Informal Transfers in Social Networks

Markus Mobius
Microsoft Research, University of Michigan and NBER

Tanya Rosenblat
University of Michigan

December 24, 2015

Abstract

Social networks can facilitate informal lending and risk-sharing in situations where formal institutions such as banks and insurance companies do not exist. The social collateral approach provides an analytically tractable framework that can be used to analyze a wide range of informal transfers. Moreover, the approach is easily amenable to empirical analysis.

JEL Classification: C91, C93, D83

1 Introduction

Many fundamental economic institutions facilitate transfers across time. Banks take deposits and repackage them as loans. Insurance companies mitigate idiosynchratic risk across a large pool of agents. These institutions can only function if they are able to control moral hazard: debtors have to repay their loans and victims of adverse shocks have to be able to rely on their insurance policy to make transfers. Developed economies usually rely on the legal system to protect lenders and policy holders.

However, developing countries often lack a reliable legal system which raises the transaction costs of providing loans or insurance. Even developed economies often require lenders to own physical collateral that can be used to secure loans and thus lower transaction costs. Certain groups of borrowers, such as entrepreneurs starting a business, frequently lack physical collateral. In these situations, *informal lending* within close-knit communities can substitute for formal transfers even in the absence of physical collateral. But what mechanisms make such informal arrangements work and control for moral hazard? How can a lender trust that an informal loan will be repaid in the future? How much assistance can a farmer with a bad harvest realistically expect from his extended family and friends? Social networks provide a natural structure to study these phenomena: intuitively we expect that households receive greater assistance from socially close neighbors with whom they share stronger ties. Moreover, agents in a social networks play a repeated supergame with their friends and neighbors that can help support informal lending and risk-sharing arrangements.

In this Chapter we describe the *social collateral* approach introduced by **?** and **?** as a simple way to model informal transfers within social networks. This approach views social links as "collateral" that can be used to control moral hazard: if a borrower does not repay a loan (in case of borrowing) or if a an agent refuses to help a neighbor in need (in case of risk-sharing), they risk losing the social link and its associated benefits. In this class of models, the function of social capital (which is exactly the social network) is analogous to the role function of physical collateral in formal transfer arrangements.

The social collateral approach makes two key simplifying assumptions that keep the model analytically tractable and empirically relevant. First of all, the entire network supergame is collapsed into a two-period game: borrowing or risk-sharing arrangements are implemented in period 1 and out-of-equilibrium punishments occur in period 2. Second, there is no explicit group punishment (such as ostracism) – not repaying a loan, for example, only jeopardizes the direct link between the borrower and the lender (or intermediary). However, it can be shown that group punishments that are robust to *coalitional deviations* can reduce the set of implementable outcomes (under certain conditions) to the same set that is implementable through direct punishment schemes.

2 Empirical Facts

Before discussing the social collateral approach in detail, it is useful to summarize some of the main findings of the empirical literature on informal transfers.

First of all, informal transfers are remarkably successful in mitigating income risk on the village level. In particular, **?** finds that the full insurance model where agents all agents pool their income provides a good benchmark for understanding consumption in Indian villages.[1]

Second, transfers are highly *local* and tend to occur between socially close households. **?** find that transfers take place primarily through networks of friends and relatives rather than at the village level. Similarly, **?** collect detailed data on insurance networks within a single village in Tanzania. They find that networks are not clustered but largely overlapping. They also confirm the effectiveness of these insurance networks in mitigating income risk: they cannot reject full insurance at village level for food insurance, and find partial insurance of non-food consumption in the networks. **?** exploit a natural experiment where the Mexican government made cash transfers to subset of households in Mexican villages. They find that receiving households share these transfers with relatives in the village (but not with non-relatives). This paper contributes by using exogenous income variation to examine the risk sharing network and extent of full insurance. **?** use a field experiment with 70 Columbian villages show that pairs of participants who are friends or relatives are more likely to form risk-sharing groups which suggests that social obligations to neighbors in the social network are an important determinant of informal transfers.

While social proximity is typically highly correlated with geographic distance remittances between relatives who live far apart are an important exception because they can help mitigate local shocks. **?** finds evidence that insurance is an important motive in remittances. **?** show that remittances respond to incomes shocks which is consistent with an insurance motivation. They cannot reject the hypothesis that households with a migrant worker are fully insured against income shocks, and can strongly reject for households without a migrant worker. Technological advances have reduced the transaction costs of remittances. **?** show that M-pesa cash transfers through mobile phones insulate households from shocks thanks to an increase in remittances from a diverse pool of senders.

Third, informal transfers often take the form of loans. **?** was one of the first papers to demonstrate that informal credit contracts play a direct role in pooling risk, as repayments owed by borrowers depend on the random shocks faced by both the borrowers and lenders. **?** use detailed data on gifts, loans, and asset sales to examine which methods are used to cope

[1] However, some papers have also documented substantial deviations from the full risk-sharing benchmark. **?** find little evidence of consumption smoothing during a period of severe drought in Burkina Faso. **?** find that informal insurance against several illness is very imperfect. **?** show that risk-sharing does not always occur even *within* households as women in rural Ethiopia bear the brunt of adverse shocks. **?** also reject full risk-sharing within families using PSID data.

Figure 1: Informal borrowing in some sample social networks

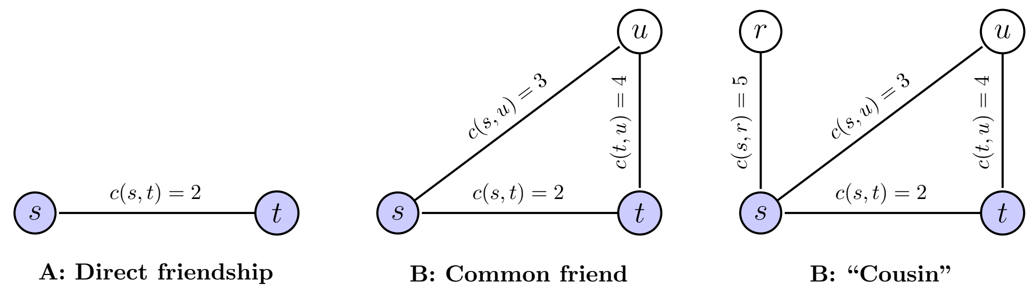

with income shocks. They find that gifts are loans are the most common mechanisms.

3 Informal Lending and Trust

We consider a situation where a borrower needs the asset of a lender to produce social surplus.[2] In the absence of legal contract enforcement, borrowing must be secured by an informal arrangement supported by the social network: connections in the network have associated consumption value, which serve as *social collateral* to enable borrowing.

3.1 Motivating Example

To understand the basic logic of the model, consider the examples in Figure **??**, where agent s would like to borrow an asset, like a car, from agent t, in an economy with no formal contract enforcement. In Figure **??**A, the network consists only of s and t; the value of their relationship, which represents either the social benefits of friendship or the present value of future transactions, is assumed to be 2. As in standard models of informal contracting, t will only lend the car if its value does not exceed the relationship value of 2.

More interesting is Figure **??**B, where s and t have a common friend u, the value of the friendship between s and u is 3, and that between u and t is 4. Here, the common friend increases the borrowing limit by $\min[3,4]=3$, the weakest link on the path connecting borrower and lender through u, to a total of 5. The logic is that the intermediate agent u vouches for the borrower, acting as a guarantor of the loan transaction. If the borrower chooses not to return the car, he is breaking his promise of repayment to u, and therefore loses u's friendship. Since the value of this friendship is 3, it can be used as collateral for a payment of up to 3. For the lender t to receive this amount, u must prefer transmitting the payment to

[2]This asset might represent a factor of production, such as a farming tool, a vehicle or an animal; it could also be an apartment, a household durable good or simply a cash payment.

losing the friendship with him, explaining the role of the weakest link.

Finally, Figure **??**C illustrates the limits of ostracism under a coalitional refinement. Assume that the lender also has a "cousin" r with whom he has a relationship valued at 5. If the cousin could also act as a guarantor for s then the borrowing limit might increase by an additional 5 to a total of 10. However, the cousin's threat to break off her relationship with the borrower is not credible: for any loan amount exceeding 5 the borrower could propose a "side-deal" to intermediary u and her cousin such that u can reimburse the lender for the her guaranteed amount (which is at most 3) while transferring 0 to the cousin in case of default. The cousin and intermediary are not worse off as a result of this side-deal. The borrower will use her friendship and therefore incur a combined loss of at most 5 - but since she borrowed an amount exceeding 5 she is strictly better off under such a side-deal. Hence, group-level punishment of the borrower that involves agents that are unconnected to the lender (such as the cousin in panel C) is not credible under coalitional refinement.

3.2 Model

Formally, a social network $G = (W, E)$ consists of a set W of agents (vertices or nodes) and a set E of edges (links), where an edge is an unordered pair of distinct vertices. Each link in the network represents a friendship or business relationship between the two parties involved. We formalize the strength of relationships using an exogenously given capacity $c(u, v)$.

Definition 1. *A capacity is a function* $c : W \times W \to \mathbb{R}$ *such that* $c(u, v) > 0$ *if* $(u, v) \in E$ *and* $c(u, v) = 0$ *otherwise.*

The capacity measures the utility benefits that agents derive from their relationships. For ease of presentation, we assume that the strength of relationships is symmetric, so that $c\,(u, v) = c\,(v, u)$ for all u and v.

The model consists of five stages.

Stage 1: Realization of needs. Two agents s and t are randomly selected from the social network. Agent t, the lender, has an asset that agent s, the borrower, desires. The lender values the asset at V, and it is assumed that V is drawn from some prior distribution F over $[0, \infty)$. The identity of the borrower and the lender as well as the value of V are publicly observed by all players.

Stage 2: Borrowing arrangement. At this stage, the borrower publicly proposes a *transfer arrangement* to all agents in the social network. The role of this arrangement is to punish the borrower and compensate the lender in the event of default. A transfer arrangement consists of a set of transfer payments $h\,(u, v)$ for all u and v agents involved in the arrangement. Here $h\,(u, v)$ is the amount u promises to pay v if the borrower fails to return the asset to the lender. Once the borrower has announced the arrangement, all agents involved have the opportunity to accept or decline. If all involved agents accept, then the asset is borrowed and the borrower earns an income $\omega(V)$, where ω is a non-decreasing function with $\omega(0) = 0$. If some agents decline, then the asset is not lent, and the game moves on directly to stage 5.

Stage 3: Repayment. Once the borrower has made use of the asset, he can either return it to the lender or steal it and sell it for a price of V. If the borrower returns the asset then the game moves to the final stage 5.

Stage 4: Transfer payments. All agents observe whether the asset was returned in the previous stage. If the borrower did not return the asset, then the transfer arrangement is activated. Each agent has a binary choice: either he makes the promised payment $h(u, v)$ in full or he pays nothing. If some agent u fails to make a prescribed transfer $h(u, v)$ to v, then he loses his friendship with agent v (i.e., the (u, v) link "goes bad"). If (u, v) link is lost, then the associated capacity is set to zero for the remainder of the game. We let $\widetilde{c}(u, v)$ denote the new link capacities after these changes.

Stage 5: Friendship utility. At this stage, agents derive utility from their remaining friends. The total utility enjoyed by an agent u from his remaining friends is simply the sum of the values of all remaining relationships, i.e., $\sum_v \widetilde{c}(u, v)$.

3.3 Analysis

We are interested in characterizing the *maximum* amount $T^{st}(c)$ that agent s can borrow from lender t for a given social network that is characterized by the capacity function c. We will refer to $T^{st}(c)$ as the *borrowing limit*.

The model is a multi-stage game with observed actions. We focus on the set of pure strategy subgame perfect equilibria. In order to rule out non-credible equilibria (as shown in Figure **??**C) we require that all equilibria are "side-deal proof".

Consider the subgame starting in stage 2, after the identities of the borrower and the lender and the value of the asset are realized, and for any pure strategy σ, let $U_u(\sigma)$ denote the total utility of agent u in this subgame. We formalize the idea of a side-deal as an alternative transfer arrangement $\widetilde{h}(u, v)$ that s proposes to a subset of agents $S \subset W$ after the original arrangement is accepted. If this side-deal is accepted, agents in S are expected to make transfer payments according to $\widetilde{h}$, while agents outside S continue to make payments described by h. In order for the side-deal to be credible to all participating agents, it must be accompanied by a proposed path of play that these agents find optimal to follow. This motivates the following definition.

Definition 2. *A side-deal with respect to a pure strategy profile σ is a set of agents S, a transfer arrangement $\widetilde{h}(u, v)$ for all $u, v \in S$, and a set of continuation strategies $\{\widetilde{\sigma}_u | u \in S\}$ proposed by s to agents in S at the end of stage 2, such that*

(i) $U_u\left(\widetilde{\sigma}_u, \widetilde{\sigma}_{S\backslash u}, \sigma_{-S}\right) \geq U_u\left(\sigma'_u, \widetilde{\sigma}_{S\backslash u}, \sigma_{-S}\right)$ *for all* σ'_u *and all* $u \in S$,
(ii) $U_u\left(\widetilde{\sigma}_S, \sigma_{-S}\right) \geq U_u\left(\sigma_S, \sigma_{-S}\right)$ *for all* $u \in S$,
(iii) $U_s\left(\widetilde{\sigma}_S, \sigma_{-S}\right) > U_s\left(\sigma_S, \sigma_{-S}\right)$.

Condition (i) says that all agents u involved in the side-deal are best-responding on the new path of play, i.e., that the proposed path of play is an equilibrium for all agents in S conditional on others playing their original strategies σ_{-S}. Condition (ii) says that if any agent $u \in S$ refuses to participate in the side-deal, then play reverts to the original path of play given by σ. Finally, (iii) ensures that the borrower s strictly benefits from the side-deal.

Definition 3. *A pure strategy profile σ is a side-deal proof equilibrium if it is a subgame perfect equilibrium that admits no side deals.*

We are now almost ready to state the main result in **?**. The final pre-requisite is a definition of *maximum flow* which is a well known concept in optimization theory and computer science **?**.

Definition 4. *An $s \to t$ flow with respect to capacity c is a function $f : G \times G \to \mathbb{R}$ that satisfies*
(i) Skew symmetry: $f(u, v) = -f(v, u)$.
(ii) Capacity constraints: $f(u, v) \leq c(u, v)$.
(iii) Flow conservation: $\sum_w f(u, w) = 0$ *unless* $u = s$ *or* $u = t$.

The value of a flow is the amount that "leaves" the borrower s, given by $|f| = \sum_w f(s, w)$. Let $T^{st}(c)$ denote the maximum value among all $s \to t$ flows.

The maximum flow captures the intuitive notion of "sum of weakest links" for all distinct paths that connect a borrower and a lender. For example the maximum flow between borrower and lenders in Figure **??** is equal to 2, 5 and 5 in panels A to C. It turns out that the maximum flow exactly characterizes the borrowing limit.

Theorem 1. *There exists a side-deal proof equilibrium that implements borrowing between s and t if and only if the asset value V satisfies*

$$V \leq T^{st}(c). \tag{1}$$

3.4 Empirical Application

The social collateral model can be easily applied to empirical applications where the social network is known. In particular, the maximum flow can be efficiently calculated using the Ford-Fulkerson algorithm (**?**). **?** report on one empirical application that uses data from two Peruvian shantytowns. In 2005, the authors collected social network data on 299 households.

Figure 2: Propensity to borrow as a function of direct and indirect flow

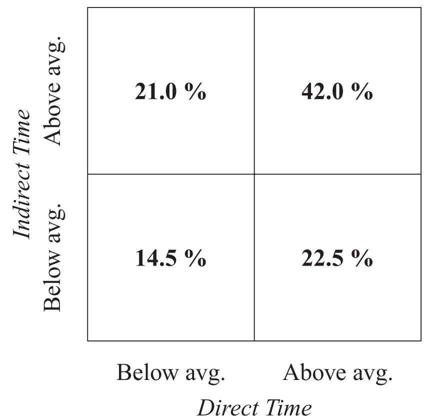

In particular, the authors recorded, for each link, how much time the subjects spends on average with the friend or acquaintance per week and whether the subject ever borrowed money from each social link.

The amount of time provides a convenient proxy for the strength of a relationship. In the data, the distribution of time spent together is skewed: the average responder spends fewer than 6 minutes with the bottom 10 percent of his/her friends and more than 3 hours with the top 10 percent. To obtain a more homogenous measure, the authors define *normalized time* for two connected agents u and v as the value, for the amount of time they spend together, of the empirical cumulative distribution function of time spent together in their community. With this definition, the empirical distribution of normalized time $\tau\left(u,v\right)$ across all connected pairs is a discretized uniform distribution on the unit interval in each community.

The authors also assume that link capacities are created by an increasing production function g such that $c\left(u,v\right) = c\cdot\tau\left(u,v\right)$, i.e., spending more time together results in stronger links. They also restrict attention to the subgraph that includes all direct links of s and t (hence borrowing arrangement can only involve common friends). This allows for a simple decomposition of the trust flow between s and t as

$$T^{st}\left(c\right) = c\cdot\tau(s,t) + c\cdot\sum_{v\in N_s\cap N_t}\min(\tau(s,v),\tau(v,t)), \tag{2}$$

where the first term represents the *direct flow* and the second term is the *indirect flow*. Here N_s is the set of direct friends of agent s.

Table **??** group all social links of each borrower into four categories along two dimensions: whether the direct flow between borrower and friend is below or above the average direct flow, and whether the indirect flow between borrower and friend is below or above the average indirect flow. The authors then calculate the share of loans that fall into each of the resulting four categories. About 14.5 percent of loans involve borrower/lender pairs with both below-average direct ow and below-average indirect flow. Almost double as many loans involve borrower/lender pairs with either above-average direct or above-average indirect ow. About three times as many loans involve borrower and lenders with both above-average direct and above-average indirect flow. Indirect paths appear to play an important role in creating social collateral for borrowing.[3]

4 Consumption Risk-sharing

We now turn our attention to risk-sharing (**?**). The application of the social collateral model in this context is very similar to borrowing.

4.1 Motivating Example

To gain some intuition consider the three networks in Figure **??**. We assume that with probability $\frac{1}{2}$ agent s (previously the "borrower") experiences a negative endowment shock $-x$ while agent t (previously the "lender") experience a positive shock $+x$. With probability $\frac{1}{2}$ the shocks are reversed. All other agents in the economy experience no shocks. If agents have standard concave utility over consumption, the egalitarian social planner would optimally ensure that everyone in the economy consumes 0.

However, the planner's ability to redistribute endowments might be limited by the social network. For example, consider panel A of Figure **??** and the state of the world where t has the positive endowment shock. Intuitively, the planner should not expect agent t to agree to any transfer that exceeds 2 since the worst punishment that could be inflicted on her would be to lose her link which is worth 2. More generally, we should not expect that t will ever transfer more than the maximum flow between s and t: by the Ford-Fulkerson theorem any set of agents that includes t and excludes s will "cut" a set of links whose sum is greater or equal to the maximum flow $T^{st}(c)$. Moreover, there is at least one such set where the sum of cut links is exactly equal to $T^{st}(c)$. Therefore, whenever the social planner requires agent t to transfer more than $T^{st}(c)$, then agent t could assemble a coalition of agents such that the cost of potentially lost links (a proxy for the worst punishment that can be imposed by the planner on the group of deviators) is lower than t's transfer. In other words, agent t could reimburse members of her coalition for lost links instead of making a requested payment. This limits the extent of transfers between agents s and t in panels B and C of Figure **??** to a maximum of 5.

[3] **?** find a similar result in their analysis of data from Indian villages – however, they rely on a different model.

4.2 Risk-sharing Arrangements

We now turn to the formal model which allows us to analyze risk-sharing when more than 2 agents receive shocks. In our model, agents face income uncertainty due to factors such as weather shocks and crop diseases. We denote the vector of endowment realizations by $e = (e_i)_{i \in W}$, which is drawn from a commonly known joint distribution. The vector of endowments is observed by all agents.

A *risk-sharing arrangement* specifies a collection of bilateral transfer payments $t^e = \left(t^e_{ij}\right)$, where t^e_{ij} is the net dollar amount transferred from agent i to agent j in state of the world e, so that $t^e_{ij} = -t^e_{ji}$ by definition. The risk-sharing arrangement t^e implements a consumption allocation x^e where $x^e_i = e_i - \sum_j t^e_{ij}$. For simplicity, we suppress in notation the dependence of the transfers t^e_{ij} and consumption allocation x^e on e.

An agent who consumes x_i enjoys utility $U_i(x_i, c_i)$, where $c_i = \sum_j c(i,j)$ denotes the total value that agent i derives from all his relationships in the network, and U is strictly increasing and concave. The case where consumption and friendship are perfect substitutes is analytically convenient but the qualitative results can be extended to the case of imperfect substitutes. The agent's ex-ante expected payoff is $EU_i(x_i + c_i)$, where the expectation is taken over the realization of endowment shocks.

We say that a risk-sharing arrangement is *incentive compatible* if every agent i prefers to make each of his promised transfers t_{ij} rather than lose the (i,j) link and its associated value. Because consumption and friendships are perfect substitutes, incentive compatibility implies $t_{ij} \le c(i,j)$.

By construction, risk-sharing arrangement are robust to coalitional deviations. To see this, we need some definitions. For any group of agents F, we define the *perimeter* $c[F]$ of F to be sum of the values of all links between the group and the rest of the community:

$$c[F] = \sum_{i \in F,\ j \notin F} c(i,j) \tag{3}$$

Intuitively, the perimeter is the maximum extent to which the rest of the community could punish group F using ostracism. Similarly, we define the total endowment of the group as e_F and their total consumption under a risk-sharing arrangement as x_F.

Definition 5. *A consumption allocation x is coalition-proof if $e_F - x_F \le c[F]$ holds for all groups of agents F.*

It is easy to see that a risk-sharing arrangements implements a consumption allocation which is coalition-proof. Hence, no group of agents has an incentive to deviate: the net transfer between any group of agents and the rest of the community, defined as the difference between the group's total endowment and total consumption, does not exceed the sum of the values of all links connecting the group and the rest of the community.

4.3 Equivalence Result

The definition of an informal risk-sharing arrangement looks at first quite restrictive because the social network not just constrains the feasible consumption allocations but also serves as a conduit for transfers. For example, could a village elder (as a stand-in for the constrained social planner) achieve superior consumption allocations by simply taxing households with positive endowment shocks and redistribute the proceeds to households with negative shocks? This elder would not have to worry about finding a set of bilateral transfers to implement her preferred allocation and instead would only have to ensure that the final allocation is coalition-proof.

Surprisingly, the answer to this question is negative.

Theorem 2. *A consumption allocation x that is feasible ($\sum x_i = \sum e_i$) and coalition-proof can be implemented by an incentive-compatible informal risk-sharing arrangement.*

The theorem states any the elder implements *exactly the same* insurance arrangements as is possible with link-level punishment. The proof builds again on the mathematical theory of network flows.[4]

Theorem **??** has two main implications. First of all, the results shows that it is sufficient to study risk-sharing arrangements. Links matter not because they act as conduits for transfer, but because they define the costs of deviations, and hence the pattern of obligations in the community. A second implication of the theorem is that it relates the *geometry* of the network to its effectiveness for risk-sharing.

4.4 Limits of Risk-Sharing

How effective are typical social networks in sharing risk? Can local obligations to help close neighbors, relatives and friends aggregate to effective risk-sharing on the village level? The research of **?** suggests the answer should be affirmative since the full insurance model provides a surprisingly good benchmark even though it is typically rejected in the data.

It turns that *full* risk-sharing under any endowment realization is generally impossible unless the social network is extremely *expansive.* To measure expansiveness, we define the

[4] In particular, (**?**) show that finding a transfer representation for a coalition-proof allocation is equivalent to finding a flow in an auxiliary network with two additional nodes s and t added. According to the theorem of Ford and Fulkerson (1956), the maximum flow equals to the value of the minimum cut, i.e., the smallest capacity that must be deleted so that s and t end up in different components. They prove that each cut in the flow problem corresponds to a coalition, and then the coalition-proofness condition ensures that the cut values are high enough so that the desired flow can be implemented.

Figure 3: Expansion properties of three sample networks

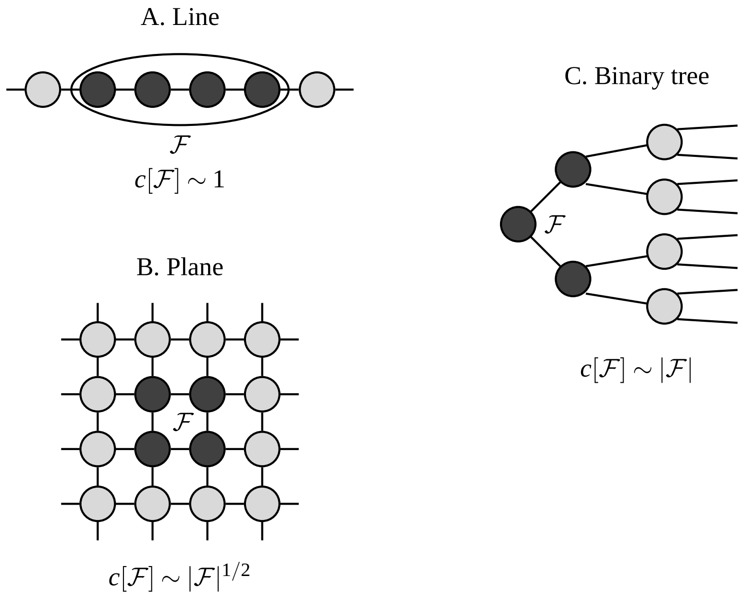

perimeter-area ratio for a set of agents F as $a[F] = c\,[F]\,/\,|F|$, where area stands for the number of agents in F. Intuitively, $a\,[F]$ represents the group's maximum obligation to the community relative to the group's size. Figure **??** shows typical sets F for three distinct geometries. Panel A shows a line which has very low expansion properties: a large connect set will have perimeter-area ratio equal to $\frac{2}{|F|}$. In contrast, the plane in panel B has significantly better expansion properties as its typical perimeter is of order $\sqrt{|F|}$. Finally, the binary tree in panel C is an expander graph whose perimeter-area ratio is always bounded away from 0 for arbitrary sets F.

Intuitively, we expect that networks with better expansion properties allow for more risk-sharing because it is more difficult to find a blocking coalition as alluded to in theorem **??**. The next proposition makes this precise. To simplify the exposition, we focus from now on the special case of i.i.d uniform endowment shocks over the interval $[-1, 1]$.[5]

Proposition 1. *[Limits to full risk-sharing] Under the aboive assumptions, equal sharing is supported by an incentive-compatible risk-sharing arrangement* ***if and only if*** *for every subset of agents F the perimeter-area ratio satisfies $a\,[F] \geq 2\left(1 - \frac{|F|}{|W|}\right)$.*

The condition implies that $a\,[F]$ must be greater than the constant $S/2$ for any set of size at most half the community. In particular, an implication for large networks is that $a\,[F]$ must be bounded away from zero since the members of F must be willing to provide resources to the

[5]The results of **?** apply to general endowment distributions as long as the tails of the distributions are not too thick and shocks are not too correlated.

rest of the community even when they all get the highest possible realization while everyone outside gets the minimum. This implies that essentially the only type of graph that allows full risk-sharing for *any* endowment realization are essentially expander graphs such as the binary tree.

4.5 Line and Plane

We now show that risk-sharing on the plane and similar networks is very good, and substantially better than on the line. It is helpful to first develop an intuition for this result

Plane networks turn out to be just sufficiently well-connected to generate very good risk-sharing in most states of the world. The key insight is that with a two-dimensional structure, outcomes in which the coalitional constraint binds under equal sharing become rare. To see the logic, consider the regular plane with the i.i.d. $[-1, 1]$ shocks. As we have seen, equal sharing fails because households in a large n by n square F would need to give up n^2 resources if all of them get a positive $+1$ shock, which is an order of magnitude larger than the perimeter $c[F] \sim n$.

The key is that for large n, such extreme realizations are unlikely, and in typical realizations the required transfers do not exceed the perimeter. With i.i.d. shocks, the standard deviation of the group's endowment is only n, which is only of order n even though it is the sum of n^2 random variables – intuitively, a lot of the idiosyncratic shocks cancel out within the group. Thus the *typical shock* in F has the same order of magnitude as the maximum pledgeable amount, and hence potentially deviating coalitions are rare. By way of contrast, the argument breaks down for the line, since the perimeter of even large interval sets is only 2, a constant.

To formulize these ideas, we assume that agents have quadratic utility function such that we can express the average utility loss relative to the benchmark of equal sharing as

$$SDISP(x) = \left[\mathrm{E}\frac{1}{|W|}\sum_{i \in W}(x_i - \overline{e})^2\right]^{1/2}, \tag{4}$$

which is the square-root of the expected cross-sectional variance of x. For non-quadratic utilities, $SDISP(x)$ can be interpreted as a second order approximation of the utility based measure.

Proposition 2. *There exist positive constants K, K' and K'' such that*

(i) On the infinite line with capacities c and i.i.d. shocks, we have $SDISP(x) \geq K/c$ for all incentive-compatible risk-sharing arrangements.

(ii) On the infinite plane with capacities c, we have $SDISP(x) \leq K' \exp\left[-K'' c^{2/3}\right]$ for some incentive-compatible risk-sharing arrangement.

This Proposition characterizes the rate of convergence to full risk-sharing as capacities increase. The contrast between the line and plane is remarkable. Risk-sharing is relatively poor on the line: $SDISP$ goes to zero at a slow polynomial rate of $1/c$ as c goes to infinity. In contrast, the rate of convergence for the plane is exponentially fast,confirming our intuition that agents are able to share typical shocks due to the more expansive structure.

The difference in the rates of convergence become quickly apparent in simulations. Figure **??** compares risk-sharing on equally-sized line and plane (100 agents each) while fixing initial endowments and the total capacity per agent across both networks.[6] SDISP declines rapidly on the plane and full risk-sharing is already achieved at capacity 1.4 per agent. In contrast, SDISP declines more slowly on the line and full risk-sharing requires capacity per agent far exceeding 2.

? extend the result for the plane to *geographic networks* which exhibit a two-dimensional sub-structure but are less regular than the plane. Using data from Peruvian social networks they show that real-world networks are geographic networks but not expander graphs. Hence, the social collateral model can explain very good risk-sharing in real-world social networks where agents only have local obligations to a small subset of the population (such as close neighbors, friends and relatives).

4.6 Risk-sharing Islands

? also characterize the micro-structure of specific risk-sharing arrangements. Intuitively, the network is partioned into a set of contiguous "risk-sharing islands" as shown in Figure **??** such that agents within the same island consume the same amount while agents in neighboring island consume either more or less. Moreover, the IC constraints for transfer across islands bing while they are slack within islands.

This phenomenon has two important implications.

Local sharing. When an agent in the interior of an island receives an endowment shock she will share the risk first *locally* with other neighbors in her own risk-sharing island. If consumption within the island increases (or decreases) sufficiently so that IC constraints to neighboring islands no longer bind, then the shock will also be shared with neighboring islands. In that sense, risk-sharing in the social collateral model is local.

Endogenous socialization. Compare two geometries, such as line and plane, where risk-sharing islands have the same average size (which translates into comparable risk-sharing in both networks). Then the more expansive network will tend to have a higher share of agents at the boundary of a risk-sharing island. Therefore, the incentive of agents to invest in socialization is at the margin greater in the more expansive network. Hence, the very

[6] If we think of capacities as social collateral then this exercise compared a linear and planar network with the same amount of social capital per agent.

Figure 4: Risk-sharing simulations on line and plane for different capacities

Panel A: initial endowments (uniform over $[-1, 1]$)

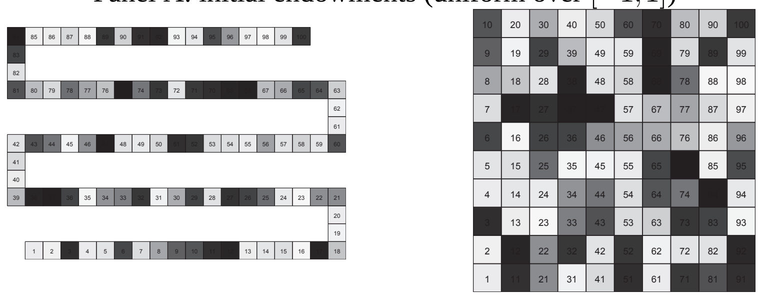

SDISP= 0.556

SDISP= 0.556

Panel B: risk-sharing with total capacity 1 per agent

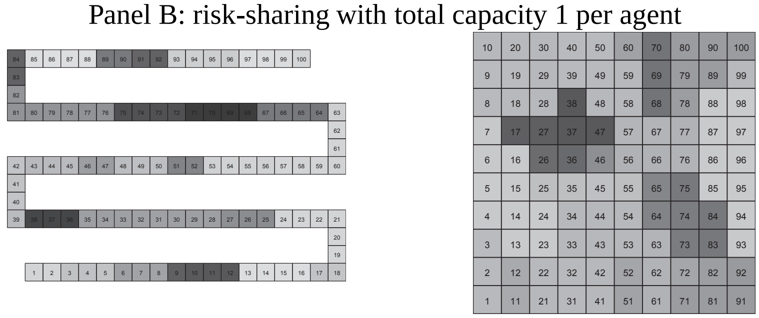

SDISP= 0.246
30 islands

SDISP= 0.127
17 islands

Panel C: risk-sharing with total capacity 1.4 per agent

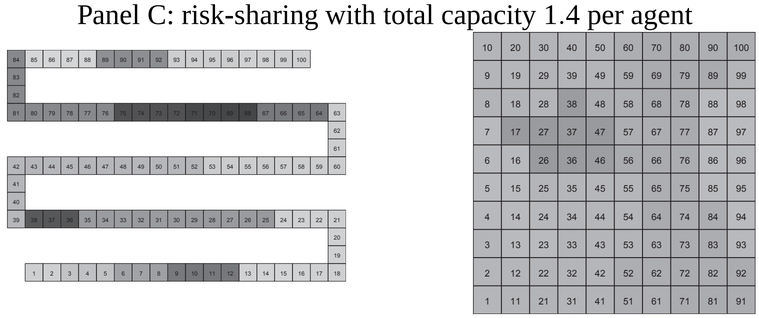

SDISP= 0.199
17 islands

SDISP= 0.035
4 islands

Panel D: risk-sharing with total capacity 2 per agent

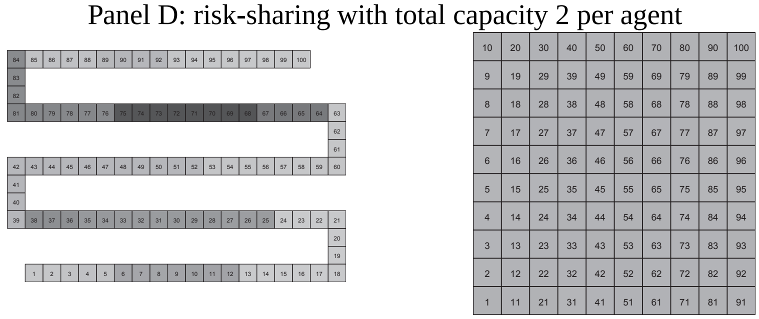

SDISP= 0.148
13 islands

SDISP= 0
1 island

features which create good risk-sharing on expansive networks such as the plane *also* make these networks more stable because they increase agents' incentives to invest in socialization.

5 Other Mechanisms

At this point, it is useful to contrast the social collateral approach to other theory frameworks.

5.1 Altruism

The social collateral approach assumes that agents are selfish and attributes any lending or risk-sharing to repeated game effects. However, we might expect that people feel particularly altruistic towards family, friends and neighbor and might for this reason alone provide help. **?** analyze this question by matching student subjects to direct and indirect friends at various social distances and have these pairs of subjects play a series of dictator games. They use a within-subject design where the recipient either finds out or does not find out the dictator's identity. This allows them to distinguish *directed altruism* (being nice to one's friends due to "warm glow") from repeated game effects. There is substantial directed altruism towards direct friends but indirect friends are treated similarly to a randomly picked nameless individuals within the network. Hence, directed altruism might not be able to explain transfers within the network that do not involve direct friends. In a closely related paper, **?** have villagers in rural Paraguay play variants of the dictator game to examine motives for sharing. They correlate behavior measured in the experiment with the real world sharing outside the experiment to find that repeated-game motives seem to better explain sharing in the real world.

Earlier work on lifetime (inter-vivos) transfers also rejects the altruism hypothesis. For example, **?** find that parents increase transfers to children by only 13 cents for every dollar that is redistributed from child to parent. **?** and **?** analyze the patterns of intergenerational transfers and also find that they are consistent with exchange motives rather than altruism.

5.2 Sharing Rules and Bargaining

? abstract away from the enforcement problem within connected networks and instead assume that agents within a connected component share resources equally (for example, social neighbors might repeatedly pool and share their income which eventually results in equal sharing across the component). **?** extend this basic idea and analyze an environment where agents within a connected component bargain over the joint surplus. They show that the surplus is allocated according to the Myerson value where more central agents receive higher shares. In these papers, the network defines the bargaining position of agents rather than enforcement as in the social collateral approach.[7]

5.3 Other Repeated Game Models

There are a number of papers in the literature which also study enforcement in a repeated game context.

? formulate a theory of limited commitment to explain deviations from full risk-sharing in village economies. They test their model with data from three Indian villages and find that the model can fully explain the dynamic response of consumption to income, but cannot explain the distribution of consumption across households (also see **?** for empirical evidence). The social collateral model can be thought as a special case of the limited commitment model where the social network defines the set of constraints on bilateral transfers.

? analyze a model where agents' needs and transfers are not perfectly observed (unlike in the social collateral model). This raises a new set of questions such as whether agents have an incentive to communicate deviations and and how quickly information spreads within the network. They show that equilibria with severe group punishments (permanent ostracism) are difficult to sustain because cheated-upon agents might not have an incentive to communicate truthfully.

Finally, **?** propose a "social quilt" model where agents play continuous-time Prisoner's dilemma games with their social neighbors under perfect information. The authors focus on equilibria that are renegotiation-proof and characterize the set of stable networks. They show that these network only include "supported links" such that any two friends who exchange favors have a common friend. The empirical predictions of their model are in fact very similar to the ones reported in Section **??**. The biggest difference is that the social collateral approach takes the social network as exogenously given and then finds the set of feasible borrowing or risk-sharing arrangements within that context.

5.4 Endogenous Network

Another strand of the literature analyzes network formation when agents form links to mitigate risk. Endogenous network formation can be inefficient because agents do not fully internalize the benefits and costs of forming links (**?**, **?**). **?** provide evidence from rural India that families arranges marriages of daughters to distant locations to mitigate income shocks.[8] **?** find that geographic proximity is a strong correlate of risk sharing networks, likely because it facilitates monitoring and enforcement (also see **?**).

[7] In related empirical work, **?** use PSID data to show that risk-sharing within the extended family is not independent of the distribution of resources.

[8] However, **?** argues the existence of sub-caste networks that provide mutual insurance to their members restricts marriage mobility.

6 Conclusion

The social collateral approach provides an analytically tractable and empirically relevant way of modeling informal transfers in social networks. This approach has been applied to analyze (1) borrowing and trust in networks and (2) consumption risk-sharing.