uv} denote the weight on edge $(u, v) \in E$; if $(u, v) \notin E$ we say that $c_{uv} = 0$. The quantity c_{uv} represents the amount of transaction that agent u has committed to allow from v . For example, it could be the dollar amount of services that u would permit v to purchase with IOUs, or the number of messages u would admit from v . Henceforth we adopt the IOU interpretation, recognizing that the model is considerably more flexible.

The commitment to transact inherent in credit relationships includes an obligation to route transactions on paths that can be composed of transactions on edges. If node u wishes to purchase a unit of service from node w , and both $c_{uw} \geq 1$ and $c_{vw} \geq 1$, then v must enable such a transaction on request by passing one of its IOUs to w (drawing on its credit with w) in exchange for an IOU from u .

To illustrate the maintenance of credit balances on a network, consider the credit network of Figure 1. Suppose agent s wishes to transact with agent t . Although s has no direct credit with t , it can access credit indirectly through several paths. Overall, it may transact up to six units, as that is the maximum flow from t to s afforded by this network. To transact the maximum allowed, it would route two units each through the paths $s \rightarrow u \rightarrow t$, $s \rightarrow u \rightarrow w \rightarrow t$, and $s \rightarrow v \rightarrow w \rightarrow t$.

Whenever a quantity q of payment is routed across an edge, the credit capacity of that edge is decremented by q , and the inverse edge is incremented by q . To illustrate, Figure 2 shows the result of s routing a payment of six units to t . The updates must be performed for each of the paths composing the total flow. For example, a total of four units

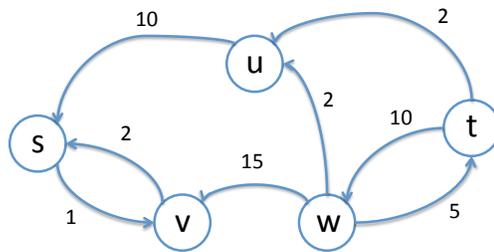


Fig. 1. A five-node credit network.

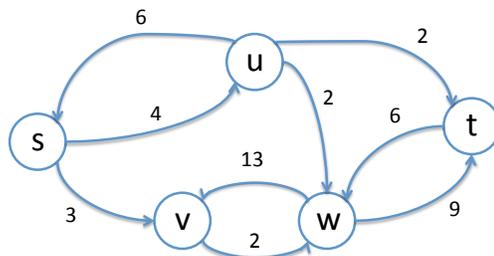


Fig. 2. The updated credit network after s transacts six units with t .

of payment pass from s to u , thus the credit c_{us} decreases from ten to six, and the credit c_{su} increases to four from zero (thus introducing a new link from s to u). Updating credit between s and v results in removing a link, as the full two available from v to s was used for this transaction. Observe that the total credit between any pair of nodes is conserved; what changes is the division of the total credit between the two directions.

C. Properties

One of the original instances of the credit network model, TrustDavis [5], was motivated by the problem of *Sybil* attacks on reputation systems, where a malicious entity would create copies of itself to amplify the impact of its endorsements. Credit networks resist *Sybil*s by requiring endorsements to be backed by actual credit exposure. Thus, enlisting collusive or duplicate identities helps only to the extent that these are also backed up by real commitments. Resnick and Sami [12] augment the basic scheme with a protocol for increasing credit limits over time based on positive experience. Viswanath et al. [13] 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-scale networks.

The effectiveness of credit networks in supporting distributed transactions was most powerfully demonstrated by Dandekar et al. [11], who established several propositions indicating that transaction failures are unlikely given sufficient network connectivity. That is, credit networks provide a high degree of *liquidity*: the ability to transact at any time at prevailing terms. Their analysis posits a probability distribution $\Lambda = \{\lambda_{uv} : u, v \in V\}$, where λ_{uv} is the probability that the next transaction request will be from agent u to agent v .

In particular they showed for several classes of graphs and with symmetric transaction probabilities (i.e., $\lambda_{uv} = \lambda_{vu}$ for all u, v), the transaction failure probability goes to zero as either network size, link density, or credit capacity increases, holding the other two parameters constant. Computational experiments further demonstrate that even networks small in size or overall credit capacity exhibit high transaction success rates if they are sufficiently well-connected.

III. CREDIT NETWORK FORMATION

The fact that networks with ample initial credit provide effective liquidity raises immediately the question: will agents indeed issue enough credit to form a viable credit network? We are concerned with scenarios where agents autonomously decide how much credit to issue and to whom, based on their own self-interest and available information. We framed the issue of *strategic formation* of credit networks in our initial investigation [14], which the current study extends. That formation of viable networks is in question at all derives from two characteristics of the problem:

- 1) The credit decision revolves around a fundamental *tradeoff*. The *benefit* of issuing credit is that it increases the prospect that the agent will be able to participate in beneficial transactions, with or through the agent given the credit. The *cost* is incurring a risk of loss, in case the credited agent defaults on its obligation. There is also a secondary risk of deficit, due to serving transactions but not obtaining the opportunity for reciprocal benefits within the scenario horizon.
- 2) Decisions about credit entail *network externalities*. Granting credit provides some benefit to the issuer, but also enhances the liquidity of the entire network. That is, the issuer incurs all the costs but accrues only some of the benefits of the credit grant. In the presence of such positive externalities, we generally expect socially suboptimal outcomes.

To address this question, we define a credit network formation *game*, which proceeds as follows. At the outset, agents are given information about default probability, and the probabilities and values of transactions with other agents. Based on this information, they decide what credit to issue, that is, agent u assigns c_{uv} , $v \in V \setminus \{u\}$. Next, agents default according to their designated probabilities. When an agent u defaults, it absconds with the credit it received (we can interpret this as having maximized its transactions with the credit issuers). Thus, a non-defaulting agent v incurs a payoff deduction c_{vu} associated with u 's default.¹ Finally, we play out a set number of unit transaction events according to the transaction rate matrix Λ . For each event, the pair (u, v) is drawn with probability λ_{uv} . If neither agent has defaulted and the max flow from v to u is at least one, the transaction succeeds. The payoff to agent u is incremented by x_{uv} : u 's value for the transaction. Correspondingly, v 's payoff is decremented by one, reflecting its cost for serving

¹Since defaulting is not a strategic choice, we assign zero payoff to agents who default.

the transaction. The credit balances are updated to reflect the transaction, as described above. (If there are multiple feasible paths or combinations for routing credit, the allocation is chosen arbitrarily.)

We consider two models for the information agents receive about default probabilities. In both, the default probability δ_u for agent u is drawn from a Beta distribution, $\delta_u \sim \text{Beta}(\alpha, \beta)$. Under the *global risk* model, each agent is informed about all the default probabilities. Under the *graded risk* model, the information available to agent u about δ_v depends on u and v 's proximity in an underlying social network. Specifically, agent u receives a signal in the form of S_{uv} samples, each a realization of the binary event with probability δ_v . The number of samples S_{uv} decreases exponentially with social network distance.

The prior study [14] produced several theoretical results related to the game described above. For example, that work considered a *dichotomous risk* model, where all transactions are between neighbors in an underlying social network, and the credit network is restricted to allow payments only on direct edges (i.e., no routing along paths). In such scenarios, agents in equilibrium issue credit up to their allowed budgets, and do so in a socially optimal manner. In a version of the global risk model (as above but with constraints on credit budgets and further assumptions about the transaction probabilities Λ), the equilibrium credit networks are star-shaped structures where everyone issues credit to the same agent.

Covering more relaxed scenarios, such as those with more flexible Λ , unconstrained budgets, 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 multi-dimensional. 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 resorted in that prior study to simulation analysis of environments relaxing the conditions for which we had theorems. Here we extend that simulation-based analysis, confirming most of its initial conclusions and producing further results, based on a more extensive and accurate empirical modeling exercise.

IV. EMPIRICAL GAME-THEORETIC ANALYSIS

Our investigation of the credit network formation game employs an approach called *empirical game-theoretic analysis* (EGTA) [15]. In EGTA, techniques from simulation, search, and statistics combine with game-theoretic concepts to characterize strategic properties of a domain.

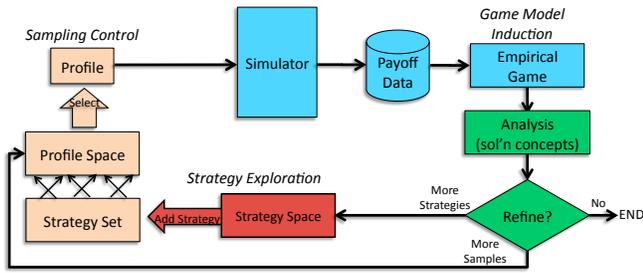


Fig. 3. Iterative procedure for empirical game-theoretic analysis.

A. Iterative EGTA Process

A high-level view of the EGTA process is presented in Figure 3. 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 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 IV-B.

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 the profile space grows exponentially with the lesser of number of players n and number of strategies S . There are $\binom{n+S-1}{n}$ distinct profiles, to be precise. For even moderate n and S , therefore, we generally cannot afford to evaluate every profile through simulation. We thus require analysis techniques that operate on incompletely specified games.

The equilibrium analysis procedure we employ in this study works as follows. We first perform an initial set of simulations, including a modest number of strategies in all combinations. We then iterate the following steps.

- 1) Identify *complete subgames*, defined as sets of strategies for which we have simulated all profiles.
- 2) For each of these, search for symmetric mixed-strategy Nash equilibria (SMSNE). We employ replicator dynamics for this purpose, from a diverse set of starting points.

²More sophisticated approaches may generalize from simulation data using regression or other machine learning techniques [16], [17].

- 3) For each SMSNE candidate found, check the strategies s' outside the subgame for which we have evaluated all profiles where one player plays s' and the other $n - 1$ play strategies in the subgame. If we find a case where s' is a beneficial deviation, the candidate is *refuted*. If all possibly deviating strategies have been evaluated without refutation, the candidate is *confirmed*.
- 4) If an SMSNE candidate is neither refuted nor confirmed, simulate the profiles necessary to check another strategy s'' outside the subgame. Repeat from Step 1.
- 5) If there exists a subgame with a refuted SMSNE candidate, such that the support of this candidate plus the best-response refuting strategy is not subsumed by any complete subgame, simulate the profiles necessary to complete that subgame. Repeat from Step 1.

B. 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 1953 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 [19].

In the current study, we employ a recently introduced technique called *deviation-preserving reduction* (DPR) [18]. 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. Figure 4 displays an example, where the payoff to the first player in a four-player DPR profile is derived from its payoff in a 100-player profile where the other players are divided 33:66 among the other two represented strategies.

In the current study, we focused our analysis on a six-player reduced game derived from simulations on 61-agent credit networks. 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 $\langle 1 \times s_1, 36 \times s_2, 24 \times s_3 \rangle$,

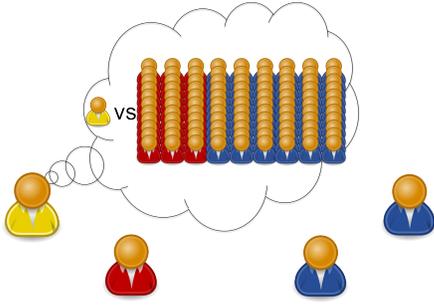


Fig. 4. Correspondence between a 100-player full game profile and the payoff to the first player in a four-player DPR profile.

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 of $n = 61$ for this study). The DPR technique, however, is defined more generally for non-divisible reduction factors, as well as for games that are symmetric only within roles [18].

V. CREDIT NETWORK EGTA STUDY

The current EGTA investigation of credit network formation extends the earlier study in three primary respects.

- 1) It incorporates a larger set of strategies, including variation on the extent and quantity of credit issued. We omit a couple of hybrid strategies, which makes the results somewhat more straightforward to interpret.
- 2) The current study employs deviation-preserving reduction (Section IV-B) rather than hierarchical reduction [19]. As argued above, we expect the DPR technique to produce a more reliable approximation of the full game.
- 3) Here we measure welfare and credit allocation for identified equilibria, and compare them to central (socially optimal) solutions.

A. Setup

As noted above, we consider a population of 61 agents. 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 is drawn uniformly, $x_{uv} \sim U[1, \bar{x}]$, with \bar{x} set to either 1.2 (*low value*) or 2 (*high value*). Cost to the seller is constant: one. The average surplus per transaction is thus either 0.1 or 0.5. Default probabilities δ_v for each agent are drawn from a Beta distribution: $Beta(1, 9)$ (average default probability 0.1) in the *low default* setting, and $Beta(1, 1)$ (average 0.5) in the *high default* setting. In the global risk environment these default probabilities are revealed to all, whereas in the

TABLE I
ENVIRONMENTAL SETTINGS AND PROFILES COLLECTED.

Risk model	Default prob	Buyer surplus	Full-game profiles	DPR profiles
Global	low	low	4619	1497
Global	low	high	2179	765
Global	high	low	3134	1045
Global	high	high	3202	1101
Graded	low	low	991	394
Graded	low	high	5377	1824
Graded	high	low	656	261
Graded	high	high	430	201

graded risk environment each agent gets sample data from the default distribution of others, with the number of samples S_{uv} determined by the social network distance between u and v . The social network itself is an Erdős-Rényi graph. We take $S_{uv} = 100$ if u and v are neighbors, $S_{uv} = 10$ if they are linked through one other node, $S_{uv} = 1$ if they have a shortest-path of length three, and $S_{uv} = 0$ otherwise.

We explored environments with high or low default and value, for each of global and graded risk. The eight environments are listed in Table I, along with the number of profiles we ended up simulating, in both the full and reduced games.

We considered a range of heuristic strategies available to agents. A 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 below, defined from the perspective of agent u 's evaluation of credit prospect v :

- **Default probability:** lowest known default (δ_v) for global risk, or lowest estimated default based on samples S_{uv} for graded risk. $(k, q) : (1, 1), (2, 2), (3, 2), (5, 2)$
- **Buy rate:** highest probability of transacting (λ_{uv}). $(k, q) : (1, 1), (1, 2), (2, 2)$
- **Sell rate:** highest probability of serving a transaction (λ_{vu}). $(k, q) : (2, 2)$
- **Trade value:** highest expected value of transaction per event ($\lambda_{uv}x_{uv}$). $(k, q) : (2, 2), (5, 2)$
- **Trade profit:** highest difference, expected value of transaction minus expected value of served transaction ($\lambda_{uv}x_{uv} - \lambda_{vu}$). $(k, q) : (2, 2), (6, 2), (8, 1)$
- **Index:** lowest node number (arbitrary global labeling). $(k, q) : (1, 1), (2, 2)$
- **Random:** uniform choice. $(k, q) : (2, 2)$

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

Altogether, we considered 17 strategies: four instances of **Default**, three **Buy rate**, etc., as described above. There are 1.4×10^{16} distinct strategy profiles for the full 61-player game, and 74,613 for the six-player DPR game. As indicated in Table I, we evaluated only a very small fraction of these

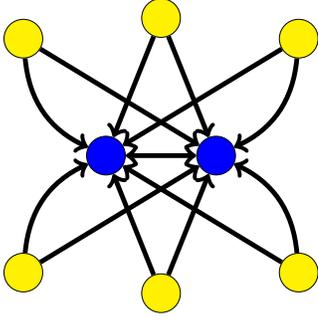


Fig. 5. A star-like credit network arising from the mutual application of default-based strategies in the global risk model.

profiles, but nevertheless are able to identify equilibria in each environment.

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 defaulters, they are creating edges to the same target agents. This leads to a centralized or star-like credit network, as illustrated in Figure 5. 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. By comparing **Default** and **Index** strategies we can separate the pure benefits of coordination from the benefits of avoiding defaulters.

Altogether we evaluated over 20,000 strategy profiles across the eight credit network environments. Each 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 [20].

B. Results

Through the process described in Section IV-A, we successfully derived equilibria for each of the eight credit network games. Specifically, we identified between one and three SMSNEs for the reduced six-player DPR games corresponding to each environment. All candidate subgame equilibria were either confirmed or refuted by the process,

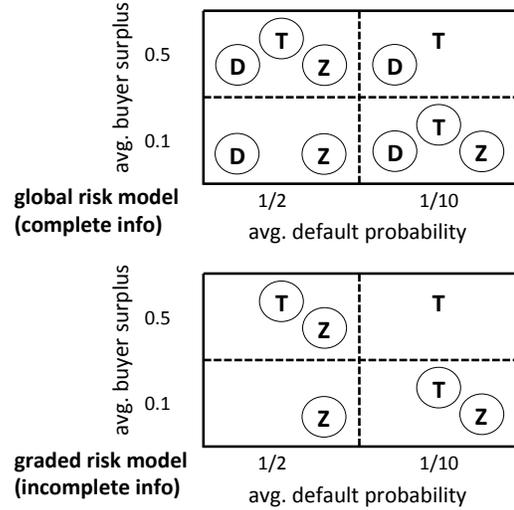


Fig. 6. Equilibria found for the eight credit network environments. Letters denote the strategy classes represented in equilibrium, with circled letters indicating pure-strategy equilibria.

and the subgames covering best responses to all candidates were completed.

The strategies **Sell rate**, **Index**, and **Random** are not supported in any equilibria. 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 6. 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 that class was confirmed as a pure strategy Nash equilibrium.

From the figure, we see that there is a no-credit equilibrium in six of the eight environments: all but those with low default and high buyer value. The least favorable environment—graded risk with high default and low value—has 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. The weakened information about defaults plus the lack of coordinating power render this strategy ineffective in graded risk environments.

For completeness, we list the equilibria found. Environments are coded by risk model (C[omplete information] for global risk, I[ncomplete] for graded risk), default probability (L[ow]/H[igh]), and buyer value (L[ow]/H[igh]). Groups in brackets with probabilities represent SMSNEs, and ungrouped strategies indicate pure equilibria.

CLL **Default**(2,2); **Buy rate**(1,1); **Zero**

CLH **Default**(3,2); [**Trade profit**(2,2), 0.248; **Trade profit**(8,1), 0.457; **Trade value**(2,2), 0.295]

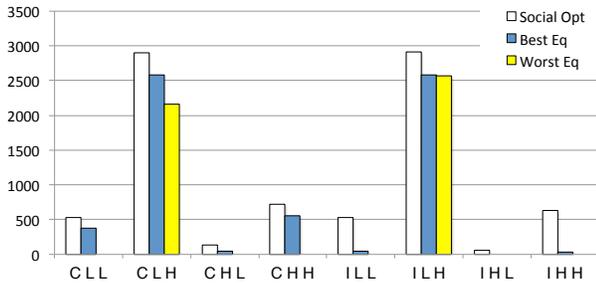


Fig. 7. Welfare at social optimum compared to welfare at equilibrium.

CHL Default(1,1); Zero
 CHH Default(2,2); Buy rate(1,1); Zero
 ILL Buy rate(1,1); Zero
 ILH [Trade profit(2,2), 0.576; Trade profit(8,1), 0.424];
 [Buy rate(2,2), 0.271; Trade profit(8,1), 0.453;
 Trade value(2,2), 0.276]
 IHL Zero
 IHH Buy rate(1,1); Zero

We next turn to the question: How well do the credit networks generated in equilibrium perform? Figure 7 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 six of eight 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), the equilibrium network formation does a good job of obtaining most of it. For less favorable environments (the least favorable global risk environment, and all but the best graded risk environment), 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 8, 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.

All of these results are of course relative to the particular selection of 17 strategies 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. We expect our qualitative conclusions will prove fairly robust, but have no doubt that additional strategy options would lead to changes or additions to the set of equilibria.

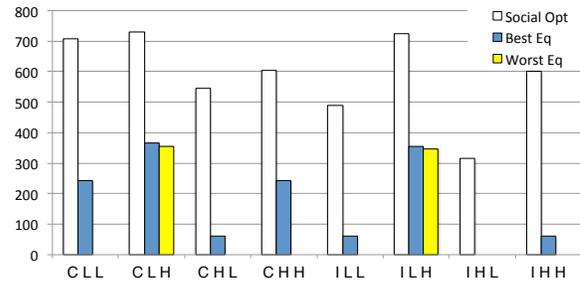


Fig. 8. Total units of credit issued at social optimum compared to equilibrium.

VI. CONCLUSION

Our EGTA study of credit network formation suggests that creation of viable credit networks by self-interested autonomous agents is indeed plausible, if the environment is sufficiently favorable, in terms of default rates and profitability. When it is not so favorable, the failure to create a viable credit network does not sacrifice a great deal of welfare anyway. The generated networks fall short of socially optimal, which is expected given the positive externalities of credit issuance. Whereas some suboptimality may be acceptable, it is also worth considering the design of incentives for credit issuers, for example the possibility of interest or other payments from the receivers and users of credit.

There remain many other interesting issues in credit network formation for future research. One is to explore dynamic models, where credit balances may be raised or lowered (perhaps with restrictions) over time, based on changes in information or actual circumstances. Another is to apply this framework to modeling actual financial credit. How to accomplish this is yet to be worked out, though some of the prior works [5], [6] expressly contemplate interest payments and distinguishing credit limits and credit balances. Combining these two thrusts, a dynamic model of financial credit networks could potentially shed light on systemic risks and their relation to financial crisis phenomena.

ACKNOWLEDGMENT

This research was supported in part by the Army Research Laboratory, under Cooperative Agreement W911NF-09-2-0053. Pranav Dandekar and Ashish Goel collaborated in the conception and design of this study, and Ben-Alexander Cassell extended our EGTA infrastructure to facilitate this work.

REFERENCES

- [1] V. Bala and S. Goyal, "A noncooperative model of network formation," *Econometrica*, vol. 68, pp. 1181–1229, 2000.
- [2] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker, "On a network creation game," in *22nd ACM Symposium on Principles of Distributed Computing*, 2003, pp. 347–351.
- [3] M. O. Jackson and A. Wolinsky, "A strategic model of social and economic networks," *Journal of Economic Theory*, vol. 71, pp. 44–74, 1996.
- [4] C. Borgs, J. Chayes, J. Ding, and B. Lucier, "The hitchhiker's guide to affiliation networks: A game-theoretic approach," in *Innovations in Computer Science*, Beijing, 2011, pp. 389–400.

- [5] D. B. DeFigueiredo and E. T. Barr, "TrustDavis: A non-exploitable online reputation system," in *Seventh IEEE International Conference on E-Commerce Technology*, Washington, DC, 2005, pp. 274–283.
- [6] A. Ghosh, M. Mahdian, D. M. Reeves, D. M. Pennock, and R. Fugger, "Mechanism design on trust networks," in *Third International Workshop on Internet and Network Economics*, 2007, pp. 257–268.
- [7] D. Karlan, M. Mobius, T. Rosenblat, and A. Szeidl, "Trust and social collateral," *Quarterly Journal of Economics*, vol. 124, pp. 1307–1361, 2009.
- [8] A. Mislove, A. Post, P. Druschel, and K. P. Gummadi, "Ostra: Leveraging trust to thwart unwanted communication," in *Fifth Usenix Symposium on Networked Systems Design and Implementation*, San Francisco, 2008, pp. 15–30.
- [9] Z. Liu, H. Hu, Y. Liu, K. W. Ross, Y. Wang, and M. Mobius, "P2P trading in social networks: The value of staying connected," in *29th IEEE International Conference on Computer Communications*, 2010, pp. 1–9.
- [10] F. Limpens and D. Gillet, "A competence bartering platform for learners," in *Advances in Web-Based Learning*, ser. Lecture Notes in Computer Science, vol. 7048. Springer, 2011, pp. 148–153.
- [11] P. Dandekar, A. Goel, R. Govindan, and I. Post, "Liquidity in credit networks: A little trust goes a long way," in *Twelfth ACM Conference on Electronic Commerce*, San Jose, 2011, pp. 147–156.
- [12] P. Resnick and R. Sami, "Sybilproof transitive trust protocols," in *Tenth ACM Conference on Electronic Commerce*, Stanford, CA, 2009, pp. 345–354.
- [13] B. Viswanath, M. Mondal, K. P. Gummadi, A. Mislove, and A. Post, "Canal: Scaling social network-based Sybil tolerance schemes," in *European Conference on Computer Systems (EuroSys)*, 2012, pp. 309–322.
- [14] P. Dandekar, A. Goel, M. P. Wellman, and B. Wiedenbeck, "Strategic formation of credit networks," in *Twenty-First International WWW Conference*, Lyon, 2012, pp. 559–568.
- [15] M. P. Wellman, "Methods for empirical game-theoretic analysis (extended abstract)," in *Twenty-First National Conference on Artificial Intelligence*, Boston, 2006, pp. 1552–1555.
- [16] Y. Vorobeychik, M. P. Wellman, and S. Singh, "Learning payoff functions in infinite games," *Machine Learning*, vol. 67, pp. 145–168, 2007.
- [17] P. R. Jordan and M. P. Wellman, "Generalization risk minimization in empirical game models," in *Eighth International Conference on Autonomous Agents and Multi-Agent Systems*, Budapest, 2009, pp. 553–560.
- [18] B. Wiedenbeck and M. P. Wellman, "Scaling simulation-based game analysis through deviation-preserving reduction," in *Eleventh International Conference on Autonomous Agents and Multi-Agent Systems*, Valencia, 2012, pp. 931–938.
- [19] M. P. Wellman, D. M. Reeves, K. M. Lochner, S.-F. Cheng, and R. Suri, "Approximate strategic reasoning through hierarchical reduction of large symmetric games," in *Twentieth National Conference on Artificial Intelligence*, Pittsburgh, 2005, pp. 502–508.
- [20] B.-A. Cassell and M. P. Wellman, "EGTAOnline: An experiment manager for simulation-based game studies," in *AAMAS-12 Workshop on Multi-Agent Based Simulation*, Valencia, 2012.