Consumption Risk-sharing in Social Networks*

Attila Ambrus
Duke University

Markus Mobius
Iowa State University, Microsoft Research, NBER

Adam Szeidl
Central European University, NBER, CEPR

November 2012

Abstract

We develop a model of informal risk-sharing in social networks, in which relationships between individuals can be used as social collateral to enforce insurance payments. We characterize incentive compatible risk-sharing arrangements and obtain two results. (1) The degree of informal insurance is governed by the expansiveness of the network, measured by the number of connections that groups of agents have with the rest of the community, relative to group size. Two-dimensional networks, where people have connections in multiple directions, are sufficiently expansive to allow very good risk-sharing. We show that social networks in Peruvian villages satisfy this dimensionality property; thus, our model can explain Townsend's (1994) puzzling observation that village communities often exhibit close to full insurance. (2) In second-best arrangements, agents organize in endogenous "risk-sharing islands" in the network, where shocks are shared fully within, but imperfectly across islands. As a result, network based risk-sharing is local: socially closer agents insure each other more. We also discuss how our results extend to an environment in which social collateral is endogenous.

Keywords: informal insurance, social collateral, coalition-proofness, perimeter-area ratio, geographic networks, risk-sharing islands

*E-mails: attila.ambrus@duke.edu, markusmobius@gmail.com, szeidla@ceu.hu. We are grateful to Dean Karlan and Tanya Rosenblat for allowing us to use data they collected jointly with Markus Mobius. We thank Daron Acemoglu, In Koo Cho, Erica Field, Drew Fudenberg, Andrea Galeotti, Matthew Jackson, Eric Maskin, Stephen Morris, Gabor Pete, Debraj Ray, Laura Schechter and seminar participants for helpful comments and suggestions, and the National Science Foundation (Award #0752835) for financial support. Szeidl thanks the Alfred P. Sloan Foundation and the European Research Council (StG #283484) for financial support. Most of the proofs as well as extensions and additional material to the paper are contained in a supplementary appendix available at `http://www.socialcollateral.org/risksharing/supplementary_appendix.pdf` (Ambrus, Mobius and Szeidl 2012).

In much of the developing world, people face severe income fluctuations due to weather shocks, diseases affecting crops and livestock, and other factors. These fluctuations are costly because households are poor and lack access to formal insurance markets. Informal risk-sharing arrangements, which help cope with this risk through transfers and gifts, are therefore widespread. For example, Figure 1 depicts financial and in-kind transfers between relatives and friends in a rural village in the Huaraz province of Peru.[1]

Development economists have studied both the pattern of informal transfers and their effectiveness in sharing risk. Two seemingly contradictory findings have been documented. On the one hand, these arrangements often seem to be based on *local obligations*, as people mainly help out close neighbors, relatives and friends (Udry 1994). On the other hand, these local mechanisms often achieve almost full *global* insurance on the village level. For example, Townsend (1994) argues that the full insurance model provides a surprisingly good benchmark even though it is typically rejected in the data.[2]

How do local obligations and transfers aggregate up to good global risk-sharing? To shed light on this question, in Section 1 we build a simple model of risk-sharing in social networks. In our model, full insurance is difficult to obtain because it requires a high level of connectedness that we do not observe in real social network data. However, consistent with the evidence, we also show that close to perfect risk-sharing *can* be achieved for the type of more loosely connected social networks that we do observe. Our model also allows us to study the nature of informal risk-sharing arrangements. We show that households' consumption will comove more strongly with that of socially closer households, a prediction consistent with the empirical findings in Angelucci, Giorgi and Rasul (2012), who therefore provide indirect evidence for our model.

We model the social network as a set of pre-existing relationships, such as friendships and family ties. These links have utility values, which represent either the direct consumption value of relationships, or indirect benefits from future transactions. We define a risk-sharing arrangement as a set of transfers between direct neighbors in the social network in every

[1] The data used in constructing this Figure were collected by Karlan, Mobius and Rosenblat (2007). See Appendix B for details.

[2] Also see Ogaki and Zhang (2001) and Mazzocco (2007).

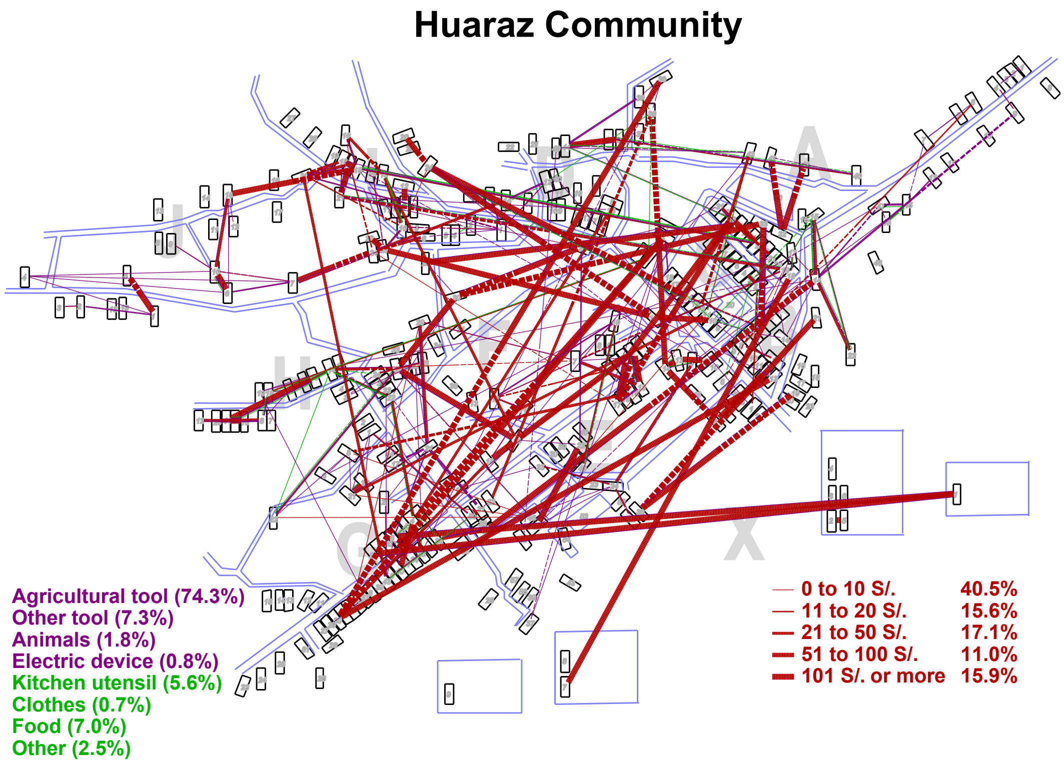

Figure 1: Financial and real transactions between relatives and friends in a rural community in Peru, represented as lines between transacting parties in the village map. Thickness of line measures value of transaction in Peruvian New Soles.

state of the world. This arrangement is subject to moral hazard: ex post, an agent who is expected to make a transfer to a network neighbor may prefer to deviate and withhold payment. In our model, such deviations result in the loss of the affected link. Intuitively, network links serve as social collateral ensuring that agents live up to their obligations under the informal risk-sharing arrangement.

In Section 2 we state our basic theoretical result, establishing an equivalence between this simple model in which an individual deviation is punished by the loss of a link with the cheated friend, and a more realistic model in which a group deviation is punished, through ostracism, by the rest of the community. In this more realistic model with ostracism and group deviations, a consumption allocation can be implemented if the net transfer from any group of agents to the rest of the community does not exceed the sum of the values of all links between the group and the community. Then, the intuition for the equivalence with link-level punishments is that individual obligations embedded in the value of links build up to group obligations represented by the total value of links connecting the group with the larger community.

The equivalence between individually rational arrangements with link-level enforcement and coalition-proof arrangements with ostracism has two implications. First, it shows that decentralized insurance arrangements with link-level enforcement can also be implemented in a centralized fashion through intermediaries such as trusted village elders, who respect the obligations of each group (e.g., extended family) in the community. Second, the result relates the *geometry* of the network to its effectiveness for risk-sharing, allowing us to study how local links aggregate to social capital at the community level.

The key property of network structure identified by our equivalence result is called *expansiveness*, and measures the number of connections that groups of agents have with the rest of the community relative to group size. To gain intuition about this property, consider the three example networks in Figure 2. Among these networks, the infinite line in Figure 2A is the least expansive, because any connected set of agents always has only two links with the rest of the community. The infinite "plane" network of Figure 2B is more expansive, while the infinite binary tree of Figure 2C is the most expansive network of all, where the number of outgoing links for any set grows at least proportionally with its size.

We show that full insurance requires highly expansive networks like the infinite binary tree. However, we do not find that real-world social networks in rural villages in Peru exhibit this large degree of expansiveness. Instead, these social networks are more similar to planar networks, possibly because people tend to have connections in multiple directions at close geographic distance. We next show that a two-dimensional structure, such as the one found in our Peruvian data, is sufficient to ensure very good risk-sharing in most states of the world. For an intuition, consider a connected group of agents in the plane network. With idiosyncratic shocks, the standard deviation of the total endowment of the group is proportional to the square root of group size. But on the plane, the number of outgoing links from the group is also at least proportional to the square root of size (the worst case would be when the group has a square shape). Thus group obligations with the rest of the community – links connecting the group with the network – are of the same order of magnitude as group shocks. Since this holds for every group, it follows that "almost" full risk-sharing

Figure 2: Expansion properties of three example networks. The parameter-area ratio c[F] is defined as the number of links leaving the set F (perimenter) divded by the number of agents inside the set (area). The perimeter-area ratio of a typical set in the network describes the expansiveness of the geometry.

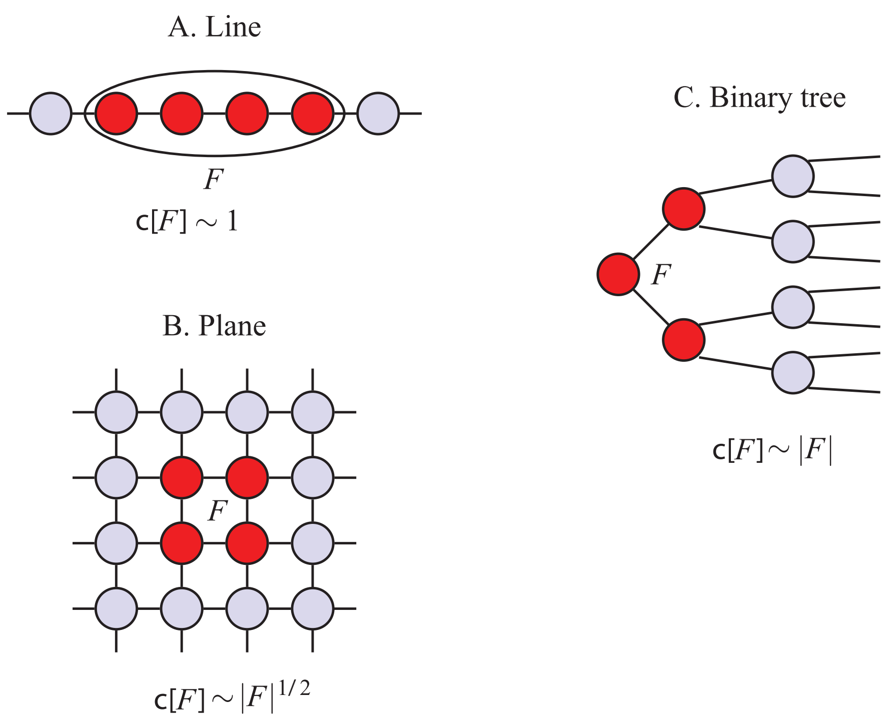

can be implemented in the network. This argument applies not just for the regular plane network, but for any social network which has a two-dimensional sub-structure. We call these networks *geographic networks* and we show that our Peruvian village networks fall into this class. As a result, our model provides a potential explanation for the informal insurance puzzle highlighted by Townsend.

The above results constitute a quantitative analysis of informal risk-sharing. Section 3 presents our second main contribution, a qualitative analysis of constrained efficient "second-best" arrangements. We show that in these arrangements, for every realization of uncertainty the network can be partitioned into endogenously organized connected groups called "risk-sharing islands". This partition has the property that shocks are completely shared within, but only imperfectly across islands. The island structure can be understood in terms of "almost deviating coalitions," that are indifferent between staying in the network and deviating as a group. Islands are maximal connected sets subject to the constraint that they are not divided by any almost deviating coalition; therefore, insurance *across* island boundaries is limited, but insurance *within* islands is complete. The size and location of these risk-pooling islands is endogenously determined by the social structure and the realization of endowment shocks, consistent with evidence documented by Attanasio, Barr, Cardenasy, Genicot and Meghir (2009), and distinguishing our model from theories with exogenously specified risk-sharing groups.

A key implication of the islands result is that an agent's consumption will comove more with the consumption of closely connected neighbors. This follows because islands are connected subgraphs: agents who are socially closer are more likely to belong to the same island and thus provide more insurance. This observation helps characterize informal insurance as a function of shock size. Risk-sharing works well for relatively small shocks: sharing islands are large, and both direct and indirect friends help out. As the size of the shock increases, only close friends help with the additional burden; and risk-sharing completely breaks down for large shocks. Some of these predictions are confirmed in the empirical work of Angelucci et al. (2012).

In Section 4 we examine how our qualitative findings extend to a setting in which the network structure is given, as before, but link capacities are determined endogenously through costly socializing. A basic intuition we highlight is that the marginal value of extra socializing is related to the likelihood that an agent is at the boundary of a risk-sharing island, because the it is only in such events that the agent's transfer constraints are binding. This logic implies that for low capacity levels—that is, when socializing is costly— the incentives to socialize are increasing the likelihood of having islands with large boundaries, i.e., the expansiveness of the network, further strengthening the results obtained in our basic model. At higher capacity levels—that is, when socialization is inexpensive—this relationship is eventually reversed because the better insurance provided by expansive networks also reduces the benfits of further insurance. We demonstrate with simulations the implication of this logic that for costly socialization, equilibrium link capacities are higher in the (more expansive) plane than on the line, amplifying our basic result that plane-like networks yield significantly better risk-sharing.

In the concluding Section 5 we discuss some further research directions and caveats with our model. Proofs are delegated to Appendix A and a supplementary appendix.

Our paper builds on a growing literature studying informal insurance in networks. Bloch, Genicot and Ray (2008) develop a model with both informational and commitment constraints, and characterize network structures that are stable under certain exogenously specified risk-sharing arrangements. We conduct the opposite investigation: taking the network as given, we study the degree and structure of informal risk-sharing. Bramoulle and Kranton (2006) also study insurance arrangements in networks, but in their model there are no enforcement constraints. Our modeling approach builds on Karlan, Mobius, Rosenblat and Szeidl (2009), who explore informal borrowing in networks.[3] Empirical work in this area includes De Weerdt and Dercon (2006), Fafchamps and Lund (2003) and Fafchamps and Gubert (2007), who use data on village networks, Attanasio et al. (2009) who document the importance of social ties for risk-pooling, while Mazzocco (2007) emphasizes the role of within-caste transfers.[4]

1 A model of risk-sharing in the network

1.1 Model setup

In our model, agents face income uncertainty due to factors such as weather shocks and crop diseases. In the absence of a formal insurance market, agents can agree on an informal risk-sharing agreement that specifies transfers between pairs of agents in each state of the world. These transfers are secured by the social network: connections in the network have an associated consumption value that is lost if an agent fails to make a promised transfer.

Formally, a social network $G = (W, L)$ consists of a set W of agents (vertices) and

[3] See also Ali and Miller (2008), who study network formation with repeated games and Dixit (2003), who compares relational and formal governance in a circle network.

[4] More broadly, our work contributes to the growing literature on informal institutions. Kandori (1992), Ellison (1994) and Greif (1993) develop game-theoretic models of community enforcement, and Kranton (1996) studies the interaction between relational and formal markets. In the context of consumption insurance, Ligon (1998), Coate and Ravaillon (1993), Kocherlakota (1996) and Ligon, Thomas and Worrall (2002) explore related models with limited commitment, while Mace (1991) and Cochrane (1991) are influential empirical studies of consumption insurance. These papers do not study the effects of network structure.

a set L of links, where a link is an unordered pair of distinct vertices. Unless otherwise stated, we assume that the network is finite; the supplementary appendix discusses how to extend our setup to infinite networks. Each link in the network represents a friendship or business relationship between the two parties involved. We assume that the strength of these relationships is determined outside the model, and that they are measured by a capacity.

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

The capacity of an (i, j) link measures the benefit that i derives from his relationship with j. These benefits can represent the direct utility that agents derive from interacting with each other, or the utility or monetary value of economic interaction in the present or in future periods. For ease of presentation, we assume that the strength of relationships is symmetric, so that $c(i, j) = c(j, i)$ for all i and j. All our results extend to the case with asymmetric capacities.

Agents in this economy face uncertainty in the form of endowment risk. 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 the dependence of the transfers t^e_{ij} and consumption allocation x^e on e for the rest of the paper.

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. To simplify exposition, in the body of the paper we focus on the analytically convenient case where consumption and friendship are perfect substitutes, so that the utility of i is $U_i (x_i + c_i)$. In the Supplementary Appendix we extend the model to the case when consumption and friendship are imperfect substitutes, and show that under weak conditions, our qualitative conclusions extend. 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 $t_{ij} \leq c(i, j)$.

1.2 Discussion of modeling assumptions

Risk-sharing arrangement. The most literal interpretation of these arrangements, in the spirit of Arrow and Debreu, is that agents choose an *ex ante* informal contract, which specifies payments for every conceivable realization of uncertainty. Alternatively, the consumption allocation may also be determined *ex post* by a social norm that specifies how to reallocate goods among connected agents. For example, Fafchamps and Lund (2003) describe how informal insurance is implemented through a collection of bilateral "quasi-loans," where households borrow from neighbors, who expect their kindness returned when they themselves are hit by adverse shocks.

Exogenous capacities. We analyze a one-time risk-sharing arrangement in a network where links and capacities are determined outside the model. The most direct interpretation of this framework is that link values are generated by a number of social activities and services besides risk-sharing. In this interpretation, the links themselves may be created through a long term network formation process largely shaped by factors outside our model, such kinship and geographic proximity. An alternative view is that link capacities are shaped endogenously by the insurance benefits that they generate. One approach to modeling this effect is to allow agents to invest in socializing: higher socializing leads to higher capacities and hence greater insurance. We explore this extension of our framework in Section 4. In an even richer environment with explicit dynamics, the value of a network connection might be determined in part by the ability to conduct insurance transactions through the link in future periods. As Bloch et al. (2008) show in a related model, this leads to restrictions on the equilibrium network structure and link values. We leave the investigation of such a framework for future research.

Incentive compatibility. Our notion of incentive compatibility is motivated by Karlan et al. (2009). In their model of informal borrowing, a link between two agents is destroyed if a promised transfer is not made. They develop explicit micro-foundations for this assumption where the failure to make a transfer is a signal that the agent no longer values his friend, in which case these former friends find it optimal not to interact with each other in the future.[5] An alternative justification is that people break a link for emotional or instinctive reasons when a promise is not kept; Fehr and Gachter (2000) provide evidence for such behavior.

Full information. Our model assumes that agents in the community can observe the vector of endowment realization so that they know what transfer payments to expect from their neighbors and how much to send. Full information about endowments seems reasonable in many village environments, in which individuals can easily observe the state of livestock or crops. For example, Udry (1994), shows that asymmetric information between borrowers and lenders is relatively unimportant in villages in Northern Nigeria.

1.3 Equivalence of Link-Level Punishments with Individual Deviations and Ostracism with Coalitional Deviations

This Section establishes our main theoretical result: that our model of link-level punishments is equivalent to ostracism-based enforcement in the presence of coalitional deviations. A plausible and commonly explored way of enforcing cooperation in social interactions is ostracism, in which a deviator is punished by all his network neighbors cutting their links with him.[6] It is easy to see that, absent other constraints, this type of enforcement mechanism—because the potential punishment following a deviation is larger—can implement higher levels of sharing than our basic model.

Yet, by only considering individual deviators, this form of ostracism abstracts away from the possibility of people siding with their close friends, and hence seems implausibly strong. For example, it seems unlikely that a person would punish a cousin or a sister just because she defected on a common acquaintance. To address this issue, we propose a version of ostracism which allows not only individuals, but also coalitions to deviate.[7] To illustrate why coalitional

[5] In the supplementary appendix we develop similar foundations for the present model, in which the value of connections is earned in a "friendship game." See Ambrus et al. (2012) (available at `http://www.socialcollateral.org/risksharing/supplementary_appendix.pdf`).

[6] Versions of this idea are explored in Greif (1993), Kandori (1992) and Dixit (2003).

[7] Genicot and Ray (2003) follow a similar approach in a model of group formation.

deviations help in this matter, suppose that i, j and k form a triangle network, and that k is a weak friend of both i and j, who in turn are strongly connected cousins. In this network, ostracism against individual deviators could enforce a large transfer from i to k, because, in the event that i defaults on that transfer, she would be badly punished by the loss of both her links. But if we allow for coalitional deviations as well, then—because the strong connection with a cousin is more valuable than a weak connection to an acquaintance—agents i and j may collectively find it more profitable to cut their weak links to k and redistribute the required payment among themselves. Thus, coalitional deviations, by allowing people to side with their close social contacts, impose additional plausible restrictions on the set of arrangements.

To formalize the idea of ostracism in the presence of coalitional deviations, 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:

$$c[F] = \sum_{i \in F,\ j \notin F} c(i,j) \tag{1}$$

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 2 *A risk-sharing arrangement is coalition-proof if $e_F - x_F \leq c[F]$ holds for all groups of agents F.*

The arrangement is coalition-proof if no group 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. In this definition we only look at the incentives of the coalition as a whole; but in the supplementary appendix we show that, in our context, the simple notion of coalition-proofness we use above is equivalent to coalition-proofness along the lines of Bernheim, Peleg and Whinston (1987), i.e., allowing only for credible coalitional deviations that are not prone to further credible deviations by subcoalitions. The intuition behind this is that in our framework any such further deviation by a subcoalition is also a profitable coalitional deviation in the first place (i.e., even in the absence of the original deviation). Also note that the extent of ostracism we allow for in this definition—given the possibility of coalitional deviations—is the harshest possible. More limited ostracism, such as punishing a coalition by only those who are within a given social distance of the agents who have been defected on, would therefore yield lower risk-sharing.

Theorem 1 *A consumption allocation x that is feasible ($\sum x_i = \sum e_i$) is suppported by ostracism in the presence of coalitional deviations **if and only if** it can be implemented by an incentive-compatible informal risk-sharing arrangement.*

The theorem states that ostracism, when combined with coalitional deviations, implements *exactly the same* insurance arrangements as link-level punishment. In essence, we have two opposing forces: while ostracising individual deviators increases the set of enforceable allocations, allowing for coalitional deviators reduces it. In the perfect substitutes environment these two forces exactly cancel. To understand the intuition for the Theorem, first note that one direction is immediate. Any arrangement that can be implemented by link-level punishments can also be implemented by coalitional ostracism: since each transfer is bounded by the capacity of the link, the same inequality must also hold when transfers are added up along the perimeter of a group.

Showing the converse—that coalition-proof ostracism cannot implement more than link level punishments—is more difficult, and builds on the mathematical theory of network flows. For an intuition, consider a feasible and coalition-proof consumption allocation x. To implement this allocation with link-level punishments, we need a set of transfers which—respecting the capacity constraints over links—move money from those who, in autarky, have "too much" ($e_i > x_i$) to those who, in autarky, have "too little" ($e_i < x_i$) relative to the target level of consumption. To build intuition for why such transfers exist, imagine that t is the transfer arrangement that gets "closest" to implementing x. Given the allocation implemented by t, let F denote the set of all agents to whom, respecting the capacity constraints, additional consumption goods from agents with $e_i > x_i$ can still be transferred through the network. They key insight is that unless t implements x, the set F forms a blocking coalition for arrangement x, contradicting the assumption that x is coalition-proof. This follows because—by its construction as the *maximal* set of agents to whom resources can still flow—no additional amount can be sent through the perimeter of F, violating the coalitional constraint $e_F - x_F \leq c[F]$ unless x is already implemented by t.

A natural question about the Theorem is whether a weaker version of coalition-proofness, in which only a smaller set of coalitions—e.g., those with a limited number of participants—are allowed to deviate is sufficient for the equivalence. The answer to this question is negative. To see why, consider the "island" network in Figure 3, which is a complete network which consists of two equal-sized communities. For concreteness, suppose that there are 100 agents in each community, that all within-community links have equal capacities of 100, and that all cross-community links also have equal capacities of 0.01. Consider the arrangement which sends, from the first to the second community, 0.01 units of the consumption good over every cross-community link. This arrangement transfers in total 100 units of consumption: each agent in the first community contributes one unit which is equally distributed to all agents in the second community. Because capacity constraints are satisfied, this is an incentive-compatible transfer arrangement; but because all links used in the arrangement are operating at full capacity, no additional transfer from the first to the second community would be incentive compatible. When looking at this arrangement from the perspective of coalitions, the binding constraint which does not permit additional transfers corresponds to the coalitional deviation of the first community. Thus, in this example, the "local" link-level constraints map into a "global" coalitional constraint in which the blocking coalition corresponds to half of the entire network.

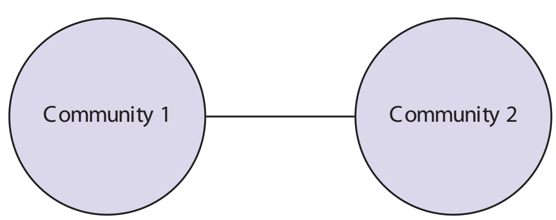

Figure 3: Network of two "islands" with strong intra-island and weak inter-island links.

The theorem has two main implications. First, it shows how individual obligations aggregate up to social capital at the community level. 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. In particular, a coalition-proof arrangement does not have to be implemented by transfers over links: intermediaries such as village elders could also collect and distribute resources, as long as they respect the obligations of each group of households, i.e., coalition-proofness.[8] Hence our model need not predict long chains of transfers in practice: these chains are likely to be shortened by intermediaries.

A second implication of the theorem is that it relates the *geometry* of the network to its effectiveness for risk-sharing. This connection forms the basis of our analysis in the following section.

2 The limits to risk-sharing

In this section we use the equivalence between incentive compatibility and coalition-proofness to explore how much risk-sharing can be obtained in a given network. Our central finding is that good risk-sharing requires social networks to have good "expansion properties"; that is, all groups of agents should have enough connections with the rest of the community, relative to group size.

2.1 Limits to full risk-sharing

We first use Theorem 1 to establish a negative result: full risk-sharing cannot be achieved unless the network is extremely expansive, because coalitions with a relatively low "group obligation" $c[F]$ will choose to deviate in some states.

To build intuition, consider the infinite line, plane and binary tree networks depicted in Figure 2, where all link capacities are equal to a fixed number c.[9] For these examples, we assume that endowment shocks are independent across agents, and take values $+\sigma$ or $-\sigma$ with equal probability. We focus on implementing equal sharing, i.e., an arrangement where

[8] At the extreme, a single trusted intermediary could implement the allocation by collecting a "tax" of $e_i - x_i$ from each agent i for whom this is positive, and use these funds to pay the unlucky agents for whom $e_i - x_i$ is negative.

[9] We consider infinite networks here because they are useful for building intuition.

all agents consume the per capita average endowment. This allocation is Pareto-optimal when agents have identical preferences over consumption. Since our example networks are infinite, the law of large numbers implies that the average endowment is zero; equal sharing thus requires all agents to consume zero with probability one.

Consider an interval set of consecutive agents F on the circle network (see Figure 2A). The coalitional constraint for F is most likely to bind in the positive probability event where all agents in F receive a positive shock $+\sigma$. In this event, the zero consumption profile dictates that members of F give $|F| \cdot \sigma$ to the rest of the community; but they can only commit to giving up $c[F] = 2c$. Coalition proofness thus requires $2c \geq |F| \cdot \sigma$ for all F. However, for any fixed c, this is violated for long enough intervals F. A similar negative result holds for the more expansive plane network in Figure 2B. The perimeter of a square-shaped set F is $c[F] = 4c\sqrt{|F|}$; for a large enough square, this is smaller than $|F| \cdot \sigma$, which is how much members of F would have to give up if they all get a positive shock $+\sigma$.

However, these perimeter bounds do not rule out equal sharing for the yet more expansive binary tree in Figure 2C. Here, the perimeter of any set F is at least $c \cdot |F|$, and so for $c \geq \sigma$, no coalition of agents has to give up more than their group obligation in any realization.

These examples suggest that equal sharing can only be incentive compatible in networks with good expansion properties, i.e., where the perimeter of sets grows in proportion with set size. To measure expansiveness, we define the "perimeter-area ratio" $a[F] = c[F] / |F|$, 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. The next result tightens the connection between expansiveness and insurance by characterizing full risk-sharing in *any network* in terms of $a[F]$, under the assumptions that (1) the support of e_i is the same compact interval of length S for all agents; and (2) the support of e_i given any realization of (e_{-i}) is the same as its unconditional support, for all i.[10]

Proposition 1 *[Limits to full risk-sharing] Under the above 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 $a[F] \geq \left(1 - \frac{|F|}{|W|}\right) S$.*

[10]Bloch et al. (2008) impose the same condition on endowment shocks in their Assumption 1.

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 for such sets as the network size grows without bound: because the the members of F must be willing to provide resources to the rest of the community even when they all get the highest possible realization while everyone outside gets the minimum. The above inequality ensures that the group has a large enough perimeter to credibly pledge the required resources even in such extreme realizations. The condition is violated for big groups on the line and plane networks because $a[F]$ can be arbitrarily small, and only holds for highly expansive graphs like the binary tree.[11]

To further illustrate the implications of the Proposition, consider the two-island network in Figure 3. This is a complete network in which each island has $N/2$ agents, each within-island link has capacity c_i and each cross-island link has capacity c_o. We assume that the island network exhibits homophily, i.e., that within-island links are stronger: $c_i \geq c_o$. We let $\bar{c} = (N/2-1)c_i + (N/2)c_o$ denote the per capita total capacity. The homophily index (Golub and Jackson 2012) of a group can be defined as the share of the capacity of within-group links relative to the capacity of all links that a group has, $H = (N/2-1)c_i/\bar{c}$. Now suppose that agents in this network are exposed to shock as above, and we attempt to implement equal sharing. Clearly the realizations in which it is the most difficult to achieve equal sharing are when all agents in one island have a positive, and all agents in the other island have a negative realization, i.e., when F is one of the islands. The condition in the Proposition for this case simplifies to $(N/2)c_o \geq \sigma$, or equivalently $\bar{c}(1-H) \geq \sigma$. Intuitively, in this network full insurance is easier to implement if either link capacities are strong ($\bar{c}$ high) or homophily is weak (H is low).[12]

Full insurance in real world networks. We use data from a village community in Huaraz, Peru to show that real-world networks are unlikely to be expansive enough to allow for full insurance.[13]

[11] Families of networks where the perimeter-area ratio is bounded below by a positive constant are called "expander graphs" in the computer science literature.

[12] Note that, unlike in the line and plane examples, here $a[F] = \bar{c}(1-H)$ *is* bounded away from zero when $H < 1$. Thus, fixing $\bar{c}$ at a high enough constant value, the islands network—and in particular the complete network where all capacities are identical ($c_i = c_o$)—is sufficiently expansive to implement full insurance for any number of agents N.

[13] The data was collected by Dean Karlan, Markus Mobius and Tanya Rosenblat and is described in

Figure 4: Expansiveness of the social network in Huaraz, Peru

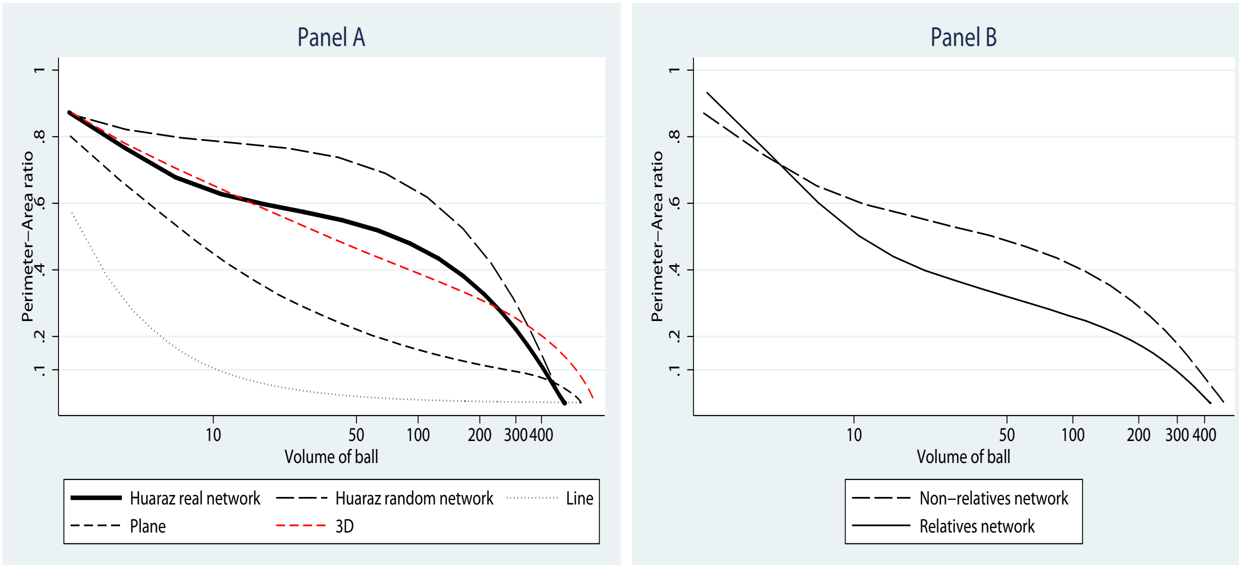

Figure 4A compares the expansiveness of the Huaraz network with the line and plane (with approximately equal number of agents) as well as a finite random network with the same degree distribution as the Huaraz network. We use the latter network as a proxy for the most expansive tree-like network that could be achieved in the Huaraz village community.[14] For all these networks, link capacities are assumed to be equal across links and normalized so that the per household average capacity is one. To measure expansiveness, we construct, for each household, a collection of "ball" sets which contain all households within a fixed social distance r. We then calculate the average of the perimeter-area ratio and set size for each r, and plot the perimeter-area ratio as a function of size for all four networks. Comparing across our three example networks illustrates our earlier discussion: the perimeter-area ratio goes to zero quickly for the line network, goes to zero more slowly for the plane, and least slowly for the random netw

The key curve in the Figure is the solid line representing the actual social network in Huaraz. This curve lies slightly above the plane but well below the random network, and approaches zero as set sizes grow, with a slope that parallels the curve for the plane. In fact,

Appendix B in more detail.

[14]There are formal results in the computer science and mathematics literature showing that the local structure of finite random network is approximately a random tree (Wormald 1999). Recent papers in the economics literature expand these results and apply them to economic models (Fainmesser and Goldberg 2012, Campbell 2010).

the Huaraz network is about as expansive as the three-dimensional "3D-cube" of approximately equal size which we have included in Figure 4A as well. It follows that the Huaraz network is less expansive than the tree-like random network, and hence our model predicts that full insurance is not coalition-proof.

The result is the same if we look at the two sub-network of relatives and non-relative friends, respectively, in Figure 4B: the non-relative network is slightly more expansive, but does not approach the expansiveness of the random network.

Figure 4 suggests that the expansion properties of the Huaraz network are similar to—somewhat better than—the plane. A plausible reason is that the Huaraz network, like many social networks in practice, is partly organized on the basis of geographic distance. For example, the average distance between two connected agents in this network is only 42 meters, while the average distance between two randomly selected addresses is 132 meters. This correlation between distance and network connections can result in expansion properties similar to the plane, if agents tend to have friends at close physical distance in multiple directions, e.g., both horizontally and vertically on a map. This logic suggests that to understand partial insurance in real world networks, exploring plane-like networks is a useful first step.

2.2 Partial risk-sharing in less expansive networks

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 where the coalitional constraint binds under equal sharing become rare. To see the logic, consider again the regular plane with the i.i.d. $+\sigma/-\sigma$ shocks. As we have seen, equal sharing fails because households in a large n by n square F would need to give up $n^2 \cdot \sigma$ resources if all of them get a positive 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 $n\sigma$, 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.[15] Thus the "typical shock" in F has the same order of magnitude as the maximum pledgeable amount, and hence potentially deviating coalitions are rare. The same logic works with correlated shocks, as long as correlation declines fast enough with distance. By way of contrast, the argument breaks down for the line, since the perimeter of even large interval sets is only $2c$, a constant.

2.3 Plane and line networks

Our intuitive analysis suggests that when shocks are not too correlated, risk-sharing on the plane should be reasonably good, and substantially better than on the line. We first formalize these ideas and then extend them to less regular networks.

Partial risk-sharing measure. We measure partial risk-sharing as the average utility loss relative to the benchmark of equal sharing where all agents consume the average endowment $\overline{e} = e_W / |W|$:

$$UDISP\left(x\right) = \mathrm{E}\frac{1}{|W|}\sum_{i\in W}\left\{U_i\left(\overline{e}\right) - U_i\left(x_i\right)\right\}.$$

This "utility-based dispersion," is simply the difference between average utility under partial and full sharing. Here we ignore the dependence of utility on link consumption to simplify notation.

If all agents have the same quadratic utility function over x, then we can express $UDISP$ as an increasing function of

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

which is the square-root of the expected cross-sectional variance of x. For non-quadratic utilities, $SDISP\left(x\right)$ can be interpreted as a second order approximation of the utility based measure. $SDISP$ is a tractable measure that inherits the intuitive properties of $UDISP$: it is zero only under equal sharing and positive otherwise, and its magnitude measures the departure from equal sharing: e.g., if e_i are $+\sigma/-\sigma$ with equal probabilities, then in autarky

[15] The sum of n^2 i.i.d. random variables has variance $n^2\sigma^2$ and hence standard deviation $n\sigma$.

$SDISP\,(e) = \sigma$. We use $SDISP$ as our central measure in the analysis below.[16]

Shocks with limited correlation. While we focused on i.i.d. symmetric shocks in our example, the formal result accommodates much more general endowment shocks. The key requirements are that shocks do not have fat tails and are not too correlated; we formalize these using assumptions (P1) to (P5) below.

We model the source of uncertainty as a collection of independent random variables y_j, $j = 1, ..., \infty$, which can represent both idiosyncratic shocks like illness and aggregate shocks like weather. Like in a factor model, endowments are determined as linear functions of these basic shocks: $e_i = \sum_j \alpha_{ij} y_j$.where α_{ij} measures the extent to which agent i is exposed to shock j. We assume that e_i and y_j satisfy the following.[17]

(P1) [Thin tails] y_j are independent, have zero mean and unit variance, and satisfy that there exists $K > 0$ such that $\log[\mathrm{E}(\exp[\theta y_j])] \leq K\theta^2/2$ for all $\theta > 0$.

(P2) [Bounded variance] There exists $K > 0$ such that $\sum_j \alpha_{ij}^2 < K$ for all i.

(P3) [Limited correlation] Endowments satisfy $\sigma_F / |F| \leq K \cdot |F|^{-1/2}$ for some $K > 0$, where σ_F is the standard deviation of e_F.

(P4) [More people have more risk] For all $G \subseteq F$, we have $\sigma_G \leq \sigma_F$.

(P5) [Sharing with more people is always good.] For all $G \subseteq F$, we have $\sigma_F / |F| \leq \sigma_G / |G|$.

Here (P1) is a uniform bound on the moment-generating function of y_j, which allows us to use the theory of large deviations to bound the tails of e_i. (P1) is satisfied for example if y_j are i.i.d. normal, or if they have a common compact support. Property (P3) requires that shocks are not too correlated, so that aggregate uncertainty disappears at the same rate as the square root of set size. This condition considerably relaxes the i.i.d. assumption; for example, on the line or plane, (P3) is satisfied if the correlation between e_i decays geometrically with network distance.

Formal results. We now turn to a formal result on risk-sharing on the plane and line networks. Although the formal result assumes that all links have equal capacities c, it would

[16] Equation (2) only defines $SDISP$ for finite networks. For infinite networks, we define it to be the lim sup of (2), taken over an increasing sequence of ball sets centered around some agent i. For the line and the plane, the choice of i does not affect this lim sup.

[17] From now on we use the convention that K denotes a positive constant, the value of which at each occurance of the phrase "there exists K" may be different; and that the same holds for K' and for K''.

continue to hold—with different constants—if all link capacities are from a bounded interval $[c/k, ck]$ for some $k > 0$. We focus on infinite networks because they are more convenient for stating our asymptotic result.

Proposition 2 *Under properties (P1)-(P5), 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 proof of (i) essentially builds on our earlier arguments: for long enough intervals, much of the interval-specific shock must remain trapped in the set, because the perimeter is only $2c$. Even if agents perfectly smooth inside the interval, overall dispersion remains high.

The result for the plane is much more difficult, and requires going beyond our previous intuition: even though the coalitional constraint is rarely violated for any *particular* set F, we need an allocation that satisfies the constraints of *all* sets. Equivalently, we need to construct a transfer arrangement such that the typical flow on every link meets the capacity constraint. The key idea is to construct this arrangement from the ground up. First we partition the plane into 2 by 2 squares of agents and implement equal sharing in each of these. Then we implement fully sharing in 4 by 4 squares, then in 8 by 8 ones, and so on. After n iterations, we obtain full sharing of endowments in 2^n by 2^n "super-squares". Because each link is used once in every round, the construction uses every link at most n times. By our earlier intuition, each time a link is used, the required transfer is typically of order one, resulting in a total flow per link of order n. This is the uniform bound on the flow over every link that we require for exponentially good risk-sharing. Since the arrangement does not yet account for capacity constraints, we use the theory of large deviations to bound the exceptional event when incentive compatibility is violated, obtaining the bound in the proposition.

Simulations. Numerical simulations suggest that the asymptotic results of the Proposition provide a good description of behavior for finite c as well. Figure 5 shows constrained optimal allocations for finite line and plane networks, for a typical realization of uniform shocks with support $[-1, 1]$.[18] Figure 5A shows the endowment realizations for both the line and the plane network: darker red (green) squares correspond to lower (higher) endowments. We use the same vector of realizations for both networks. The $SDISP$ of these realizations is 0.55 in the absence of any insurance. Now consider Figure 5B, where we assume that the average capacity per agent is 1: thus each link has value $c = 0.5$ in the line network and $c = 0.25$ in the plane. For these capacities, the figure depicts the optimal, $SDISP$ minimizing incentive compatible allocation. The contrast between the line and the plane is remarkable: for the line, we see substantial color variation reflecting imperfect risk-sharing ($SDISP = 24\%$), while the plane achieves better insurance ($SDISP = 12\%$). As capacities increase, the contrast becomes sharper. In Figure 5C, the per capita capacity in both networks is assumed to be 1.4, $SDISP$ on the line is still 20%, while on the plane it falls to 3%. Finally, in Figure 5D, when the per capita capacity is 2, dispersion on the line falls to 14% while full risk-sharing is achieved on the plane ($SDISP = 0$). We conclude that the asymptotic results of the Proposition provide a good characterization of insurance behavior in finite networks and for finite c as well.

2.4 Geographic networks

If real world networks are similar to the plane, Proposition 2 suggests that they should allow for reasonably good risk-sharing. Many papers, in various contexts, show that geographic proximity is a major determinant of interpersonal relationships (see for example Conley and Udry (2010), Fafchamps and Gubert (2007) in development contexts, and Lee, Mancini and Maxwell (1998,1995), Topa (2001) in other contexts). This motivates our investigation below

[18] In the simulations opposing edges of the networks are connected, so the line is in fact a circle and the plane a torus.

Figure 5: Risk-sharing simulations on the line and the plane for increasing capacities

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

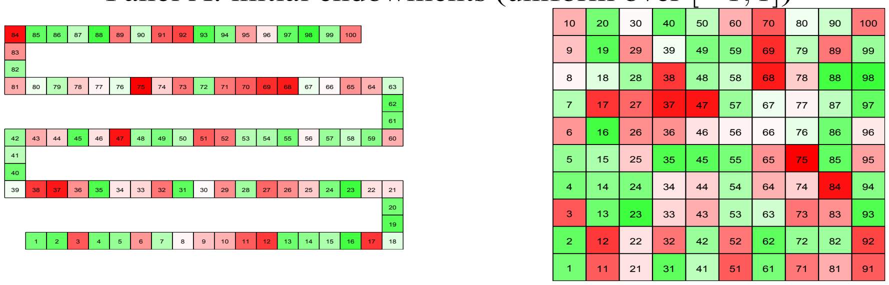

SDISP= 0.556 SDISP= 0.556

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

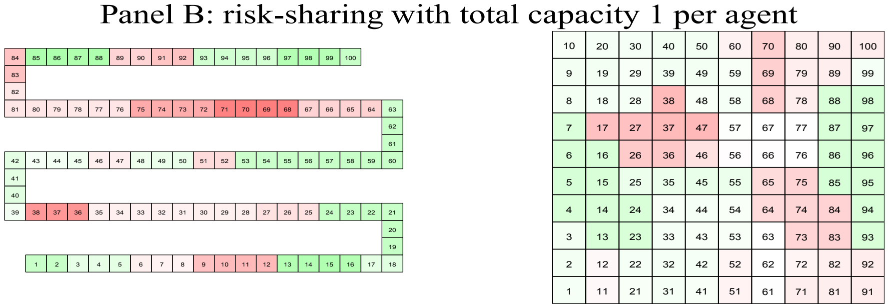

SDISP= 0.246 SDISP= 0.127
30 islands 17 islands

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

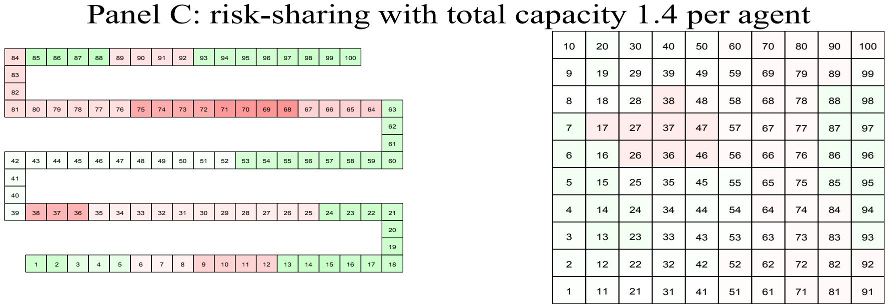

SDISP= 0.199 SDISP= 0.035
17 islands 4 islands

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

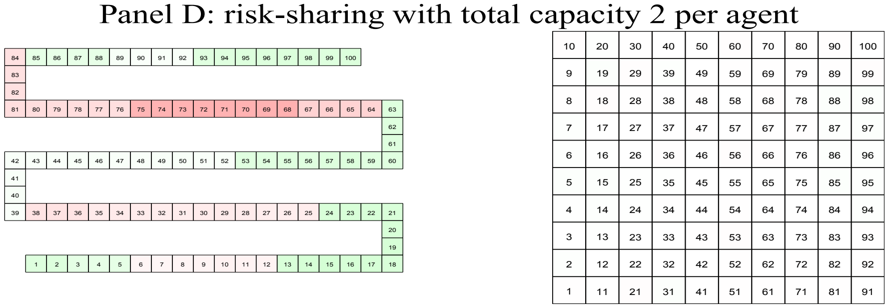

SDISP= 0.148 SDISP= 0
13 islands 1 island

to define plane-like networks in a spatial context.

As Figure 1 illustrates, real-world social networks have a much less regular structure than the plane. Nevertheless, these networks can often be represented in a way that closely resembles a regular plane, because in the physical map of the community, households tend to have social connections at close distances and in multiple directions. Intuitively, if a sufficiently accurate representation of this sort does exist, then our results on good risk-sharing are likely to carry over to real world social networks.

To formally define what makes a representation "sufficiently accurate," we consider (1) a function $\pi : W \to \mathbb{R}^2$ that maps agents in a social network to locations in $\mathbb{R}^2$; and (2) a two dimensional grid that divides $\mathbb{R}^2$ into squares of side length A. This pair constitutes an *even* representation if the number of households inside each grid cell is uniformly bounded by positive constants from below and above. The representation is *local* if geographically close agents are connected through a path that is also geographically close: for any $d > 0$ and i and j at geographic distance $d(\pi(i), \pi(j)) \leq d$, there is a path connecting i and j such that for all agents h in the path, $d(\pi(i), \pi(h))$ is bounded from above by a constant that only depends on d. Finally, the representation exhibits *no separating avenues* if the sum of capacities of links between any two neighboring squares is uniformly bounded away from zero; this is the key condition that guarantees plane-like expansion properties.

A network is called a *geographic network* if it has a representation that is even, local, and has no separating avenues, and all link capacities are bounded away from zero.[19]

Corollary 1 *In a geographic network, if (P1)-(P5) is satisfied, then there exist positive constants K' and K'' such that $SDISP(x) \leq K' \exp\left[-K'' c^{2/3}\right]$ for some incentive-compatible risk-sharing arrangement.*

Thus the risk-sharing properties of geographic networks are similar to the plane. The proof combines Proposition 2 with a renormalization argument. We take a geographic network, and superimpose on its planar representation a grid with A by A squares. We then

[19]A geographic network is by assumption infinite; we define $SDISP$ for these networks as the lim sup of (2) over a sequence of increasing squares in the map representation. The exact sequence does not matter for the results.

merge all people within each square to create a new network. Because of the key no separating avenues condition, this new network is essentially a plane, and hence Proposition 2 (ii) can be applied to yield a bound for $SDISP$ in the new network. We then pull this bound back to the old network using the fact that the embedding is even and local.

Geographic networks in practice. Because real-world networks are finite, they cannot satisfy the conditions required for geographic networks, which are by definition infinite. Nevertheless, it is possible to evaluate whether concrete finite networks share some of the features required for geographic networks. Here, we develop an embedding to show that the Huaraz network gets close to satisfying the key conditions of evenness and no separating avenues, suggesting that the same properties that generate good risk-sharing for geographic networks are also at work in the Huaraz case. Figure 6A shows the natural geographic map of household locations, referred to as lots, in this village. In Figure 6B the horizontal and vertical coordinates of the map are re-scaled to fit the community into the unit square, and a grid of 16 squares is also depicted. As is clear from Figure 6B, this representation is unlikely to satisfy the geographic networks condition, because there are empty squares and the distribution of agents is quite heterogeneous. To construct a "geographic" representation of this Huaraz community, we transform the map using a diffusion algorithm described in detail in the supplementary appendix. The basic idea is to stretch the network uniformly over the unit square using a procedure in which nearby lots "repel" each other and hence lots will tend to escape to empty spaces. Figures 6C and 6D depict the result after one and five rounds of iteration: the distribution of lots becomes gradually more homogenous. After 23 iterations (Figure 6E), the distribution of lots is almost completely uniform. Figure 6E also shows the number of lots in each of the 16 squares, confirming that we have an even embedding.

To evaluate the key "no separating avenues" condition, Figure 6E also shows the number of links crossing the sides of each square.[20] The agreement with our theoretical condition is very good: except for one side of the square in the lower right corner, there are no separating avenues between any two neighboring squares. The average number of nodes in each grid cell

[20] Opposing sides of the large square are assumed to be geographically next to each other, generating the topology of a torus.

Figure 6: Stretching a real-world network to construct a geographic representation

Panel A: original map of Huaraz community

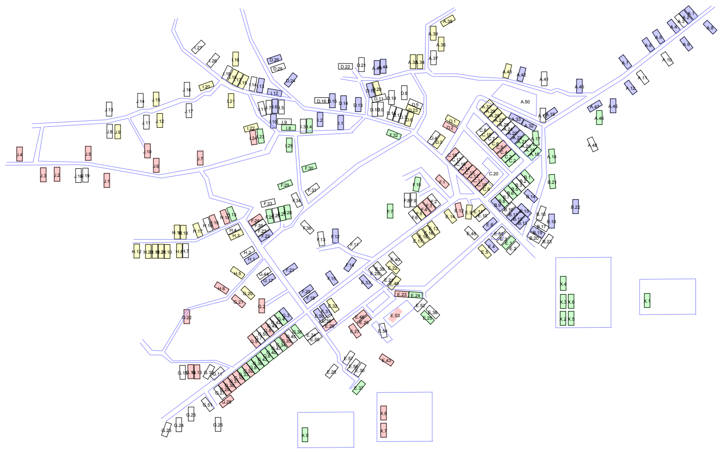

Panel B: iteration 0

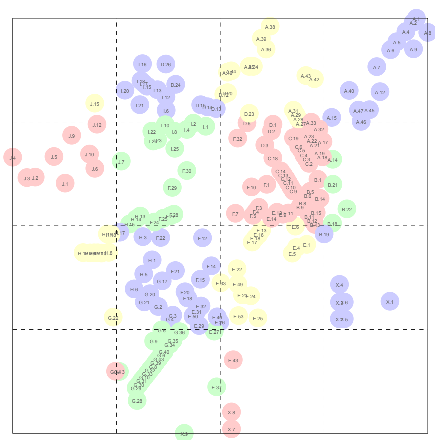

Panel C: iteration 1

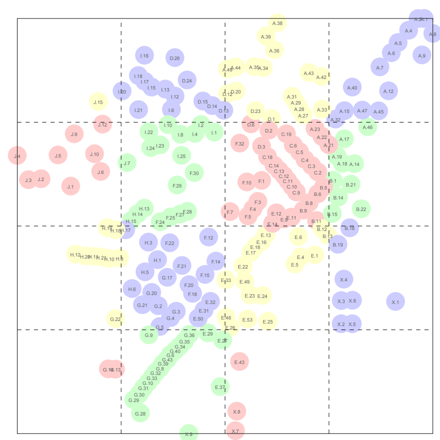

Panel D: iteration 5

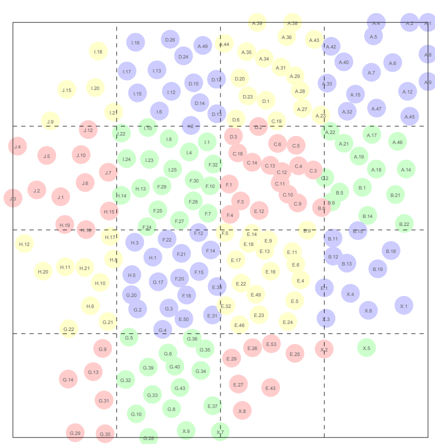

Panel E: iteration 23

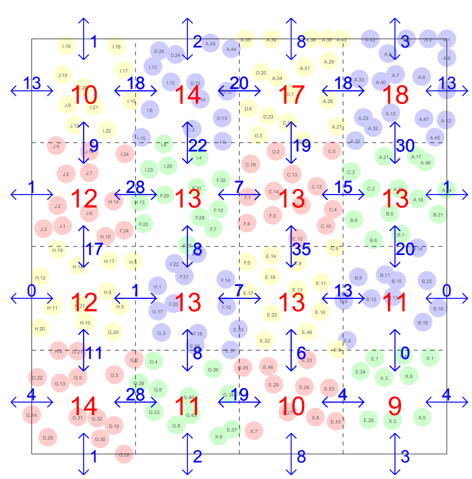

is 12.7 and the number of connections to neighboring squares is 49.4. To better understand what drives the success of this embedding, note that in Figure 6E each of the 16 squares is differently colored, and the corresponding households are represented by the same colors in panels A to D as well. In the original image (Figure 6A), households are geographically concentrated by color; hence the reason why the Huaraz network has similar expansion properties as the plane is that households tend to have friends in multiple directions at close distance in the original map.[21]

Numerical risk-sharing simulations suggest that the Huaraz social network in fact behaves very much like the plane network: we calculate $SDISP$ for uniform shocks with support $[-1, 1]$ and per capita capacities 1, 1.4 and 2. We obtain $SDISP$ equal to 0.20, 0.11 and 0.02, respectively, which tracks the rapid decline of $SDISP$ on the plane. The finding that the Huaraz community resembles a "geographic network", in part because connections are correlated with physical distance suggests that village networks in developing countries may be similarly expansive. Our results then imply that typical village networks should facilitate high, although imperfect, levels of informal risk-sharing – consistent with the empirical findings of Townsend (1994), Ogaki and Zhang (2001), Mazzocco (2007) and others.

3 Constrained efficient risk-sharing

In this section, we study *constrained efficient* arrangements which are Pareto-optimal given the enforcement constraints imposed by the network. Such second-best arrangements are a natural benchmark because they achieve the highest possible level of risk-sharing in a given network. Such arrangements can either be proposed and implemented by a village leader, or attained in ex ante coalitional bargaining, possibly through multiple rounds of renegotiation (see Gomes (2000) and Aghion, Antras and Helpman (2007) that such bargaining procedures lead to efficient agreements). In the Supplementary Appendix we also illustrate how a decentralized sharing procedure between neighboring agents, as in Bramoulle and Kranton

[21] In contrast, when we apply the same diffusion procedure to a finite circle network with the same number of nodes and equivalent average degree, we find that the representation is far from satisfying the no separating avenues condition. In particular, Figure 10 in the Supplementary appendix has many more gaps, especially in the center; and the average number of neighboring square connections is now only 23.0 which is less than half the number of neighboring connections in Figure 6E.

(2006) can achieve any constrained efficient arrangement.

3.1 Risk-sharing islands

Our main result is that constrained-efficient insurance arrangements exhibit an "island structure." For every realization of endowments, connected islands of agents emerge endogenously, such that risk-sharing is perfect within each island, while links between different islands are "blocked" in the sense that transfers equal the link capacities. This result follows from the equivalence between constrained efficient arrangements and a planner's problem formalized below.

The intuition for islands can be seen by focusing on a utilitarian social planner who maximizes average expected utility. Whenever two agents consume different amounts, this planner can increase welfare by shifting a small amount from the agent with higher- to the one with lower consumption. But in the optimum, such shifts must violate the enforcement constraints. Hence linked agents either consume the same amount and belong to the same "island", or consume different amounts and are connected by a blocked link that does not allow for further transfers. Panels B-D of Figure 5 depict constrained efficient allocations corresponding to such a social planner: islands within which consumption is equalized are indicated by different colors.

For a formal analysis, let (λ_i) be a set of positive weights, and define the *planner's problem* as

$$\max_{(t)} \sum_{i \in W} \lambda_i \cdot \mathrm{E}U_i\left(x_i\right) \tag{3}$$

subject to the constraint that all transfers respect the capacity constraints of the social network.

Proposition 3 *Every constrained efficient risk-sharing arrangement is the solution to a planner's problem with some set of weights (λ_i). Conversely, any solution to the planner's problem is constrained efficient.*

The proof of this result parallels a similar equivalence result for risk-sharing in syndicates by Wilson (1968). Because the set of coalition-proof payoff vectors is convex—when two transfers satisfy a capacity constraint, so does their convex combination—efficient allocations, which by definition lie on the boundary of this set, can be supported by tangent hyperplanes. The normal vector (λ_i) associated with the suporting hyperplane gives the appropriate planner's problem.[22,23]

Maximizing the planner's expected utility $\mathrm{E}\sum \lambda_i U_i$ is equivalent to maximizing realized utility $\sum \lambda_i U_i$ independently for each state. This yields a set of intuitive first-order conditions for each realization. To state these conditions, recall that a link from i to j is *blocked* in a given realization if $t_{ij} = c(i, j)$, i.e., if the link is used at full capacity.

Proposition 4 *An incentive-compatible arrangement* (t_{ij}) *is constrained efficient if and only if there exist positive weights* $(\lambda_i)_{i \in W}$ *such that for every* $i, j \in W$ *one of the following conditions hold:*

1) $\lambda_i U_i'(x_i) = \lambda_j U_j'(x_j)$

2) $\lambda_i U_i'(x_i) > \lambda_j U_j'(x_j)$ *and the link from* j *to* i *is blocked*

3) $\lambda_i U_i'(x_i) < \lambda_j U_j'(x_j)$ *and the link from* i *to* j *is blocked.*

This result generalizes our earlier intuition for arbitrary welfare weights. Sufficiency and uniqueness of the first-order conditions follow from the strict concavity of the planner's objective function and the convexity of the domain. The Proposition also implies that for any pair of agents i and j, if $\lambda_i U_i' < \lambda_j U_j'$, then along every all path connecting i and j, at least one link must be blocked. Therefore, in any realization agents can be partitioned into connected risk-sharing islands such that within an island agents share risk perfectly, while cross-island insurance is limited because boundary links operate at full capacity.

Proposition 5 *[Risk-sharing islands] In any realization* e *the set of agents can be partitioned into connected components* W_k *such that* $\lambda_i U_i' = \lambda_j U_j'$ *if* $i, j \in W_k$*, and* $|t_{ij}| = c(i, j)$ *if* $i \in W_k$*,* $j \notin W_k$*.*

Sharing islands partition the network in each realization. Using the coalitional interpretation, these islands can be thought of in terms of "almost-deviating coalitions." For

[22] See the Supplementary Appendix for extending this result to imperfect substitutes.

[23] All simulations in Section 2 compute the constrained-efficient arrangement with equal λ weights under quadratic utility.

example, if all links on the boundary of an island are blocked in the outward direction, then members of this are transferring the highest amount they can credibly pledge to the community, and hence are indifferent to deviating as a coalition. More generally, it can be shown that the island decomposition obtains by splitting the network along the boundaries of all almost-deviating coalitions. In effect, almost deviating coalitions act as "bottleneck groups" limiting the flow of resources in a way parallel to the bottleneck agents emphasized in Bloch et al. (2008). The emergence of network-based risk-pooling islands is consistent with evidence documented by Attanasio et al. (2009) about the importance of social ties in the formation of insurance groups in Colombian villages.

When link capacities increase, the planner becomes less constrained and risk-sharing islands tend to grow in size. This is illustrated by Figure 5, panels B to D. In Figure 5B, where per capita capacity is one, insurance is fairly local: there are 30 islands on the line and 17 on the plane. As the per capita capacity goes up to 1.4, in Figure 5C there are 17 islands on the line and only 4 on the plane; and in Figure 5D where average capacity is 2 per agent, there are 13 islands on the line and just one, fully insured island on the plane. In these simulations, the number of islands closely tracks the degree of insurance.

As is clear from Figure 5, in the island partition the size and location of islands, and hence the set of agents who fully share each others' shocks, is endogenous to the realization and the network. This result differentiates our model from group-based models of risk-sharing, where insurance groups are exogenous and do not vary with the realization.

3.2 Spillover effects and local sharing

The island result also helps us characterize how shocks propagate in the network as a function of social distance. We show that shocks are shared to a greater degree with socially close agents, and hence network-based insurance is *local*: the consumption of socially close agents comoves more strongly than that of socially distant ones.

To formalize this point, we introduce a slightly stronger definition of risk-sharing islands. Fix an endowment realization (e_i), and let $W(i)$ denote the sharing island containing i. We now define $\widehat{W}(i)$ to be the maximal connected set of agents j such that there exists a path between i and j along which no links are blocked in either direction. With this definition,

$\widehat{W}(i) \subset W(i)$ because Proposition 5 implies that links connecting different islands are all blocked. Except for knife-edge cases when the transfer constraint is reached but does not bind yet—which have zero probability when the distribution of shocks is absolutely continuous—the two definitions are equivalent: $\widehat{W}(i) = W(i)$.

We now explore the effects of an idiosyncratic shock to one agent's endowment on the consumption of others. Fix a constrained efficient arrangement, and consider two realizations $e = (e_i)$ and $e' = (e'_i)$, where $e'_i < e_i$ for some i but $e'_j = e_j$ for all others $j \neq i$. Effectively, agent i is experiencing an idiosyncratic negative shock in e' relative to e (or a positive shock like aid in e relative to e'). We can measure the impact of this negative shock on another agent j by computing the ratio of marginal utilities of j before and after the shock. Formally, let x and x' denote the consumption vectors associated with e and e', then we can define

$$MUC_j = \frac{U'_j\left(x'\right)}{U'_j\left(x\right)}$$

which measures the marginal utility cost of the shock for agent j. A larger MUC_j corresponds to a higher increase in marginal utility and hence a greater consumption drop.

Proposition 6 *[Spillovers and local sharing] Consider two realizations $e = (e_i)$ and $e' = (e'_i)$, where $e'_i < e_i$ for some i but $e'_j = e_j$ for all $j \neq i$. Then in any second best arrangement x:*

(i) [Monotonicity] $x_j(e') \leq x_j(e)$ for all j, and if $j \in \widehat{W}(i)$ then $x_j(e') < x_j(e)$.

(ii) [Local sharing] There exists $\Delta > 0$ such that $|e_i - e'_i| < \Delta$ implies $MUC_i = MUC_j$ for all $j \in \widehat{W}(i)$, and $x_j\left(e'\right) = x_j\left(e\right)$ for all $j \in W \backslash W(i)$.

(iii) [More sharing with close friends] For any $j \neq i$, there exists a path $i \rightarrow j$ such that for any agent l along the path, $MUC_l \geq MUC_j$.

Part (i) shows that spillovers are monotone: If one agent receives a negative shock, the consumption of everybody else either decreases or remains constant. Moreover, the agent is partially insured by all others in the same risk-sharing island, who all reduce their consumption by a positive amount. Thus unless i is in a singleton island, he has access to at least some insurance. Intuitively, links within $\widehat{W}\left(i\right)$ are not blocked, and hence all members of the island can help out a little. As part (ii) shows, for small shocks, the set of agents who insure i is exactly $\widehat{W}(i)$. All these agents share an equal burden measured in terms of the marginal utility cost MUC. Agents outside of $W(i)$ do not reduce their consumption at all.[24] Finally, (iii) shows how the utility cost of agents varies by social distance. Indirect friends provide less insurance to i than direct friends: for any agent $j \neq i$, there exists some direct friend of i, denoted l, who shares at least as much of the burden of the shock as j does.

The results of Proposition 6 are consistent with the empirical findings in Angelucci and De Giorgi (2009), who show that Progresa, a conditional cash transfer program in rural Mexico, leads to an increase in the consumption of the non-treated, which they attribute to the spillover effect of aid through the social network of the village. This is the logic of part (i) in the Proposition. Angelucci et al. (2012) also show that much of the increase in the consumption of the non-treated is due to the consumption increase of households who are relatives of the treated, consistent with (ii) and (iii). The agreement between our results and existing evidence suggests that a calibrating our model may be useful for quantifying the welfare effects of development aid taking into account network-based spillovers.

4 Endogenous Link Strength and Stability

This section presents an extension of our basic model in which the strength of social connections is endogenously determined. The preceding sections, by assuming that capacities are determined outside the model, take the view that link strength depends primarily on benefits of socialization which are unrelated to informal insurance. We now consider the alternative that agents choose their level of socialization to obtain better informal risk-sharing. In this context, we use a very simple model to explore whether the difference in insurance outcomes between the line and plane networks is reduced, because people in the plane choose to socialize relatively less, or amplified, because people in the plane choose to socialize relatively more. We leave a fuller analysis of insurance with endogenous link strength for future research.

[24] In the knife edge case where $\widehat{W}(i) \neq W(i)$, agents in $W(i) \setminus \widehat{W}(i)$ may or may not share.

Setup. We consider an exogenous network which is symmetric in the sense that for any pair of agents i and j there exists an automorphism of the network $b(.)$ such that $b(i) = j$.[25] We assume that, before shocks are realized, each agent chooses effort a_i to socialize with her set of neighbors N_i. Effort is spread equally across all links of the agent, and, denoting the degree of agents by d, for and given the vector of efforts $a = (a_i)$ capacities are determined as

$$c(i, j|a) = \min\left(\frac{a_i}{d}, \frac{a_j}{d}\right). \tag{4}$$

We assume that agent i's incentives to socialize are determined by the utility function $EU_i(x_i - \widehat{c}_i) - \alpha \cdot a_i$, where α captures the marginal cost of socializing, and $\widehat{c}_i = c(i, j|a)$ if the agent defects on an obligation with j, and zero otherwise. Slightly differently from the previous sections, this formulation assumes that link capacities enter utility not as positive, but as potentially negative terms, which are activated by a deviation. This specification, by removing the direct utility effect of increased socialization, allows us to isolate the insurance-based incentive to invest in social links. Allowing link capacities to enter positively would introduce a non-insurance-based motive to socialize.

We call the pair (a, t) a symmetric feasible social arrangement if t is an incentive-compatible risk-sharing arrangement when capacities are given by $c(i, j|a)$, i.e., if each agent chooses socialization level a. We think of the pair (a, t) as a social norm which specifies a suggested level of socialization and a suggested risk-sharing arrangement for society; and from now on we focus on the case in which t is the equal-weighted constrained-efficient arrangement given capacities $c(i, j|a)$.

We are interested in social norms that are stable with respect to individual deviations in socialization. To define stability, we first need to specify what happens when an agent chooses $\widetilde{a}_i \neq a$. Equation (4) immediately implies that no agent would want to set $\widetilde{a}_i > a$. When i sets $\widetilde{a}_i < a$, we assume that in the resulting new network, required transfers are

[25] An automorphism b is a bijection $b : W \to W$ such that $(u, v) \in L$ if and only if $(b(u), b(v)) \in L$. For example, the circle or torus satisfy this criterion.

Figure 7: Maximal supported total capacity c per agent for line and plane - binary shocks (left panel) and uniform shocks (right panel)

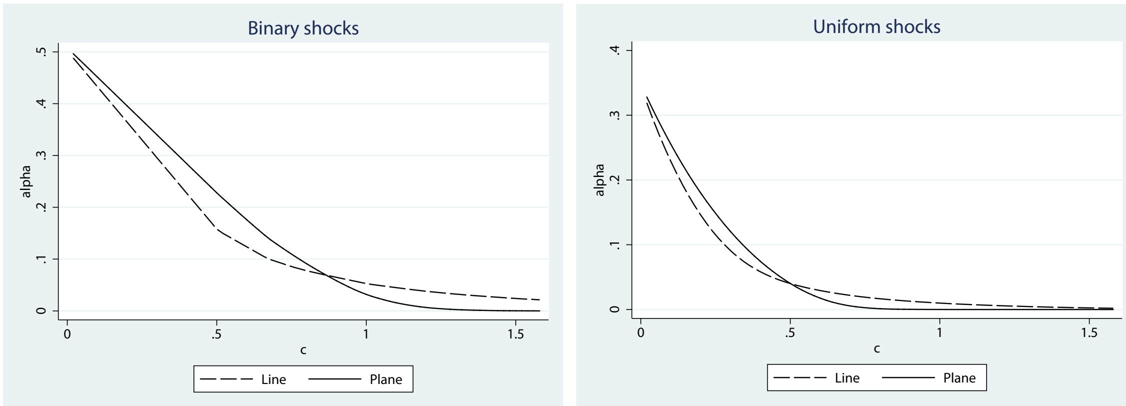

specified by the *truncated risk-sharing arrangement* $\widetilde{t}^i$ defined as

$$\widetilde{t}_{ij}^{i,e} = \begin{cases} \min\left(t_{ij}^e, c\left(i,j|\widetilde{a}_i, a_{-i}\right)\right) & \text{if } t_{ij}^e > 0 \\ -\min\left(-t_{ij}^e, c\left(i,j|\widetilde{a}_i, a_{-i}\right)\right) & \text{otherwise.} \end{cases}$$

In words, in the new network in which the links of i have lower capacity, the previously specified transfers between i and a connection j take place fully if they meet the new capacity constraint, but take place only partially—up to the new constraint—otherwise. Thus, truncation captures the notion that the new social structure can only support transfers t^e up to the point at which they are also incentive-compatible in the modified network. Given a distribution of endowment shocks, we call the social arrangement (a, t) *stable* if no agent can increase her expected utility by changing her socialization effort.

Analyis. It is easy to verify that each agent's expected utility is increasing and strictly concave in her effort level $\widetilde{a}_i$, provided that $\widetilde{a}_i < a$,[26] and hence the following necessary and sufficient first-order condition characterizes symmetric stable arrangements:

$$\left.\frac{\partial\_EU_i\left(\widetilde{x}_i\right)}{\partial \widetilde{a}_i}\right|_{\tilde{a}_i = a_i} \geq \alpha \quad \text{for all agents } i.$$

The left-side derivative $\partial\_EU_i\left(\widetilde{x}_i\right)/\partial\widetilde{a}_i$ represents the marginal utility loss to agent i if she slightly reduces her socialization effort. Stability requires that this utility loss is not smaller than the utility gain from having to spend less on the socialization effort.

[26] Contact the authors for a formal proof.

Figure 8: Maximal supported SDISP for line and plane - binary shocks (left panel) and uniform shocks (right panel)

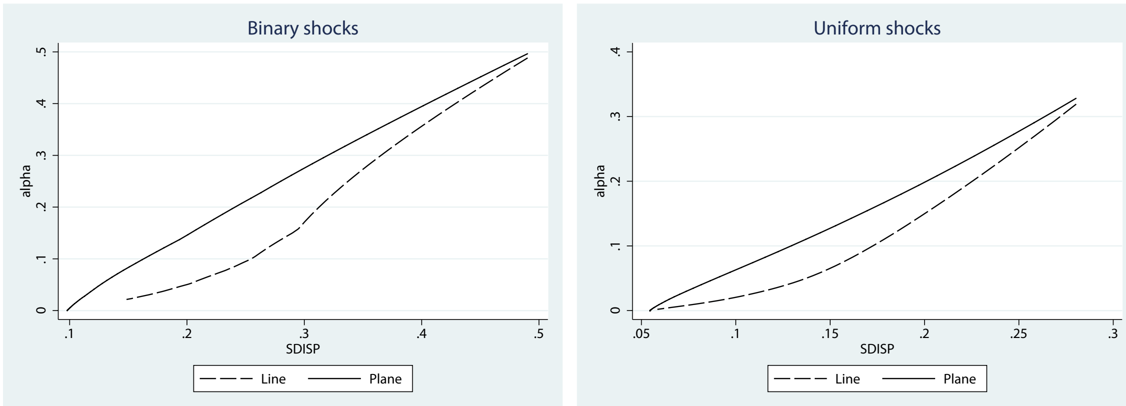

We now turn to use this model to explore how our conclusions about the line and the plane are affected with endogenous link strength. Specifically, we are interested in the highest stable socialization effort that can be supported in each network as we vary α. To begin, we numerically solve for the equilibrium for both binary and uniform shock distributions and plot, in Figure 7, the maximum stable per-capita link capacity $\overline{c}$ for a range of values of the marginal cost of socialization α. The lesson from the Figure is that for large and intermediate α the plane provides more incentives to socialize than the line, while this ordering is reversed for small α. Thus, in the range of α where insurance is not yet close to perfect, our basic conclusion that risk-sharing is better on the plane is amplified. As the plane reaches close to full insurance sooner than the line the relationship is eventually reversed as the marginal benefit of insurance decreases more quickly on the plane. However, risk sharing is always better on the plane compared to the line for all values of α as figure 8 demonstrates.

A partial intuition for how the incentives to invest vary with α comes from noting that an agent is affected by a marginal reduction in his investment only when he is on the perimeter of a risk-sharing island—because otherwise the truncation does not bind. In turn, the frequency with which he ends up on such a perimeter is related to the average perimeter-area ratio of sharing islands. In particular, when—as in the plane for α relatively high—that ratio is large, agents are more frequently on the perimeter, and hence the incentives to invest are strong, generating relatively higher incentives to socialize.

This intuition is only partial, because the direction of flows on the boundary of a sharing island also matters, and, in general, the same island can have both inflows and outflows along its boundary. To clarify this point, let $P\left(k, r^{in}, r^{out}\right)$ denote the probability that the agent is in a sharing island of size k such that its perimeter has r^{in} links receiving transfers and r^{out} links sending transfers.[27] Denote by $\overline{U}'(x|k, r^{in}, r^{out})$ the mean marginal utility of consumption of agents across all $\left(k, r^{in}, r^{out}\right)$ islands and across all realizations under risk-sharing arrangement x. Then we can write

$$\left.\frac{\partial\_EU_i\left(\widetilde{x}_i\right)}{\partial \widetilde{a}_i}\right|_{\tilde{a}_i=a_i} = \sum_{k,r^{in},r^{out}} P\left(k, r^{in}, r^{out}\right) \overline{U}'(x|k, r^{in}, r^{out}) \frac{r^{in}-r^{out}}{kd}. \tag{5}$$

The logic behind this formula is the following. Reducing socialization affects i' utility only in those realizations in which he is on the boundary of a risk-sharing island. Because the network is symmetric, for i.i.d. shocks an agent is equally likely to take any of the k positions inside the risk-sharing island. Therefore, conditional on i being in this island, and denoting his per-link capacity by c, the expected amount of resources which flow to him from the outside equals $r^{in}c/k$, and the expected flow to the outside originating from him equals $r^{out}c/k$. Given that for $\widetilde{a}_i < a$ the per-link capacity is $c = \widetilde{a}_i/d$, the derivative of these quantities with respect to $\widetilde{a}_i$ gives the last term in the expression. These consumption effects are weighted by probabilities and by the marginal utility of consumption of agents, $\overline{U}'(x|k, r^{in}, r^{out})$.

Equation (5) links the incentives to socialize with the variable $\frac{r^{in}-r^{out}}{kd}$ which we call the normalized net flow. This random variable, although closely related, differs from the perimeter-area ratio (which equals $a\left[F\right] = c\frac{r^{in}+r^{out}}{kd}$) in that it also takes into account the direction of flows on the boundary of the island. In particular, the net flow can be either positive or negative. What matters for the incentives to socialize, according to (5), is the *dispersion* in the net flow: because positive normalized net flows tend to be associated with low consumption levels inside the sharing island and therefore large weights $\overline{U}'(x|k, r^{in}, r^{out})$ while the converse is true for negative net flows, higher net flow dispersion results in stronger incentives to invest in links.

The variability in the net flow is closely related to the average perimeter-area ratio. This is easiest to see in an environment with symmetric binary shocks and small capacities, in

[27] In particular, the total perimeter of the island is $\left(\overline{c}/d\right)\left(r^{in}+r^{out}\right)$.

Figure 9: Variance of distribution of net-flow-area ratio $\frac{r^{in}-r^{out}}{kd}$ for line and plane - binary shocks (left panel) and uniform shocks (right panel)

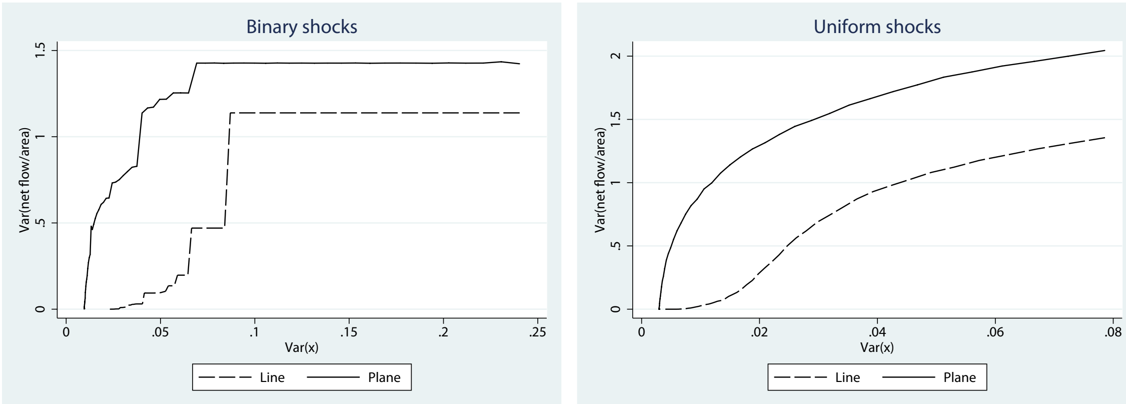

which each sharing island either has all links pointing in, or all links pointing out. In that environment, in equation (5) with probability one either $r^{in} = 0$ or $r^{out} = 0$, which implies that the expression simplifies to a weighted sum of the perimeter-area ratios of sharing islands. But even more generally, we show through simulations that geometries with high perimeter-area ratios (such as the plane or the binary tree) also have highly dispersed net-flow-area distributions (and hence high marginal utility from socialization). Figure 9 shows the variance of the net-flow-area ratio distribution for the line and plane while controlling for the variance of consumption x (i.e. controlling for the degree of risk-sharing achieved by the network).[28] For both binary and uniform shocks the net-flow-area ratio distribution always has a higher variance on the plane compared to the line. This translates into a larger marginal incentive to invest in socialization on the plane (for the same degree of risk-sharing): the very feature that creates good risk-sharing on the plane also makes these risk-sharing arrangements stable.

5 Conclusion

This paper showed that the expansiveness of a social network determines the effectiveness of informal risk-sharing. Our results provide an explanation for why many real-life social networks are likely to be sufficiently expansive to allow for good risk-sharing. We also

[28] Note, that this distribution has mean 0 and is symmetric around the mean for binary and uniform shocks.

characterized Pareto-optimal arrangements and found that resources are shared among local groups.

In future work we plan to develop a dynamic version of our model, in which the value of a social link is partly derived from the present value of future insurance benefits in the network. In such a model the values of social links, the network structure, and the risk-sharing agreement would all be endogenized.

We also plan to extend our empirical analysis. Our model is sufficiently tractable that it can be used to estimate the strength of different types of links from social network and consumption data. Such estimates could be used for policy experiments, such as (i) measuring the welfare effects of development aid, taking into account network spillovers; or (ii) comparing the network structure of communities with different degrees of ethnic heterogeneity, and exploring the implications for informal insurance.

References

Aghion, Philippe, Pol Antras, and Elhanan Helpman, "Negotiating Free Trade," *Journal of International Economics*, September 2007, *73* (1), 1–30.

Ali, Nageeb and David Miller, "Enforcing Cooperation in Networked Societies," *Working paper, UC San Diego*, 2008.

Ambrus, Attila, Markus Mobius, and Adam Szeidl, "Supplementary Material to: "Consumption Risk-sharing in Social Networks"," available online at `http://www.socialcollateral.org/risksharing/supplementary_appendix.pdf` November 2012.

Angelucci, Manuela and Giacomo De Giorgi, "Indirect Effects of an Aid Program: How do Cash Injections Affect Non-Eligibles' Consumption?," *American Economic Review*, March 2009, *99* (1), 486–508.

___ **, Giacomo De Giorgi, and Imran Rasul**, "Resource sharing witin family networks: insurance and investment," Working paper, University of Michigan 2012.

Attanasio, Orazio, Abigail Barr, Juan Camilo Cardenasy, Garance Genicot, and Costas Meghir, "Risk Pooling, Risk Preferences, and Social Networks," working paper, UCL Department of Economics August 2009.

Bernheim, B. Douglas, Bezalel Peleg, and Michael D. Whinston, "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory*, 1987, pp. 1–12.

Bloch, Francis, Garance Genicot, and Debraj Ray, "Informal Insurance in Social Networks," *Journal of Economic Theory*, 2008, *143* (1), 36–58.

Bramoulle, Yann and Rachel E. Kranton, "Risk-Sharing Networks," *Journal of Economic Behavior and Organization*, 2006, *forthcoming.*

Campbell, Arthur, "Tell Your Friends! Word of Mouth and Percolation in Social Networks," working paper, Yale University September 2010.

Coate, S. and M. Ravaillon, "Reciprocity Without Commitment: Characterization and Performance of Informal Insurance Arrangements," *Journal of Development Economics*, 1993, *40*, 1–24.

Cochrane, John H., "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 1991, *99* (5), 957–976.

Conley, Timothy G. and Christopher R. Udry, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, March 2010, *100* (1), 35–69.

Cormen, Thomas H., Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein, *Introduction to Algorithms*, Cambridge, MA: MIT Press, 2001.

De Weerdt, Joachim and Stefan Dercon, "Risk-sharing networks and insurance against illness," *Journal of Development Economics*, December 2006, *81* (2), 337–356.

Dixit, Avinash, "Trade Expansion and Contract Enforcement," *Journal of Political Economy*, 2003, *111*, 1293–1317.

Edmonds, Jack and Richard M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *Journal of the ACM*, 1972, *19* (2), 248ï£¡264.

Ellison, Glenn, "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 1994, *61*, 567–588.

Fafchamps, Marcel and Flore Gubert, "Risk Sharing and Network Formation," *American Economic Review Papers and Proceedings*, 2007, *97* (2).

___ **and Susan Lund**, "Risk-Sharing Networks in Rural Philippines," *Journal of Development Economics*, 2003, *71*, 261 – 287.

Fainmesser, Itay P. and David A. Goldberg, "Cooperation in Partly Observable Networked Markets," working paper, Brown University July 2012.

Fehr, Ernst and Simon Gachter, "Cooperation and Punishment in Public Goods Experiments," *American Economic Review*, September 2000, *90* (4), 980–994.

Genicot, Garance and Debraj Ray, "Group Formation in Risk-Sharing Arrangements," *Review of Economic Studies*, 2003, *70*, 87–113.

Golub, Benjamin and Matthew O. Jackson, "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *The Quarterly Journal of Economics*, 2012.

Gomes, Armando R., "Externalities and Renegotiations in Three-Player Coalitional Bargaining," Technical Report PIER Working Paper No. 01-028, Washington University, St. Louis 2000.

Greif, Avner, "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition," *American Economic Review*, 1993, *83*, 525–548.

Kandori, Michihiro, "Social Norms and Community Enforcement," *Review of Economic Studies*, 1992, *59*, 63–80.

Karlan, Dean, Markus Mobius, Tanya Rosenblat, and Adam Szeidl, "Trust and Social Collateral," *forthcoming in the Quarterly Journal of Economics*, November 2009, *124* (4).

Kazianga, Harounan and Christopher Udry, "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso," *Journal of Development Economics*, April 2006, *79* (2), 413–446.

Kocherlakota, Narayana, "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies*, 1996, *63* (4), 595–609.

Kranton, Rachel E., "Reciprocal Exchange: A Self-Sustaining System," *American Economic Review*, 1996, *86*, 830–851.

Lee, Thomas R., Jay A. Mancini, and Joseph W. Maxwell, "Sibling relationships in adulthood: contact patterns and motivatioins," *Journal of Marriage and Family*, May 1990, *52* (2), 431–440.

Ligon, Ethan, "Risk-sharing and Information in Village Ecocomies," *Review of Economic Studies*, 1998, *65*, 847–864.

——, **J. P. Thomas, and T. Worrall**, "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies*, 2002, *69* (1), 209–244.

Litwin, Howard, "The social networks of elderly immigrants: An analytic typology," *Journal of Aging Studies*, 1995, *9* (2), 155 – 174.

——, "Social network type and health status in a national sample of elderly Israelis," *Social Science & Medicine*, 1998, *46* (4-5), 599 – 609.

Mace, Barbara J, "Full insurance in the presence of aggregate uncertainty.," *Journal of Political Economy*, 1991, *99*, 928ï£¡956.

Mazzocco, Maurizio, "Household Intertemporal Behavior: a Collective Characterization and a Test of Commitment," *Review of Economic Studies*, July 2007, *74* (3), 857–895.

Milgrom, Paul and Chris Shannon, "Monotone Comparative Statics," *Econometrica*, January 1994, *62* (1), 157–80.

Ogaki, Masao and Qiang Zhang, "Deceasing Relative Risk Aversion and Tests of Risk Sharing," *Econometrica*, March 2001, *69* (2), 515–526.

Topa, Giorgio, "Social Interactions, Local Spillovers and Unemployment," *Review of Economic Studies*, 2001, *68*, 261–295.

Townsend, Robert, "Risk and Insurance in Village India," *Econometrica*, May 1994, *62* (3), 539–591.

Udry, Chris, "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria," *Review of Economic Studies*, 1994, *61* (3), 495–526.

Wilson, Robert, "The Theory of Syndicates," *Econometrica*, January 1968, *36* (1), 119–132.

Wormald, Nicholas C., "Models of random regular graphs," in J.D. Lamb and D.A. Preece, eds., *Surveys in Combinatorics: London Mathematical Society Lecture Note Series*, Vol. 276, Cambridge: Cambridge University Press, 1999, pp. 239–298.

Appendix A: Proofs

Proof of Theorem 1

The theorem can be generalized to the case where links in the network are directed, so that $c(i,j)$ and $c(j,i)$ may differ. In that environment, coalition proofness now requires that

$$e_F - x_F \leq c^{\text{out}}[F] \tag{6}$$

where $c^{\text{out}}[F] = \sum_{i \in F, j \notin F} c(i,j)$ is the maximum amount that agents in F are willing to give to the outside community. Here we present a proof of this more general result. Sufficiency follows from the discussion in the text. To prove necessity, let $g_i = e_i - x_i$ the amount that i has to transfer away, and let $g_F = \sum_{i \in F} e_i$ for any subset of agents F. Note that $g_W = 0$ by $e_W = x_W$. Let U be the set of agents for whom $g_i \geq 0$ and let $D = W \backslash U$. Define the auxiliary graph G' which has two additional vertices, s and t, and additional edges connecting s with all agents in U, and additional edges connecting t with all agents in D. For any $i \in U$, define the capacity $c(s,i) = g_i$ and $c(i,s) = 0$. Similarly, for any $j \in D$, let $c(j,t) = -g_j$ and $c(t,j) = 0$.

The auxiliary graph is useful, because implementing the desired consumption allocation with a transfer scheme that meets the capacity constraints is equivalent to finding an $s \to t$ flow in G' that has value $g_U = \sum_{g_i \geq 0} g_i$. To see why, note that in the desired allocation, exactly g_i must leave each agent $i \in U$. The capacities on the new links ensure that in any $s \to t$ flow, at most g_i can leave agent i. Similarly, to implement the target, exactly $-g_j$ must flow to each agent $j \in D$, and the capacity on the (j,t) link ensures that this is the maximum that can flow to j. As a result, any flow with value $\sum_{g_i \geq 0} g_i$ must, by construction, take exactly g_i away from i and deliver exactly g_j to j.

We have reduced our implementation problem to a flow problem. To compute the maximum $s \to t$ flow, we instead compute the value of the minimum cut. Fix a minimum cut, let S be the set of agents in W that are still connected to s after the cut, and let $T = W \backslash S$. Clearly, if we consider the restriction of the cut to the original network G, there will be no surviving paths connecting some agent in S with some other agent in T.

Let $U_1 \subseteq U$ denote those agents whose link with s is cut in the minimum cut of G', and let $D_1 \subseteq D$ denote those in D whose link with t is cut. Let $U_2 = U \backslash U_1$ and $D_2 = D \backslash D_1$ be the sets of agents whose link with s respectively t remains; then $U_2 \subseteq S$ and $D_2 \subseteq T$, because otherwise there would be surviving path in G' connecting s and t after the cut. This also implies that $g_S \geq g_{U_2} + g_{D_1}$, because

$$g_S = g_{S \cap U} + g_{S \cap D} \geq g_{U_2} + (g_D - g_{D2}) = g_{U_2} + g_{D_1} \tag{7}$$

where we used that $g_i \geq 0$ when i is in U and negative when i is in D.

The value of the cut in G' can be bounded as

$$\text{cut value} \geq g_{U_1} - g_{D_1} + c^{\text{out}}[S]$$

where the first two terms count the total capacity of links with s and t that have been deleted, and the final term is a lower bound for links deleted from the original network G. By assumption (6), $c^{\text{out}}[S] \geq e_S - x_S = g_S$, and using (7) we obtain

$$\text{cut value} \geq g_{U_1} - g_{D_1} + g_{U_2} + g_{D_1} = g_{U_1} + g_{U_2} = g_U.$$

It follows that the value of the maximum flow is at least g_U, as desired.

Appendix B: Data

Dean Karlan, Markus Mobius and Tanya Rosenblat conducted a survey in November 2006 in a rural village close to Huaraz (Peru). The heads of households and spouses (if available) of 223 households were interviewed. The survey consisted of two components: a household survey and a social network survey. The household survey recorded a list of all members of the household and basic demographic characteristics including gender, education, occupation and income.

The social network component of the survey asked the head of household and the spouse to list up to 10 non-relatives in the community with whom the respondent spends the most time with in an average week. Respondents were also asked separately to list their first and second-degree relatives (excluding relatives related through marriage). We use this data to construct an undirected social network where two agents have a *friendship link* if one of them names the other as a friend and as a *relative link* if one of them lists the other as relative. We also added *intra-household links* between all members of a household which are assumed to be of unlimited strength. Individuals have, on average, 1.84 relative links and 1.95 non-relative links.

In the survey, individuals were also asked whether they borrow or lend money or object across each link. This data was aggregated on the household level and used to construct figure 1.