

Strategic Payment Routing in Financial Credit Networks

FRANK CHENG, University of Michigan

JUNMING LIU, University of Michigan

KAREEM AMIN, University of Michigan

MICHAEL P. WELLMAN, University of Michigan

Credit networks provide a flexible model of distributed trust, which supports transactions between untrusted counterparties through paths of intermediaries. We extend this model by introducing interest rates (prices on lines of credit), both as a means to incentivize credit issuance and to provide a framework for modeling networks of financial relationships. Including interest rates poses a new constraint on transactions, as intermediaries will route payments only if the interest received covers any interest paid. We account for these constraints in an efficient algorithm for finding the maximum transaction flow between two agents in a financial network. There are generally many feasible payment paths serving a given transaction, and we show that the policy for selecting among such paths can have a substantial effect on liquidity, as measured by steady-state probability of transaction success. Finally, we consider the situation where the transaction source can choose among heuristic path selection mechanisms, in order to maximize their payoff. Through empirical game-theoretic analysis, we find that routing is inefficient due to the positive externality of choices promoting network liquidity. However, agent choices do reflect some consideration of overall network liquidity, in addition to their own interest payments.

General Terms: Economics

Additional Key Words and Phrases: Credit networks, financial networks, payment routing

1. INTRODUCTION AND PRIOR WORK

The key functions of a financial system are to allocate capital to productive uses and support transactions across a heterogeneous set of agents. These functions often interact, for example through institutions (banks) that lend capital, which they can do in part by maintaining deposit accounts for which they provide payment services. At the core of these functions is the management of financial obligations between parties. We therefore view expression of such obligations as prerequisite to comprehensive financial modeling, and introduce here a model of *financial credit networks* (FCNs) based on these relations. Our focus in this paper is on payment operations, which serve as a foundation for general economic transactions such as the purchase of routine products and services, lending and saving, and capital investment.

We extend an existing abstract model of *credit networks*: weighted directed graphs that represent the capacity of agents (each represented as a node in the graph) to transact with each other. The credit network model was proposed independently by several distinct groups of researchers who were motivated to capture distributed trust in different contexts. Ghosh et al. [2007] aimed to support distributed payment and

This work was supported in part by the National Science Foundation under grant IIS-1440360.

Authors' address: Computer Science & Engineering, University of Michigan, Ann Arbor, USA; email: {frcheng,liujm,amkareem,wellman}@umich.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands.

ACM 978-1-4503-3936-0/16/07 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/2940716.2940738>

multi-user credit checking for multi-item auctions. Karlan et al. [2009] wanted to construct an economic model of informal borrowing networks. In both of these cases, the concept of credit is financial in nature. In two other cases [DeFigueiredo and Barr 2005; Mislove et al. 2008], credit serves as an accounting mechanism to limit computational actions. Moreno-Sanchez et al. [2015] recently demonstrated a privacy-preserving payment protocol for credit networks, and also suggest that some real distributed payment systems like Ripple are based on the credit network model.

The effectiveness of credit networks for distributed transactions was most powerfully demonstrated by Dandekar et al. [2011], 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. In particular they showed for several classes of graphs, 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.

A follow-up study addressed the question of whether self-interested agents would issue sufficient credit to form such high-performing networks [Dandekar et al. 2015]. Issuing credit entails a tradeoff between increasing the prospect of valuable transactions at the cost of exposure to risk of defaulting counterparties. The study found, across a range of experimental environments, that if there is sufficient transaction profit to be earned, a network will form to extract a sizable fraction of that surplus. However, the credit networks formed in equilibrium are still suboptimal, as we might expect given the positive externality in credit issuance. Namely, when an agent issues a credit line, the entire network benefits from the liquidity it provides while only the issuer bears the risk that the borrower may default.

To provide proper incentives to issue credit, a natural approach is to allow creditors to charge interest on outstanding obligations. Pricing credit is of course standard in actual financial systems, and so extending the credit network formalism to support interest charges would also make them more suitable as a modeling substrate for this domain. The extension to FCNs developed here is driven primarily by the requirement for handling interest rates. We show how the representation of obligations must be refined to accommodate interest on outstanding debt, and demonstrate how to realize payment operations in the extended model.

In §2 we introduce the basic concepts of FCNs, through an extended example. §3 presents the formal FCN model, and defines what constitutes a feasible solution to the payment routing problem. Next (§4) we give a polynomial algorithm that solves for maximum flow on financial credit networks when interest rates are restricted to the set of contract interest rates. We then show that this is equivalent to solving the problem under unrestricted interest rates. In §5 we define several payment mechanisms that select among multiple feasible routing solutions. In §6 we evaluate the liquidity of each mechanism experimentally and find that there is a positive relationship between how much interest agents pay and how much liquidity is available in the network. We find that several mechanisms exhibit liquidity performance close to the no interest rate setting. The ability to pay interest rates above that offered by lenders plays a large role in this good performance. Finally (§7), we consider strategic choice by payers over routing mechanisms. We find that the socially optimum payment mechanism is not strategically stable, but that agents are willing to pay a limited amount in exchange for maintaining liquidity on a financial credit network.

2. FINANCIAL CREDIT NETWORKS

A prerequisite for any payment is *trust*. Transactions between unknown parties are often enabled by third parties (e.g., banks or credit card issuers) that assure successful execution by mediating a transfer of obligations tantamount to a flow of funds. In the early days of Internet commerce, platforms like eBay introduced reputation mechanisms to facilitate the development of trust necessary to overcome lack of direct experience with counterparties [Resnick and Zeckhauser 2002]. Others provided more direct mediation. For example, Alibaba offered an escrow service whereby two parties that mutually trust Alibaba can transact with each other, relying on Alibaba to make them whole if the counterparty defaults [Yu and Shen 2016].

In the basic credit network formalism, a 's trust of b is represented by a directed weighted edge from a to b , where weight w denotes the **capacity** of credit a offers to b . This credit is interpreted as an obligation for a to accept up to w units of IOUs from b in exchange for commensurate service. These IOUs may be returned by a to b at a later time, in exchange for service from b . The real power of credit networks, though, comes from transactions along paths, achieving an effective transitivity of trust. If a offers credit to b and b to c , then c can transact with a by routing its payment through b : specifically, for a unit transaction c sends one of its IOUs to b , and b sends one of its IOUs to a . The net result is a payment of one unit from a to c . Node b has exchanged one its own IOUs for one of a 's, and thus its balance of obligations is unchanged.

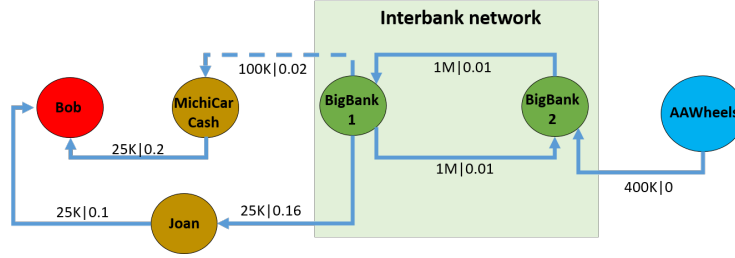
2.1. Illustrative Example

Chains of payment are common occurrences in everyday commerce. Suppose Bob wishes to buy a new car from his local dealership, AAWheels. He negotiates a deal to purchase his favorite model for \$25,000. Bob however does not carry that much cash, and AAWheels does not trust him directly. Anticipating this issue, before car-shopping Bob had applied for a credit line from MichiCarCash, a prominent consumer lender, who after some research decided to issue a credit line of \$25,000 at a healthy interest rate (20%). MichiCarCash has a checking account with BigBank1, with a current balance of \$100,000. This deposit is essentially a loan to the bank (at a relatively smaller interest rate, 2%), so we can think of MichCarCash as holding 100,000 BigBank1 IOUs. AAWheels maintains a no-interest checking account with BigBank2, currently with zero balance, however it is willing to hold up to \$400,000 there. This can be represented as a credit line from AAWheels to BigBank2.


The situation as described thus far comprises part of the FCN depicted in Figure 1a. BigBank1 and BigBank2 are connected through the interbank network, which for present purposes we model as a complete subgraph of high-capacity credit lines with nominal interest rates (1%). In the figure, credit edges are indicated with solid arrows, and holdings of IOUs (i.e., actual loans) are indicated with dashed arrows. Whereas in the original credit network formalism credit lines and IOU holdings are treated uniformly, in FCNs we must distinguish them because IOUs accrue interest payments and unused credit lines do not. The rates of interest associated with credit and debt edges are annotated along with the capacities.

To purchase the car, Bob routes a payment of \$25,000 to AAWheels. He can do so by drawing on his credit line with MichiCarCash, who in turn returns a like number of IOUs to BigBank1 (i.e., withdraws from its BigBank1 checking account), which then sends its own IOUs through the interbank network¹ to BigBank2, which credits the checking account of AAWheels (i.e., grants AAWheels 25,000 BigBank2 IOUs). As in

¹We do not model the actual network here, but capturing the structure of the broader financial system is a long-term goal motivating this research. For instance, we could model cash as IOUs from a central bank, which everyone trusts with high capacity at zero interest.



(a) Initial credit network. Bob has credit lines of 25K (\$25,000) from MichiCarCash and Joan. MichiCarCash has a checking account with 100K at BigBank1, and Joan carries a credit card with 25K limit at that bank. BigBank1 and BigBank2 offer large low-interest credit lines to each other through the interbank network, and AAWheels is willing to maintain up to 400K in a no-interest account at BigBank2.



(b) Credit network after 25K payment from Bob to AAWheels through MichiCarCash. The credit line from Joan remains unused, but is discussed as an alternative in §2.2.

Fig. 1: Financial credit networks before and after the transaction. Arrows indicate obligations: solid for credit and dashed for IOUs. An arrow from x to y with marking $c | r$ denotes an obligation from x to y with capacity c at an interest rate r . The direction of payment flow is against the arrows.

the basic credit network $a \rightarrow b \rightarrow c$ example above, the source has decreased its net capacity balance by the payment amount, and the destination accrues a corresponding increase. The FCN after executing this payment is shown in Figure 1b.

2.2. Example Continued: Interest Considerations

We can verify from the figure that intermediate nodes on the payment path experience no change in net capacity. For FCNs, however, we must also consider the effect of interest rates. MichiCarCash receives new IOUs from Bob, in exchange for returning BigBank1 IOUs. It is happy to do so though, because the Bob IOUs carry a higher interest rate than the BigBank1 (interest checking) rate. Similarly, BigBank1 accepts the payment routing because the rate on its checking IOUs exceeds the interbank rate. Finally, BigBank2 effectively gets the interbank rate on the BigBank1 IOUs, and pays no interest on the checking account of AAWheels. The payment path satisfies *interest rate monotonicity*, and so we deem it *feasible*; no agent routes this payment at a loss.

We next consider other potential payment paths in this network. Bob has a friend Joan, who also (coincidentally) trusts him for \$25,000 and requires only 10% interest. Joan, however, has no bank deposits, and her only way to route a payment through the banking system is to use her BigBank1 credit card, which carries a 16% interest rate and \$25,000 credit limit. Routing a payment through Joan at the *contract rate* (interest rates associated with issued credit lines) would violate monotonicity, and is therefore infeasible. However, if Bob were to borrow from Joan at a higher rate (say 16%), she would cover her credit card interest and still offer Bob a better payment deal

than he is getting from MichiCarCash. This keeps everyone satisfied, so we consider paths that use credit lines above contract rates to be feasible as well.

Suppose instead of this credit line, Bob had been holding 25,000 IOUs from Joan at a 10% interest rate. Joan cannot be expected to take these IOUs back to route a payment that will cost her a higher rate. Moreover, unlike the situation with credit lines, there is no obvious way to restore monotonicity by increasing a rate. Bob has no means to compensate Joan other than by returning her IOUs (e.g., he cannot route her a side payment), which merely saves her interest at the specified rate. We therefore consider interest on IOUs to be fixed at the rate at which they were incurred, and disallow payments along paths where a debt link would violate interest rate monotonicity.

3. MODEL

Credit networks are formally specified by a capacitated graph $G = (V, E)$ where vertices represent agents, and edges represent credit relationships. A directed edge $e = (u, v)$ of capacity $c(e)$ indicates that u is willing to lend up to $c(e)$ units to v . Dandekar et al. [2011] observed that to determine whether agent s can pay agent t a sum of $X > 0$ units, one needs only check whether there exists an $s - t$ flow of value X in the network derived by flipping the direction of edges in G .

To extend the basic framework with interest rates, we must first capture the additional constraint that intermediary agents not incur a net interest cost. Second, the introduction of interest requires that we distinguish capacity on outstanding credit lines from capacity based on actually incurred debts (the **IOU holdings**).

3.1. Credit Networks with Interest Rates

Formally, a *financial credit network* (FCN: credit network with interest rates) is specified by a directed network $G = (V, E)$. Each edge $e \in E$ is associated with three values. The first value is e 's obligation type, indicated by $\tau : E \rightarrow \{\text{credit}, \text{iou}\}$. Second is the edge's capacity, given by the function $c : E \rightarrow \mathbb{R}^+$. As in the interest-free setting, edge e 's capacity indicates the credit limit associated with that edge (if $\tau(e) = \text{credit}$), or the number of IOUs held, (if $\tau(e) = \text{iou}$). Finally, a third value $r : E \rightarrow \mathbb{R}^+$ represents the edge's **contract interest rate**. The values τ, c, r at an edge e can be interpreted as follows. An edge $e = (u, v)$ with $\tau(e) = \text{credit}$, means that u is willing to lend $c(e)$ units to v at an interest rate of $r(e)$. Edge e with $\tau(e) = \text{iou}$ means that u owes $c(e)$ units to v , on which v is charging an interest rate of $r(e)$.

One reason that credit edges must be distinguished from iou edges is to assess periodic interest obligations. Another is that credit lines provide greater flexibility on interest rates than do IOU holdings. For a credit line, willingness to lend at an interest rate $r(e)$ implies willingness to lend at any interest rate $r' > r(e)$. In contrast, the interest rate on debt represented by an IOU edge is fixed once the debt is incurred.

This flexibility turns out to be valuable for maximizing the liquidity of the network. In particular, our model assumes that intermediaries along a payment path will not route a payment at a loss. This amounts to requiring that payments be routed along paths of *monotonically nonincreasing* interest rates. By paying a larger interest rate on one of its immediate credit lines, an agent s may be able to route a payment to another agent t that would otherwise be infeasible.

In describing our routing algorithm, it is convenient to refer to the **reverse network** of G where edges point in the direction that payments occur, rather than the direction of credit obligations. Define $G^\dagger = (V, E^\dagger)$, the network identical to G but with edges flipped. When applied to edges in E^\dagger , the functions τ, c and r take the same values as they would on the unflipped edges in E . We also define for each $v \in V$, $\text{In}(v) = \{(u, v) \mid$

$(u, v) \in E^\dagger$ and $\text{Out}(v) = \{(v, u) \mid (v, u) \in E^\dagger\}$. Note that G is in fact a multigraph, as there may be many edges between two vertices, differing in types or interest rates.²

3.2. Payment Routing Problem

Given a financial credit network $N = \{G, \tau, c, r\}$, we now define what it means for a payment of $X > 0$ units from agent s to t to be **feasible**. We call the problem of determining whether a payment is feasible the **payment routing problem**.

As in an interest-free credit network, we demand that there exists an $s - t$ flow $f : E^\dagger \rightarrow \mathbb{R}^+$ of value at least X in G^\dagger , where f respects the capacity constraints denoted by c . In order for an FCN payment to be feasible, we also impose requirements on the interest rates. Informally, we demand that payments between s and t flow along edges of non-increasing interest rate, which we call the **monotonicity condition**. This condition is complicated by the fact that realized interest rates are not always equal to contract rates. Whereas the contract interest rates on iou edges are fixed, rates on credit edges can be increased to satisfy monotonicity. Payments that respect this invariant are said to have **valid realized interest rates**.

Let a sequence of $s - t$ paths $P = (P^{(1)}, \dots, P^{(K)})$ and corresponding nonnegative real numbers $F = (f^{(1)}, \dots, f^{(K)})$ be a **consistent decomposition** of a flow f if for any edge e , $f(e) = \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] f^{(k)}$. In other words, f can be thought of as the finite union of paths in G^\dagger , where each path $P^{(k)}$ is assigned flow value $f^{(k)}$.

For each such path $P^{(k)}$, which we take as a set of edges, we can define a function $\hat{r}^{(k)} : P^{(k)} \rightarrow \mathbb{R}^+$ which assigns a **realized interest rate** to each edge along the path. For example, if edge (u, v) belongs to path $P^{(k)}$, $\hat{r}^{(k)}((u, v))$ corresponds to the interest rate charged to u by v on the $f^{(k)}$ units of payment routed along edge e in path $P^{(k)}$. Consistent decompositions can be used to state the aforementioned monotonicity and validity conditions more formally.

DEFINITION 1 (VALIDITY). *Given a consistent decomposition P, F , for some flow f , and realized interest rates $\{\hat{r}^{(k)}\}$, we say that the realized interest rates are valid if, $\forall k \in \{1, \dots, K\}$ and $e \in P^{(k)}$, $\hat{r}^{(k)}(e) = r(e)$ if $\tau(e) = \text{iou}$ and $\hat{r}^{(k)}(e) \geq r(e)$ if $\tau(e) = \text{credit}$.*

DEFINITION 2 (MONOTONICITY). *Given a consistent decomposition P, F , for some flow f , and realized interest rates $\{\hat{r}^{(k)}\}$, we say that the realized interest rates are monotonic if, $\forall k \in \{1, \dots, K\}$ with $P^{(k)} = (e_1^{(k)}, e_1^{(k)}, \dots, e_{n_k}^{(k)})$, and $i \in \{0, \dots, n_k - 1\}$, $\hat{r}^{(k)}(e_i) \geq \hat{r}^{(k)}(e_{i+1})$.*

Suppose there exists an $s - t$ flow f of value X that admits a decomposition P, F with valid and monotonic realized interest rates $\{\hat{r}^{(k)}\}$. A payment of X units can be executed by routing, for each k , $f^{(k)}$ units along the $s - t$ path $P^{(k)}$. Assigning each edge e on that path an interest rate of $\hat{r}^{(k)}(e)$ ensures that no agent along the path $P^{(k)}$ is routing this payment at a loss, as a consequence of monotonicity. Finally, the validity of $\hat{r}^{(k)}$ ensures that the realized interest rates are consistent with existing rates on IOUs, and no smaller than the contract rate on credit lines employed.

DEFINITION 3 (FEASIBLE PAYMENT). *For FCN N with $s, t \in V$, we say that a payment of amount $X > 0$ from source s to destination t is feasible if there exist (1) an $s - t$ flow $f : E^\dagger \rightarrow \mathbb{R}^+$ of at least X , in the standard flow network defined by G^\dagger and c ,*

²Thus, the use of the functions $\tau(e), r(e), c(e)$ is a slight abuse of notation. A single pair $e = (u, v)$ may, for example, be assigned multiple types. An edge in G is more accurately written as a 5-tuple (u, v, c, r, τ) . Nevertheless, for clarity, we write τ, r and c unless this distinction is important.

and (2) a consistent decomposition P, F of f , and corresponding realized interest rates $\{\hat{r}^{(k)}\}$ that are valid and monotonic.

4. PAYMENT ROUTING ALGORITHM

Given an instance of the payment routing problem (N, s, t, X) , our goal is to find a flow f of value at least X in G^\dagger , and corresponding realized interest rates that are valid and monotonic, or output that no such flow exists. In general, realized interest rates may be any nonnegative real number. We start by restricting the rates to come from some finite set $\mathcal{I} = \{I_1, \dots, I_l\} \subset \mathbb{R}^+$. We show that this \mathcal{I} -restricted payment routing problem can be solved by via linear programming. We then observe that taking \mathcal{I} to be the set of contract interest rates $\mathcal{I} = \{r(e) \mid e \in E\}$ suffices to solve the *unrestricted* payment routing problem. In other words, payment (N, s, t, X) is feasible if and only if it is feasible with realized interest rates restricted to the set of initial contract rates.

4.1. \mathcal{I} -Restricted Payment Routing : Monotonicity

The key construct for solving the \mathcal{I} -restricted routing problem is a function $f' : E^\dagger \times \mathcal{I} \rightarrow \mathbb{R}^+$, which we call an **interest-flow**. We can think of f' as indicating for each edge $e \in E^\dagger$ and interest rate level $I \in \mathcal{I}$ the amount of payment along edge e routed at a realized interest rate level I . Cast in this way, the problem is reminiscent of a *multicommodity flow* [Even et al. 1975], where payment routed at an interest rate I corresponds to a distinct commodity. Such problems are known to be NP-hard. In our problem, however, a unit of flow entering a vertex v at interest rate I can exit as a unit of flow at any interest rate $I' \leq I$ (i.e., monotonicity). This relaxation proves crucial for developing a polynomial time algorithm.

We say an interest flow is *valid* in the analogous way to realized interest rates. Namely f' is valid if for any e, I , $f'(e, I) \geq 0$ only if $\tau(e) = \text{iou}$ and $I = r(e)$ or $\tau(e) = \text{credit}$ and $I \geq r(e)$. The main result of this section establishes that the monotonicity of the realized interest rates can be compactly represented as a collection of local inequalities on the interest-flow. We see that monotonicity is equivalent to the requirement that for every interest rate level I , and vertex v , the flow into v at interest rates $I' \leq I$ is no more than the flow out of v at interest rates $I' \leq I$.

LEMMA 1. *Let $f : E^\dagger \rightarrow \mathbb{R}^+$ be an arbitrary $s - t$ flow. f admits a consistent decomposition, along with a set of \mathcal{I} -restricted, valid, monotonic realized interest rates if and only if there exists a valid interest-flow satisfying the following conditions.*

$$\text{For every } e \in E^\dagger: \quad \sum_{I \in \mathcal{I}} f'(e, I) = f(e), \text{ and} \quad (1)$$

for every $I \in \mathcal{I}$ and vertex $v \in V, v \notin \{s, t\}$:

$$\sum_{e \in \text{In}(v)} \sum_{I' \in \mathcal{I}, I' \leq I} f'(e, I') \leq \sum_{e \in \text{Out}(v)} \sum_{I' \in \mathcal{I}, I' \leq I} f'(e, I'). \quad (2)$$

PROOF. In the first direction, we suppose that f admits a consistent decomposition with valid, monotonic interest rates, then show that there exists a valid interest-flow satisfying the conditions (1) and (2). Fix an $s - t$ flow f , a consistent decomposition F, P , and realized interest rates $\{\hat{r}^{(k)}\}$. Define $f'(e, I) \equiv \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] \mathbf{1}[\hat{r}^{(k)}(e) = I] f^{(k)}$. That is, $f'(e, I)$ simply aggregates the flow values along edges e that were assigned interest rate I . We now check condition (2). Fixing a vertex v , and interest rate I ,

$$\sum_{e \in \text{In}(v)} \sum_{I' \leq I} f'(e, I') = \sum_{e \in \text{In}(v)} \sum_{I' \leq I} \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] \mathbf{1}[\hat{r}^{(k)}(e) = I'] f^{(k)}. \quad (3)$$

If there exists some path $P^{(k)}$ containing an edge (u, v) with $\hat{r}^{(k)}((u, v)) = I'$, then the very next edge (v, w) in $P^{(k)}$ must satisfy $r^{(k)}((v, w)) \leq I'$, by monotonicity. Thus, for every summand in the right hand side of (3) equal to $f^{(k)}$, there exists a distinct summand in $\sum_{e \in \text{Out}(v)} \sum_{I' \leq I} \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] \mathbf{1}[\hat{r}^{(k)}(e) = I'] f^{(k)}$ also equal to $f^{(k)}$. Since $\sum_{e \in \text{Out}(v)} \sum_{I' \leq I} \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] \mathbf{1}[\hat{r}^{(k)}(e) = I'] f^{(k)} = \sum_{e \in \text{Out}(v)} \sum_{I' \leq I} f'(e, I')$, condition (2) is satisfied.

From the definition of consistent decomposition, we know that $f(e) = \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] f^{(k)} = \sum_{I \in \mathcal{I}} \sum_{k=1}^K \mathbf{1}[e \in P^{(k)}] \mathbf{1}[\hat{r}^{(k)}(e) = I] f^{(k)}$ which is equal to $\sum_{I \in \mathcal{I}} f'(e, I)$ by how we have defined f' , and so (1) is satisfied as well. Finally, the validity of f' follows immediately from the validity of the realized interest rates.

In the other direction, fix f , and let f' be a valid interest-flow satisfying conditions (1) and (2). We use f' to reconstruct a consistent decomposition P, F of f with valid and consistent interest rates. Consider a vertex $v \in V$. We assign interest rates to the flow entering and exiting v according to f' . That is, for each edge e containing v , with flow value $f(e)$, we take $f'(e, I)$ units of that flow and assign it a realized interest rate of I .

If interest rates are assigned in this manner, all flow entering v can be routed out of v while respecting monotonicity. In particular, order the interest rates $I_1 < \dots < I_l$. (2) implies $\sum_{e \in \text{In}(v)} f'(e, I_1) \leq \sum_{e \in \text{Out}(v)} f'(e, I_1)$ which in turn implies that all incoming flow at interest rate I_1 can be routed out of some edge at rate I_1 . Now fix some I_{k-1} and suppose for induction that there is a way to route all incoming flow at $I' \leq I_{k-1}$ out of v while respecting monotonicity. To route all incoming flow at level $I' \leq I_k$, we first assign all flow entering v at level $I' \leq I_{k-1}$ to outgoing edges, which by induction we can do while respecting monotonicity. To monotonically route the $\sum_{e \in \text{In}(v)} f'(e, I_k)$ units of flow entering at exactly level I_k there needs to be enough remaining capacity on the outgoing edges at interest rate $I' \leq I_k$. In other words, it needs to be the case that $\sum_{e \in \text{In}(v)} f'(e, I_k) \leq \sum_{e \in \text{Out}(v)} \sum_{I' \leq I_k} f'(e, I') - \sum_{e \in \text{In}(v)} \sum_{I \leq I_{k-1}} f'(e, I')$. Rearranging we get $\sum_{e \in \text{In}(v)} \sum_{I' \leq I_k} f'(e, I') \leq \sum_{e \in \text{Out}(v)} \sum_{I' \leq I_k} f'(e, I')$ which is implied by (2).

Thus, f can be decomposed into P, F , where the interest rates $\{r^{(k)}\}$ assigned along the paths are derived from the above procedure. By (1) all flow is accounted for, and since f' is valid, $\{r^{(k)}\}$ is also valid. \square

4.2. A Linear Program for \mathcal{I} -Restricted Payment Routing

With Lemma 1 in hand, we can derive a linear program for solving the \mathcal{I} -restricted payment routing problem (Algorithm MaxInterestFlowLP). The first four conditions of the LP make it so that the ordinary flow f derived from the interest-flow f' (via (1)) is both a valid flow and routes X units of payment from s to t . The remaining conditions specify that f' also induces valid, monotonic, interest rates, which is a direct consequence of Lemma 1.

THEOREM 1. *Given an instance (N, s, t, X) of the \mathcal{I} -restricted routing problem, the payment is feasible if and only if the LP solved by Algorithm MaxInterestFlowLP has a solution of value at least X . Furthermore, the LP has a total of $\mathcal{O}(|\mathcal{I}||E|)$ variables and $\mathcal{O}(|\mathcal{I}||V| + |E|)$ constraints.*

4.3. A Solution to the Unrestricted Problem

We now prove that the payment (N, s, t, X) is feasible with unrestricted interest rates if and only if there is a feasible payment for the \mathcal{I} -restricted routing problem, when taking $\mathcal{I} = \{r(e) \mid e \in E\}$, the set of initial contract interest rates. As a consequence, Algorithm MaxInterestFlowLP gives us an algorithm for the unrestricted payment routing problem.

Algorithm MaxInterestFlowLP:

$$\begin{aligned}
& \max_{f_{e,I}, X} X \text{ s.t.} && \text{(Objective)} \\
& \forall e \in E^\dagger : f_e = \sum_{I \in \mathcal{I}} f_{e,I} && \text{(Total Flow)} \\
& X + \sum_{e \in \text{In}(s)} f_e = \sum_{e \in \text{Out}(s)} f_e, \quad \sum_{e \in \text{In}(t)} f_e = X + \sum_{e \in \text{Out}(t)} f_e && \text{(Flow Value)} \\
& \forall v \notin \{s, t\} : \sum_{e \in \text{In}(v)} f_e = \sum_{e \in \text{Out}(v)} f_e && \text{(Flow Conservation)} \\
& \forall v \notin \{s, t\}, \forall I : \sum_{e \in \text{In}(v)} \sum_{I' \leq I} f_{e,I'} \leq \sum_{e \in \text{Out}(v)} \sum_{I' \leq I} f_{e,I'} && \text{(Monotonicity)} \\
& \forall e \in E^\dagger, \tau(e) = \text{credit}, \forall I < r(e) : f_{e,I} = 0 \\
& \forall e \in E^\dagger, \tau(e) = \text{iou}, \forall I \neq r(e) : f_{e,I} = 0 && \text{(Valid Interest Rates)} \\
& \forall e \in Ef, I : 0 \leq f_{e,I}, \forall e : 0 \leq f_e \leq c(e) && \text{(Capacity)}
\end{aligned}$$

First note that if there is no solution to the general (non-restricted) payment routing problem, then clearly there cannot be a solution to the \mathcal{I} -restricted payment routing problem (for any \mathcal{I}). The following lemma states that the other direction holds as well, when taking \mathcal{I} to be the original set of contract interest rates.

LEMMA 2. *Given an instance of the payment routing problem, let $\mathcal{I} = \{r(e) \mid e \in E\}$. There exists a feasible payment for the payment routing problem if and only if there exists a feasible payment for the \mathcal{I} -restricted payment routing problem.*

PROOF. As described above, one direction is immediate, and so we prove that the existence of a solution to the non-restricted payment routing problem implies a solution to the \mathcal{I} -restricted payment routing problem.

Let f be the flow solution to the non-restricted problem, with consistent decomposition P, F , and valid, monotonic realized interest rates $\{\hat{r}^{(k)}\}$. We define a new set of realized interest rates $\{\hat{r}_{\mathcal{I}}^{(k)}\}$ which take values only in \mathcal{I} , but are still valid and monotonic. f is still a flow for the restricted problem, and F, P a consistent decomposition, so this is sufficient to prove the lemma.

To define $\{\hat{r}_{\mathcal{I}}^{(k)}\}$, fix a k , and consider the path $P^{(k)} = (e_1^{(k)}, \dots, e_{n_k}^{(k)})$. For any edge e_i such that $\tau(e_i) = \text{iou}$, we leave the interest rate unchanged. That is, we set $\hat{r}_{\mathcal{I}}^{(k)}(e_i) = \hat{r}^{(k)}(e_i)$ which is also equal to $r(e_i)$ since $\{\hat{r}^{(k)}\}$ are valid.

Next consider edges e_i such that $\tau(e_i) = \text{credit}$. At a high level, we define $\hat{r}_{\mathcal{I}}^{(k)}(e_i)$ by taking $r^{(k)}(e_i)$ and increasing it to the contract rate of the preceding iou edge along the path. In detail, if e_i and e_j , for $i < j$, are consecutive iou edges in $P^{(k)}$, then for all $e_{i'}$, $i < i' < j$ we define $\hat{r}_{\mathcal{I}}^{(k)}(e_{i'})$ by letting $\hat{r}_{\mathcal{I}}^{(k)}(e_{i'}) = r(e_i)$. The validity and monotonicity of $\hat{r}^{(k)}$ tells us that $r(e_i) = \hat{r}^{(k)}(e_i) \geq \hat{r}^{(k)}(e_{i+1}) \geq \dots \geq \hat{r}^{(k)}(e_j) = r^{(k)}(e_j)$. And therefore $r(e_i) = \hat{r}_{\mathcal{I}}^{(k)}(e_i) = \hat{r}_{\mathcal{I}}^{(k)}(e_{i+1}) = \dots = \hat{r}_{\mathcal{I}}^{(k)}(e_{j-1}) \geq \hat{r}_{\mathcal{I}}^{(k)}(e_j) = r(e_j)$. Similarly, if e_i is the last iou edge, we set $\hat{r}_{\mathcal{I}}^{(k)}(e_{i'}) = r(e_i)$ for $i < i' \leq n_k$. Finally, if e_j is the first iou edge, we set $\hat{r}_{\mathcal{I}}^{(k)}(e_{i'}) = \max\{r(e_1), r(e_2), \dots, r(e_j)\}$, and therefore $\hat{r}_{\mathcal{I}}^{(k)}(e_1) = \dots = \hat{r}_{\mathcal{I}}^{(k)}(e_{j-1}) \geq \hat{r}_{\mathcal{I}}^{(k)}(e_j) = r(e_j)$. In each of these cases $\hat{r}_{\mathcal{I}}^{(k)}(\cdot)$ is monotonically nonincreasing over the subsequence in question, and is therefore nonincreasing along the entire path.

The $\{\hat{r}_{\mathcal{I}}^{(k)}\}$ are also valid since the realized interest rate for iou edges are unchanged (compared to $\hat{r}^{(k)}$), and the realized interest rate for credit edges only increase, so the validity of $\{\hat{r}^{(k)}\}$ also ensures that $\hat{r}_{\mathcal{I}}^{(k)}(e_i) \geq \hat{r}^{(k)}(e_i) \geq r(e_i)$ for any edge e_i . Observing that the construction ensures the each $\hat{r}_{\mathcal{I}}^{(k)}$ takes values in \mathcal{I} concludes the proof. \square

Lemma 2 tells us that the \mathcal{I} -restricted problem is equivalent to the non-restricted problem when $\mathcal{I} = \{r(e) \in E\}$. Thus, we can state our main algorithmic result as a corollary of this lemma and Theorem 1.

COROLLARY 1. *Let $\mathcal{I} = \{r(e) \mid e \in I\}$ be the set of contract interest rates. Given an instance (N, s, t, X) of the payment routing problem, the payment is feasible if and only if the the LP solved by Algorithm MaxInterestFlowLP finds a solution of value at least X for the \mathcal{I} -restricted routing problem.*

4.4. Routing Multiple Payments

The preceding demonstrates how to efficiently compute whether a payment is feasible for some static instance (N, s, t, X) of the payment routing problem in an FCN. To route multiple payments in sequence, we update the FCN N after each to reflect the state of obligations between agents. An illustration of such an update is provided in Figure 1a. Here we describe this update formally. In order to do so, we must be explicit about the fact that G contains multi-edges described by 5-tuples (u, v, τ, c, r) .

A feasible payment is given by a flow f in G^{\dagger} , a consistent decomposition $F = \{f^{(1)}, \dots, f^{(K)}\}$, $P = \{P^{(1)}, \dots, P^{(K)}\}$, and realized interest rates $\{\hat{r}^{(1)}, \dots, \hat{r}^{(K)}\}$. Given such a payment $N = (G, \tau, c, r)$ is updated as follows. For each $k \in \{1, \dots, K\}$, consider each edge $e = (u, v) \in P^{(k)}$.

If $\tau(e) = \text{iou}$, the payment was routed through edge e by having v relinquish $f^{(k)}$ IOUs back to u . Thus, there exists some edge (v, u, iou, c, r) in G , for some c and r . We update this edge to $(v, u, \text{iou}, c - f^{(k)}, r)$. At the same time, if there exists a credit line from u to v represented by $(u, v, \text{credit}, c', r')$, we update this edge to $(u, v, \text{credit}, c' + f^{(k)}, r')$, (allowing $c' = 0$ if the credit line between u and v is saturated).

If $\tau(e) = \text{credit}$, the payment was routed through edge e by drawing upon a line of credit that v extends to u . Thus, there exists some edge $(v, u, \text{credit}, c, r)$ in G , for some values of c and r . We update this edge to $(v, u, \text{credit}, c - f^{(k)}, r)$. Drawing upon this credit creates a debt that u owes v at realized interest rate $r^{(k)}(e)$. Thus, if there exists an edge $(u, v, \text{iou}, c', r^{(k)}(e))$, we update this edge to $(u, v, \text{iou}, c' + f^{(k)}, r^{(k)}(e))$, otherwise we create a new iou edge given by $(u, v, f^{(k)}, r^{(k)}(e))$. Note that in both these cases, the newly created IOU is given the realized interest rate for the credit line.

5. PAYMENT MECHANISMS

In general, there may be several feasible ways to route a payment between two agents. In this section, we discuss **payment mechanisms**: rules for choosing flows to achieve a designated payment.

5.1. Choice of Payment Paths

In basic credit networks, the selection of payment paths has no bearing on long-term liquidity. Dandekar et al. [2011] showed that if a sequence of unit flows defined by source/sink pairs $\{s_1, t_1\}, \dots, \{s_k, t_k\}$ is feasible using corresponding payment paths $\{P_1, \dots, P_k\}$ on a network with unit capacities, they remain feasible if P_i , for any i , is changed to an arbitrary feasible payment path P'_i . Given this invariance, and the lack of any differential costs, selecting among feasible payment paths has not been a pressing issue in the basic model without interest rates.




Fig. 2: Dotted arrows represent IOUs, and edges are identified by #num labels. The max flow from s to t is two, but a unit payment on the (red) path (#1,#4) blocks any subsequent flow.

In FCNs, however, interest rates do pose differential costs, and moreover the liquidity invariance property does not hold. Consider the simple example of Figure 2. The max flow between s and t is 2, but if a payment of one unit is first routed along the red path, no further flow is achievable. Payment on the path comprising edges #2 and #3 in the residual FCN would be infeasible due to the interest monotonicity constraint. So in general the success of a sequence of transactions may depend on which feasible payment paths are chosen. As we show below, alternative path selection mechanisms exhibit systematically different long-term liquidity properties.

5.2. Mechanism Definitions

One way to define alternative payment mechanisms is by refining the objective function of MaxInterestFlowLP. Instead of maximizing flow, we specify the requested flow as a constraint and insert criteria for choosing among feasible flows in the objective. For example, we could choose based on length of paths, or some function of the interest rates on the included paths.

The monotonicity constraint ensures that intermediate nodes accrue nonnegative net interest, but alternative paths may differ on the amount of positive net interest. The payment source generally pays positive interest. We term the interest associated with the first edge on a payment path the *originating rate*. To the extent the transaction initiator has influence over paths chosen, it may be particularly concerned with minimizing this rate.

For mechanisms below, fix \mathcal{I} to be the set of initial contract interest rates, and let $I^+ \triangleq \max(\mathcal{I})$ be the maximum possible interest rate in the network. Let s be the source, and $\text{Out}(s)$ the outgoing edges from s in the reverse network G^\dagger , as defined in §3.1. The mechanisms are defined by replacing the objective function of MaxInterestFlowLP with those exhibited, fixing the flow to the requested amount, and in one instance adding additional constraints. The optimization variables and remaining constraints in MaxInterestFlowLP are unchanged.

$$\min \sum_{e \in \text{Out}(s), I \in \mathcal{I}} f_{e,I} \times I \quad \text{Minimize source cost (MinSrc)}$$

The MinSrc mechanism minimizes the originating rate. We place a cost equal to the interest rate that would be paid on all potential outgoing flows from the source.

$$\min \sum_{e \in \text{Out}(s), I \in \mathcal{I}} f_{e,I} \times (I + (I^+ + 1) \times \mathbb{1}\{\tau(e) = \text{credit}\}) \quad \text{Minimize credit usage (MinCred)}$$

The MinCred mechanism prioritizes use of IOUs relative to credit lines. The idea is that since credit edges have flexible realized interest rates, their capacity should be preserved for future situations. The originating rate is minimized secondarily. The flexibility provided by a credit line could enable additional payment paths by raising the originating interest rate when necessary. The $I^+ + 1$ term is simply a way to ensure that IOU status has priority in the objective over originating rate.

$$\begin{aligned} \min \sum_{e \in \text{Out}(s), I \in \mathcal{I}} f_{e,I} \times I & \quad \text{Restrict to contract rates (NoMarkup)} \\ \forall e, I > r(e) : f_{e,I} = 0 & \end{aligned}$$

For flow $f_{e,I}$ on a credit line, we refer to the difference $I - r(e)$ as the **interest markup**. The markup is additional interest above the contract rate assessed on credit edges in order to make a payment path monotone. The NoMarkup mechanism minimizes originating rate, subject to constraints disallowing any agent from drawing on credit lines at anything other than their contract rates. NoMarkup is unique among the mechanisms we consider in that it imposes additional constraints beyond flow feasibility. This allows us to evaluate the effect of flexible markup policies (which all our other mechanisms allow) on liquidity.

$$\max \sum_{e \in E, I \in \mathcal{I}} f_{e,I} \times \mathbb{1}\{I = I^+\} \quad \text{Maximize total interest paid (MaxIR)}$$

Given an FCN initially consisting of credit edges, liquidity is maximized by always assessing the highest interest rate possible. In fact, this MaxIR mechanism provides liquidity equivalent to that of basic credit networks (Theorem 6.2). Note that MaxIR is sensitive to the maximum interest rate I^+ , and thus is not very robust. We can peg the interest paid to an arbitrarily high rate r^* by introducing a single credit edge e' with $r(e') = r^*$, regardless of how much credit exists at reasonable rates. We include MaxIR to provide an upper bound on liquidity and demonstrate the tradeoff between liquidity and interest rates.

6. STEADY-STATE LIQUIDITY ANALYSIS

With respect to a given transaction, any payment mechanism that does not restrict feasible payments (i.e., all those listed above except NoMarkup) offers the same prospects for success. As illustrated in §5.1, however, how a payment is routed can affect the network's configuration, which changes the prospects for subsequent transactions. Alternative payment mechanisms may affect network configurations systematically, and thus have a qualitative impact on liquidity. We measure liquidity by the long-term failure rate of transactions once the network has reached a steady-state distribution over network configurations. We call this the **steady-state transaction failure rate**. We show in §6.2 that given our experimental setup, such a steady state exists.

6.1. Experimental Setup

Our experimental setup is designed as an extension of the no-interest liquidity analysis of Dandekar et al. [2011]. We initialize an Erdos-Renyi graph [Erdős and Rényi 1959] with a chosen average degree between 5 and 35, and size of 200 nodes. Each directed edge represents a credit line with initial capacity 10. Edge orientation is a fair coin toss and interest rates are generated uniformly from a specified set \mathcal{I} . Namely, if we choose to have four interest rates, each edge is assigned contract rates of 0.01, 0.02, 0.03, or 0.04 uniformly, and so on for different $|\mathcal{I}|$. We generate transactions by selecting source and destination nodes uniformly at random and attempting to route 10 units of flow between them.

For each mechanism, we attempt to route transactions sequentially while recording any failures. Following a failure, we do nothing to the graph. After 9000 transaction attempts, we check the failure rate among the first 4500 against the second 4500 transactions. If the two failure rates are within 0.002 of each other, we stop and record the

failure rate of the entire history as the steady-state failure rate. Otherwise, we generate another transaction and move our observation window forward by one, that is, we compare observations 2 through 4501 with 4502 through 9001. This continues until we satisfy our steady-state criterion. We perform the whole process ten times, averaging the results.

Note that for purposes of this liquidity analysis, we do not consider actual interest payments. Interest plays a role here only for transaction feasibility and selection of flows by payment mechanisms.

6.2. Steady-State Liquidity of FCNs

A network configuration, or **state**, is a single specification of the tuple $\{G, \tau, c, r\}$. In our simulations, the initial state is set at $\tau(e) = \text{credit}$, c maps to a constant integer, r maps to an element of a set of constant size, and G is an Erdos-Renyi graph, for all $e \in E$. The number of reachable states is in general unbounded, since a flow on an edge may take on any value between 0 and $c(e)$, leaving any real value as a possible residual capacity. However, if $|\mathcal{I}|$ is finite, payment amounts are integers, and all flows are restricted to be integral, then the number of states is finite. In all the simulations reported here, we route constant integer payment amounts under bounded $|\mathcal{I}|$. As the linear program solver software we used (CPLEX) maintains integer solutions, the set of states is finite and discrete [ILOG 2007].

The generation of random transactions induces a **transition probability matrix** between states. The resulting stochastic process is an ergodic Markov chain and therefore has a steady-state distribution.

THEOREM 6.1. *Let $h, h' \in T$ where T is the finite state space over states of the FCN, when routing integer valued flows of amount X . Let $P(h, h')$ be the probability of moving from state h to state h' . These probabilities are induced by the transaction probability matrix Λ which picks any source and sink node pair (s, t) with positive probability, between which a payment of amount X is attempted. This stochastic process χ is an ergodic Markov chain.*

PROOF. First note that χ satisfies the Markov property. Given we are in state h , $P(h, h')$ is independent of any other future or historical state for any state h' . To calculate this probability, we can just sum the probabilities of every transaction that will get us from h to h' .

To show that χ is ergodic, first we need that h is reachable from h' with nonzero probability (irreducibility), for any $h, h' \in T$. We use the fact that any pair of nodes can be selected for a transaction with nonzero probability, since we choose the source and sink uniformly at random. So any series of pairwise transactions also has nonzero probability. In particular, neighbors in the graph can be selected with positive probability. Given h , we can always find a series of pairwise transactions between neighbors to reach h' . The reason for this is that any one-hop path can be routed ignoring the monotonicity constraint, since the source is willing to pay any originating rate. So first we can route flows on a sequence of one-hop paths that cancels all existing IOUs and returns to G to the all-credit edge graph. We can then route a sequence of transactions that result in the exact configuration of h' . So χ is *irreducible*.

A version of the ergodicity theorem says that for a finite-state irreducible Markov chain, we need only one state to be aperiodic in order for all states to be aperiodic [Serfozo 2009]. Consider a state h , with two agents s and t such that a payment between s and t is not feasible. Since there is a non-zero probability of picking s and t for the next transaction, there is a non-zero probability that h transitions to itself. Therefore, the state h is aperiodic, which implies that the Markov chain is aperiodic. Such a state

is always reachable by considering a sequence of states that fill all feasible payment paths from s to t to capacity. \square

This means that after routing enough transactions, the overall probability of being in a particular state is invariant. Since failure probability is a function of network state, this probability is also invariant. Following Dandekar et al. [2011], we use this steady-state failure probability as our measure of *liquidity*. We expect the procedure of §6.1 to yield the correct liquidity if our error tolerance is appropriately chosen. Thus, we can evaluate the liquidity when using each of our payment mechanisms. Notably, the MaxIR mechanism maximizes liquidity.

THEOREM 6.2. *Initialize an FCN $G = (V, E)$ such that $\forall e \in E, \tau(e) = \text{credit}$. If mechanism MaxIR is used for every payment in G with contract rates \mathcal{I} , each transaction is infeasible iff it is infeasible under a basic (interest-free) credit network with the same graph structure.*

PROOF. If a transaction is infeasible in the interest-free network, then adding additional constraints by including interest rates will not cause it to succeed. In the other direction, no matter how many transactions are routed on G under \mathcal{I} , there will never exist an interest rate on an IOU that is not I^+ . That is, $\forall e \in E. \tau(e) = \text{IOU} \implies r(e) = I^+$, by definition of the mechanism. Furthermore, the mechanism always opts for the maximum markup on every credit edge. Thus, realized interest rates are constant, at the value I^+ , and monotonicity is trivially satisfied. If a payment of amount X is not feasible it must be because there was no $s - t$ flow of amount X , which implies that the payment is not feasible in the interest-free network. \square

6.3. Markup Flexibility on Credit Lines

Figure 3 indicates that steady-state failure rate converges and becomes small at around degree 25 regardless of how many interest rates are available and regardless of which payment mechanism we use. The number of available interest rates is much more influential when the NoMarkup mechanism is used. Especially striking is how much allowing markups increased liquidity at lower degrees. The liquidity of MinSrc is almost as good as the basic credit network case (i.e., one interest rate) at degrees as low as 10. $|\mathcal{I}|$ has minimal effect on liquidity. For comparison, NoMarkup has almost three times the failure rate of the basic credit network at degree 10. This suggests that accounting for interest has modest effect on liquidity as long as agents are allowed to increase rates on credit lines.

6.4. Tradeoff Between High Markups and Better Liquidity

We already know that FCN liquidity is maximized by forcing every agent to pay the maximum possible interest rate on every edge. The relationship between higher interest rates and improved liquidity is also exhibited more generally. In Figure 4b, we see that compared to MinSrc, the MinCred mechanism yields a slightly higher average interest rate on IOUs. Correspondingly, it has a slightly lower steady-state failure rate (Figure 4a). Many other heuristic mechanisms were tested and the relation between interest and failure rates held true. For brevity, we omit detailed descriptions, and show only average interest and failure rates (Figure 5). Conditional on average degree of the graph, almost all variation in liquidity among mechanisms can be explained by the average interest rate paid by agents when the mechanism is employed. The tradeoff is clear: better liquidity comes with higher interest rates. This tradeoff diminishes when the graph is well-connected, as almost all transactions succeed.

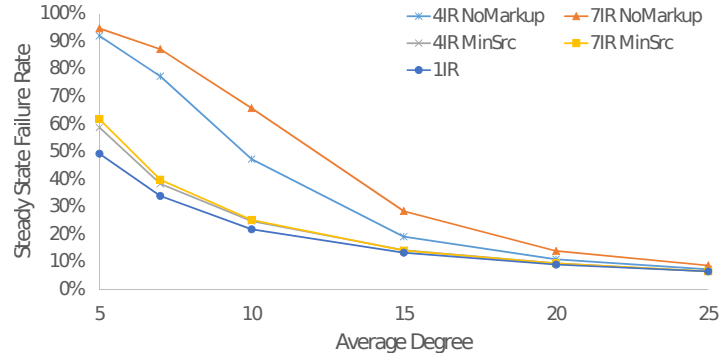


Fig. 3: Steady-state failure probabilities for NoMarkup and MinSrc mechanisms at degrees 5 to 25, Erdos-Renyi graphs with 200 nodes. “xIR” means $|\mathcal{I}| = x$.

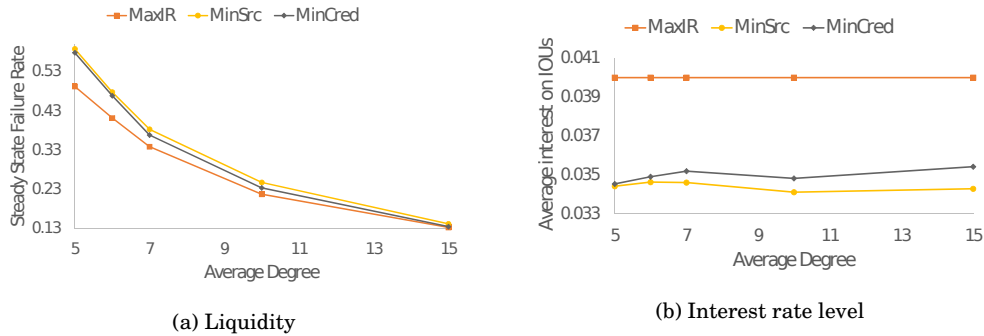


Fig. 4: Failure rate and average interest rate level over all IOUs in the final graph, obtained after reaching steady state. Starting from degree 7, difference in failure rate between MinSrc and MinCred is statistically significant.

The explanation for this is straightforward. Given the monotonicity constraint, an IOU at low interest restricts the completing paths. Thus, higher interest on IOUs promotes greater transaction capacity for the network.

7. STRATEGIC ROUTING GAME

The relationship between interest rates and liquidity presents a strategic dilemma for choice among mechanisms. A low originating interest rate minimizes costs for the payer, but imposes an externality on the rest of the network in the form of reduced liquidity. We explore the conflict between individual incentives and global effectiveness in FCNs by defining a *game*, where the source of each transaction chooses a payment mechanism. We evaluate the game by simulation, over a range of network settings and payment mechanisms similar to those explored in the liquidity study above.

Our scenario employs 100 agents initialized to a random-graph FCN, who attempt to execute a randomly generated sequence of transactions. Specifically, the steps of each simulation run are: (1) Assign strategies to agents according to a specified *strategy profile*. Strategies in this context are simply payment mechanisms as defined in §5.2. (2) Generate a directed Erdos-Renyi graph with 100 nodes and a specified average degree. Assign agents to these nodes. Set the capacity on each edge to 100, and the

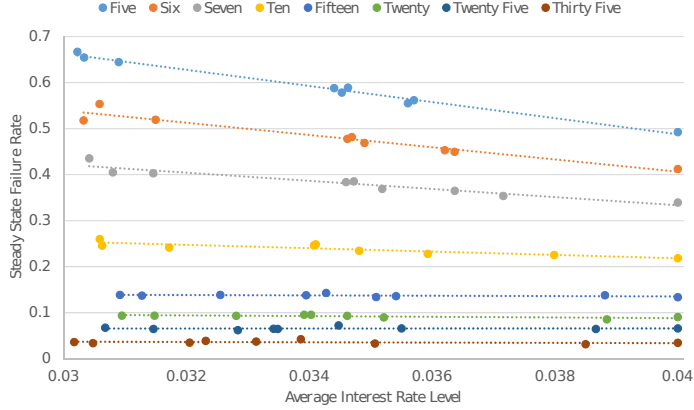


Fig. 5: Each point represents a heuristic payment mechanism’s steady state failure rate at a certain average interest rate level on Erdos-Renyi graphs of fixed average degree. Each dotted line represents a collection of mechanisms whose liquidity was tested at a fixed average degree.

contract interest uniformly at random from the set $\{0.01, 0.02, 0.03, 0.04\}$. Initialize the payoff for each agent to zero. (3) Choose a source and sink uniformly at random. (4) Attempt to route 100 units from source to sink using the source’s fixed strategy. If successful, update the FCN based on the chosen flows, and increment the source’s payoff by the amount $xVal$. (5) Repeat steps (3) and (4) 2000 times. (6) Calculate net interest income for each node, and add this to the corresponding payoff.

The interest income is defined under the assumption that IOUs resulting from the transaction sequence remain outstanding for one period. Let $IN_x = \{e_{uv} \in E : v = x, \tau(e_{uv}) = iou\}$ denote the set of incoming IOUs of node x , and similarly $OUT_x = \{e_{uv} \in E : u = x, \tau(e_{uv}) = iou\}$ the outgoing IOUs of x . Node x ’s net interest income is then $\sum_{e \in IN_x} c(e) \times r(e) - \sum_{e \in OUT_x} c(e) \times r(e)$. An agent’s overall payoff is the cumulative value from successful transactions plus net interest income.

In our analysis the only scenario parameter we vary is the average degree of the initial random graph.³ We explored three settings: 8, 15, and 22.

Since promoting liquidity comes with positive externalities, we would not expect to see social welfare maximized in equilibrium. Our hypothesis was that agents would choose to pay a lower originating rate than is socially optimal. To evaluate this hypothesis, we performed simulations over strategy profiles combining three mechanisms: MinSrc, MinCred, and MaxIR. These represent three points on the spectrum between minimizing interest cost (MinSrc) and maximizing liquidity (MaxIR).

To evaluate a profile, we average over at least 1500 simulation runs, to produce accurate payoff estimates. This takes 50-100 core-hours, depending on average node degree. Even with only three strategies, exhaustive simulation of profiles for 100 players is not feasible. Exploiting symmetry, there are 5151 profiles in the full game, which would take too long to cover with available resources. We therefore employed *deviation preserving reduction* [Wiedenbeck and Wellman 2012] to approximate the 100-player game with a reduced 4-player version. This required simulation of only 30 full-game (100-agent) profiles to estimate each game model.

Among the symmetric pure-strategy profiles (i.e., where all players choose the same strategy) shown in Figure 6, the profile with the highest social welfare (i.e., sum of

³We explored two $xVal$ settings, but observed no interesting differences, so report results only for $xVal = 10$.

Average Degree	Pure Strategy Profile Payoffs			PSNE	Dominated Strategies
	MaxIR	MinCred	MinSrc		
8	145	139	132	MinSrc, MinCred	MaxIR
15	180	178	177	MinCred	MaxIR
22	193	192	192	MinCred	MaxIR, MinSrc

Fig. 6: Payoff and symmetric equilibria information for varying average degree.

all payoffs) in every setting is that where all players play MaxIR. The MinCred profile had lower social welfare, followed by MinSrc. MaxIR was dominated at all settings. We searched for symmetric Nash equilibria using replicator dynamics [Schuster and Sigmund 1983]. We found only pure-strategy Nash equilibria (PSNE) in our experiments, reflecting a coordination benefit for adopting a uniform mechanism.

MinSrc and MinCred exhibited similar payoffs in most profiles containing both strategies. At low average degree (8), we found PSNE consisting of both strategies, but at high average degree (15, 22) only MinCred is in equilibrium. At average degree 22, the Mincred PSNE was confirmed statistically using the bootstrapped regret methodology outlined by Wiedenbeck et al. [2014]. No other games exhibited solutions with zero regret at the 95th percentile.

Though the socially optimal outcome (i.e., all agents playing MaxIR) is strategically unstable, the fact that MinCred is competitive with MinSrc is encouraging. From Figure 4 we know that MinCred trades off higher interest rate for better liquidity compared to MinSrc. This suggests that agents have some willingness to pay higher interest, in conjunction with the coordination benefit, to achieve higher liquidity.

8. CONCLUSION

We have extended the credit-network model of distributed trust with interest rates and a distinction between credit lines and debt. *Financial credit networks* provide a standard means to incentivize extending credit, and thereby also support direct modeling of real-world financial relationships. The extension to support interest rates raises new issues in routing payments over the network, in the form of constraints to assure the willingness of intermediaries to participate. These are further complicated by a distinction between credit lines and debt, namely that the former may admit flexibility in rates whereas the latter are more rigid. We formalize these constraints, and develop an efficient algorithm for determining a feasible payment flow, as well as a range of mechanisms for choosing among feasible flows.

Given the plethora of routing policies and options, we perform computational studies to explore the implications of alternative mechanisms on network liquidity. We find that performance can vary greatly across mechanisms, and moreover that there may be a strategic tension between preferences of the transaction source and global network effectiveness. We explore this issue through empirical game-theoretic analysis, and find that while this tension does exist, there is evidence that agents will not simply maximize myopic gain and instead consider overall network liquidity in equilibrium.

One direction for further work is to grant more discretion to intermediate agents on payment paths. Monotonicity ensures that no such agent loses net interest, but they could go further in choosing how to maximize payoff even when they are not the source. Another direction is to explore how agents in a setting where IOUs yield interest will actually issue credit, and what kinds of networks result. This is a core question underlying the effectiveness and stability of our global financial system. Theories about financial network dynamics should arguably be tested within a framework for mod-

eling the financial sector at high fidelity. By providing a grounded representation of financial obligations, FCNs can provide a foundation for such modeling.

ACKNOWLEDGMENTS

We thank the anonymous reviewers (of this submission and a prior version) for numerous suggestions on substance and presentation, including proposing the possibility of interest rate markups.

REFERENCES

- Dandekar, P., Goel, A., Govindan, R., and Post, I. (2011). Liquidity in credit networks: A little trust goes a long way. In *12th ACM Conference on Electronic Commerce*, pages 147–156.
- Dandekar, P., Goel, A., Wellman, M. P., and Wiedenbeck, B. (2015). Strategic formation of credit networks. *ACM Transactions on Internet Technology*, 15(1):3:1–3:41.
- DeFigueiredo, D. B. and Barr, E. T. (2005). TrustDavis: A non-exploitable online reputation system. In *7th IEEE Int'l Conf. on E-Commerce Technology*, pages 274–283.
- Erdős, P. and Rényi, A. (1959). On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297.
- Even, S., Itai, A., and Shamir, A. (1975). On the complexity of time table and multi-commodity flow problems. In *16th Annual Symposium on Foundations of Computer Science*, pages 184–193.
- Ghosh, A., Mahdian, M., Reeves, D. M., Pennock, D. M., and Fugger, R. (2007). Mechanism design on trust networks. In *Third International Workshop on Internet and Network Economics*, pages 257–268.
- ILOG (2007). CPLEX 11.0 user's manual. Technical report, ILOG SA, Gentilly, France.
- Karlan, D., Mobius, M., Rosenblat, T., and Szeidl, A. (2009). Trust and social collateral. *Quarterly Journal of Economics*, 124(3):1307–1361.
- Mislove, A., Post, A., Druschel, P., and Gummadi, K. P. (2008). Ostra: Leveraging trust to thwart unwanted communication. In *Fifth Usenix Symposium on Networked Systems Design and Implementation*, pages 15–30.
- Moreno-Sanchez, P., Kate, A., Maffei, M., and Pecina, K. (2015). Privacy preserving payments in credit networks. In *Network and Distributed Security Symposium*.
- Resnick, P. and Zeckhauser, R. (2002). Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In Baye, M. R., editor, *The Economics of the Internet and E-Commerce*. Elsevier.
- Schuster, P. and Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538.
- Serfozo, R. (2009). *Basics of Applied Stochastic Processes*. Probability and Its Applications. Springer Berlin Heidelberg.
- Wiedenbeck, B., Cassell, B.-A., and Wellman, M. P. (2014). Bootstrap statistics for empirical games. In *13th International Conference on Autonomous Agents and Multiagent Systems*, pages 597–604.
- Wiedenbeck, B. and Wellman, M. P. (2012). Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 931–938.
- Yu, Y. and Shen, M. (2016). Consumer protection as the 'open sesame' that allows Alibaba to crush the forty thieves. *Journal of Antitrust Enforcement*.