CONSUMPTION RISK-SHARING IN SOCIAL NETWORKS

Attila Ambrus
Markus Mobius
Adam Szeidl

Consumption Risk-sharing in Social Networks
Attila Ambrus, Markus Mobius, and Adam Szeidl
NBER Working Paper No. 15719
February 2010
JEL No. D02,D31,D70

ABSTRACT

We develop a model of informal risk-sharing in social networks, where 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.

Attila Ambrus
Department of Economics
Harvard University
Littauer 309
Cambridge, MA 02138
ambrus@fas.harvard.edu

Markus Mobius
Harvard University
Department of Economics
Littauer 327
Cambridge, MA 02138
and NBER
mobius@fas.harvard.edu

Adam Szeidl
Department of Economics
University of California, Berkeley
517 Evans Hall, #3880
Berkeley, CA 94720-3880
szeidl@econ.berkeley.edu

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 can local obligations and transfers aggregate up to good global risk-sharing? We build a simple model of risk-sharing in social networks that provides an explanation for this puzzle. We find that 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, Rangel and Rasul (2008), who therefore provide indirect evidence for our model.

We model the social network as a set of pre-existing relationships, like friendships and family ties. These links have utility value, which represents 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).

Huaraz Community

Agricultural tool (74.3%)
Other tool (7.3%)
Animals (1.8%)
Electric device (0.8%)
Kitchen utensil (5.6%)
Clothes (0.7%)
Food (7.0%)
Other (2.5%)

— 0 to 10 S/.   40.5%
— 11 to 20 S/.   15.6%
— 21 to 50 S/.   17.1%
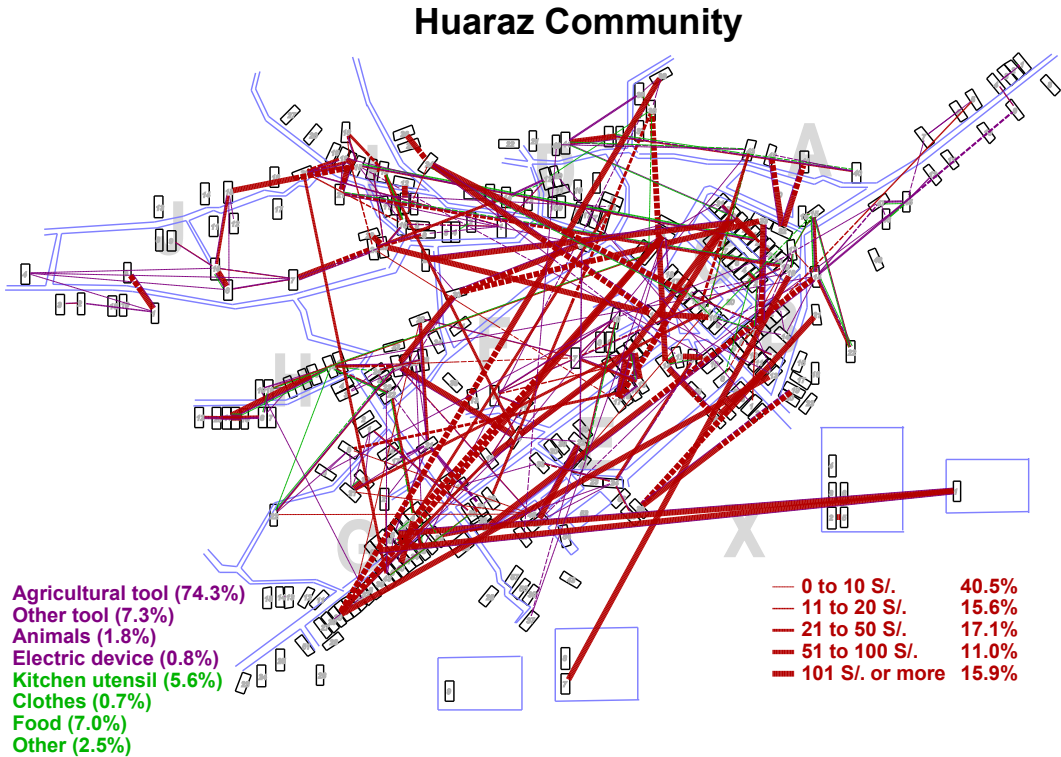— 51 to 100 S/.   11.0%
— 101 S/. or more   15.9%

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.

Our first result is that an incentive-compatible risk-sharing arrangement always gives rise to a consumption allocation that is *coalition-proof* in the following sense: the net transfer from any group of agents to the rest of the community, defined as the difference between the group's total endowment and consumption, cannot exceed the sum of the values of all links between the group and the community. Intuitively, 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.

This equivalence between coalition-proof allocations and incentive compatible risk-sharing

2

arrangements has two implications. First, it shows that decentralized insurance arrangements 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.

Figure 2: Expansion properties of three example networks

A. Line

$c[F] \sim 1$

B. Plane

$c[F] \sim |F|^{1/2}$

C. Binary tree

$c[F] \sim |F|$

3
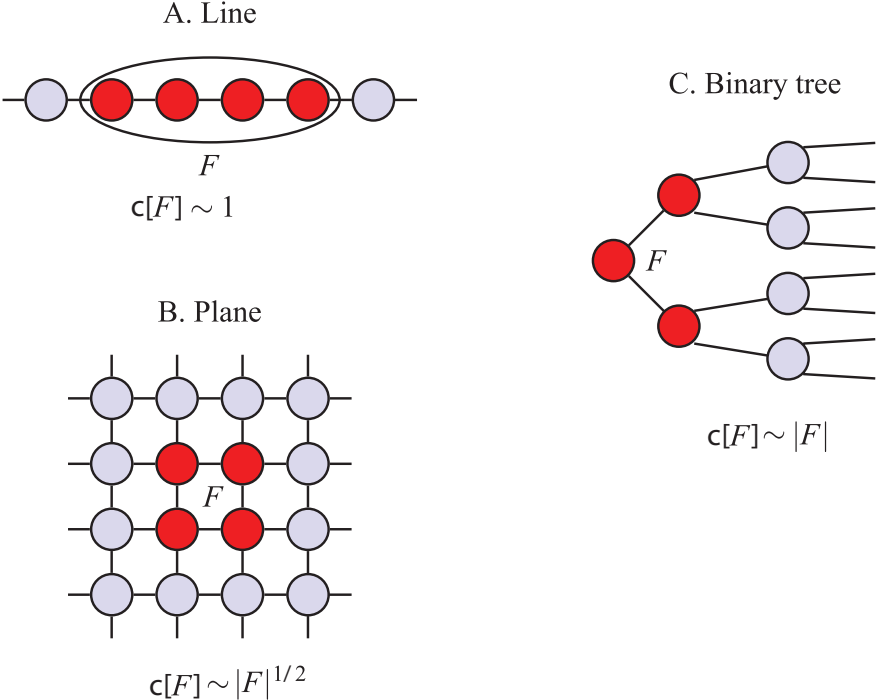
We show that full insurance requires highly expansive networks like the infinite binary tree. However, we do not find that real-world social networks in rural villages in Peru exhibit this large degree of expansiveness. Instead, these social networks are more similar to planar networks, possibly because people tend to have connections at close geographic distance. We next show that a two-dimensional structure, such as 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 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. Our second main contribution is a qualitative analysis of constrained efficient "second-best" arrangements. We show that in these arrangements, the network can be partitioned into endogenously organized connected groups called "risk-sharing islands" for every realization of uncertainty. 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," who 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. (2008).

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.

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

---

[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.

of consumption insurance. These papers do not study the effects of network structure.

The rest of this paper is organized as follows. Section 1 presents our model of informal insurance in networks. Section 2 characterizes the limits to risk-sharing, and confronts the theoretical results with data on social networks in Peru. Section 3 analyses constrained efficient arrangements. Section 4 explores a more general version of our model and Section 5 concludes. Proofs are delegated to Appendix A and a supplementary appendix.

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 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.

6

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 $t^e = (t^e_{ij})$, 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 $U_i(x_i + c_i)$. Section 4 develops the model with imperfect substitutes, and shows that under weak conditions, all 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.

Capacities and dynamic interpretation. 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. However, our setup can also be viewed as a "snapshot" of a dynamic model, where the value of a network connection is determined in part by the ability to conduct insurance transactions through the link in the future. In such a dynamic model, link capacities would be endogenized by the expected future benefits from risk-sharing. As Bloch et al. (2008) show in a similar model, this leads to restrictions on the equilibrium network structure and link values. While our static analysis applies for any set of capacities, our results could presumably be strengthened by imposing such restrictions on the network. We plan to explore the implications of dynamics more explicitly in future work.

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.[4] 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 village environments, where 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.

---

[4]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. (2010) (available at `http://www.socialcollateral.org/risksharing/supplementary_appendix.pdf`).

1.3 Coalition-proof allocations

We first show that incentive compatible risk-sharing arrangements give rise to consumption allocations that are *coalition-proof* in every state of the world in the following sense. 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, cannot exceed the sum of the values of all links connecting the group and the rest of the community. Formally, for any group F we define the *perimeter* $c[F]$ 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 "joint obligation" of the group F to the rest of the community. Similarly, we define the joint endowment of the group as e_F and the joint consumption allocation induced by the risk-sharing arrangement with x_F. Coalition-proofness then requires $e_F - x_F \leq c[F]$ for all F, i.e., the net transfer from the group to the community cannot exceed the group's joint obligation $c[F]$.[5]

Surprisingly, coalition-proofness tightly characterizes all the consumption allocations that are implementable through informal risk-sharing.

Theorem 1 *A consumption allocation x that is feasible $(\sum x_i = \sum e_i)$ is coalition-proof in every state of the world **if and only if** it can be implemented by an incentive-compatible informal risk-sharing arrangement.*

That an incentive compatible allocation is coalition proof is easy to see: 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. Proving the converse is more difficult, and builds on the mathematical theory of network flows. Recall that the *maximum flow* between nodes s and t in a network is the highest amount that can flow from s to t along the edges respecting the capacity constraints. Finding a transfer representation for a coalition-proof

---

[5]The supplementary appendix shows that this definition of coalition-proofness in our context is equivalent to defining coalition-proofness along the lines of Bernheim, Peleg and Whinston (1987), i.e., allowing only for coalitional deviations that are not prone to further deviations by subcoalitions.

allocation turns out to be equivalent to finding a flow in an auxiliary network with two additional nodes s and t added. According to the theorem of Ford and Fulkerson (1956), the maximum flow equals the value of the minimum cut, i.e., the smallest capacity that must be deleted so that s and t end up in different components. We prove in Appendix A that each cut in the flow problem corresponds to a coalition, and then the coalition-proofness condition ensures that the cut values are high enough so that the desired flow can be implemented.

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 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.[6] 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 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.

---

[6] 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.

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.[7] 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 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 line (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 $\sigma \cdot |F|$, 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

---

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

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.[8]

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 $a[F] \geq \left(1 - \frac{|F|}{|W|}\right) S$.*

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, $a[F]$ must be bounded away from zero for such sets as the network size grows without bound. The intuition builds on our earlier examples: risk-sharing between F and the rest of the community is hardest to support when everyone in F gets the maximum realization and everyone outside F 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.[9]

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.[10]

Figure 3A compares the expansiveness of the Huaraz network with the line, plane, and infinite binary tree. 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:

---

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

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

[10]The data was collected by Dean Karlan, Markus Mobius and Tanya Rosenblat and is described in Appendix B in more detail.

Figure 3: Expansiveness of the social network in Huaraz, Peru
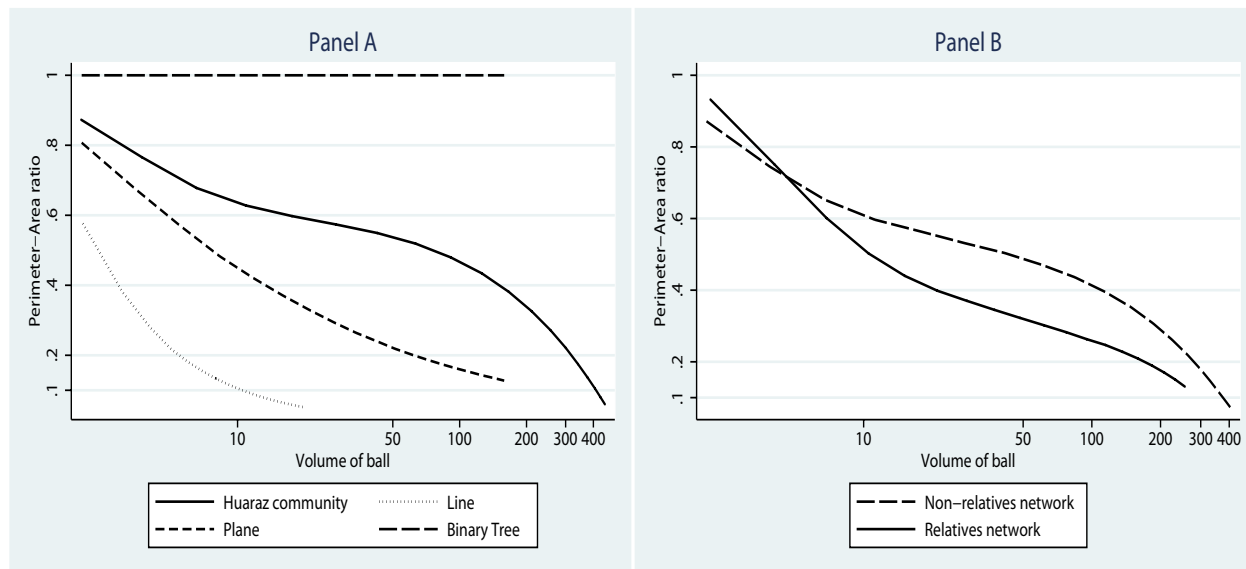
the perimeter-area ratio goes to zero quickly for the line network, goes to zero more slowly for the plane, and remains bounded away from zero for the binary tree.

The key curve in the Figure is the heavy line representing the social network in Huaraz. This curve lies slightly above the plane but well below the infinite binary tree, and approaches zero as set sizes grow, with a slope that parallels the curve for the plane. It follows that the Huaraz network is less expansive than the infinite binary tree, 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 3B: the non-relative network is slightly more expansive, but does not approach the expansiveness of the binary tree.

Figure 3 suggests that the expansion properties of the Huaraz network are similar to 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, we should focus on plane-like networks.

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.[11] 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 $\overline{e} = e_W/|W|$:

---

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

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

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

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

which is the square-root of the expected cross-sectional variance of x. For non-quadratic utilities, $SDISP(x)$ can be interpreted as a second order approximation of the utility based measure. $SDISP$ 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 $SDISP(e) = \sigma$. We use $SDISP$ as our central measure in the analysis below.[12]

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. From now on we use the convention that K, K' and K'' denote positive constants, but their values at different occurrences may be different.

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.

---

[12]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.

(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 state the formal result on risk-sharing on the plane and line networks. 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 4 shows constrained optimal allocations for finite line and plane networks, for a given realization of uniform shocks with support $[-1, 1]$.[13] Figure 4A 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 4B, 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
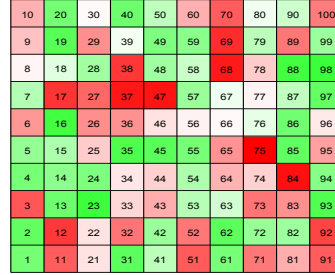
---

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

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

Panel A: initial endowments (uniform over $[-1, 1]$)
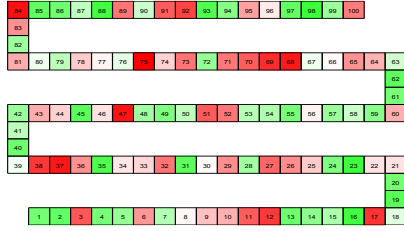
SDISP= 0.556   SDISP= 0.556

Panel B: risk-sharing with total capacity 1 per agent
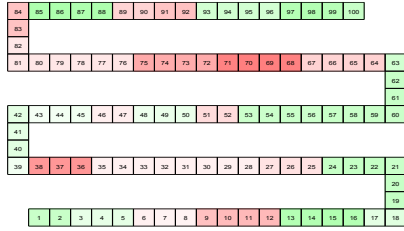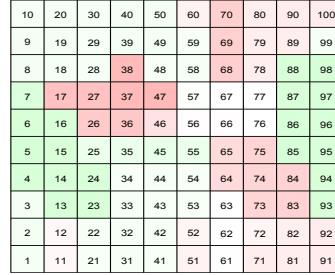
SDISP= 0.246   SDISP= 0.127
30 islands   17 islands

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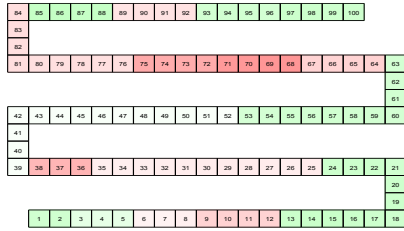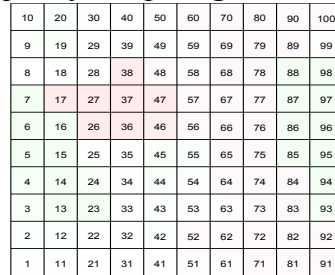

SDISP= 0.199   SDISP= 0.035
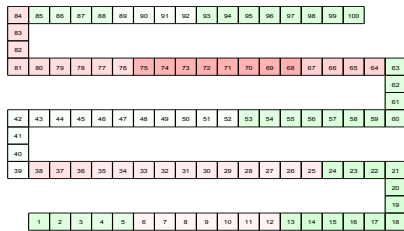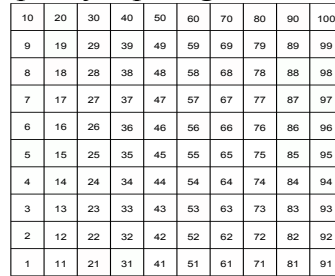17 islands   4 islands

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

SDISP= 0.148   SDISP= 0
13 islands   1 island

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 4C, 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 4D, 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. However, as Figure 1 illustrates, real-world social networks have a much less regular structure. 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 $d\left(\pi\left(i\right), \pi\left(j\right)\right) \leq d$, there is a path connecting i and j such that for all agents h in the path, $d\left(\pi\left(i\right), \pi\left(h\right)\right)$ 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.[14]

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 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. We next check whether the Huaraz village network is a geographic network. Figure 5A shows the natural geographic map of household locations, referred to as lots, in this village. In Figure 5B 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 5B, 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 5C and 5D depict the result after one and five rounds of iteration: the distribution of lots becomes gradually more homogenous. After 23 iterations (Figure 5E), the distribution of lots is almost completely uniform. Figure 5E 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 5E also shows the number

---

[14] 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.

Figure 5: Stretching a real-world network to construct a geographic representation

Panel A: original map of Huaraz community
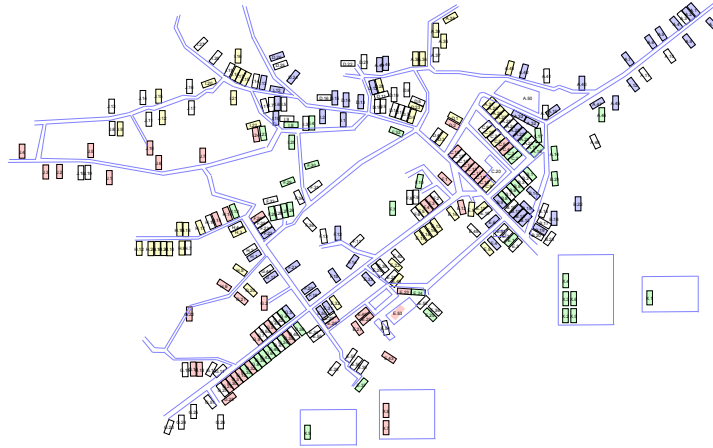
Panel B: iteration 0
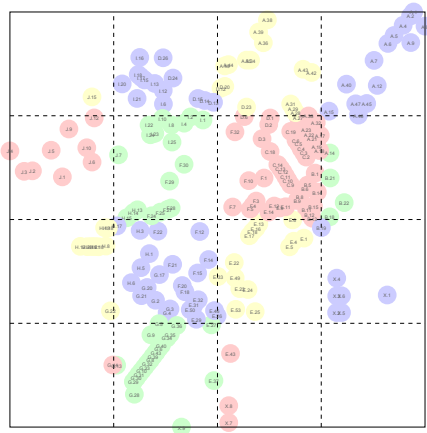
Panel C: iteration 1
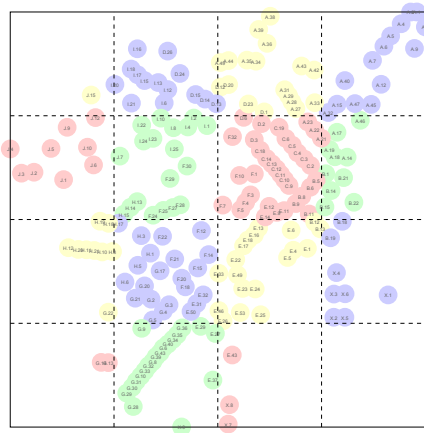
Panel D: iteration 5
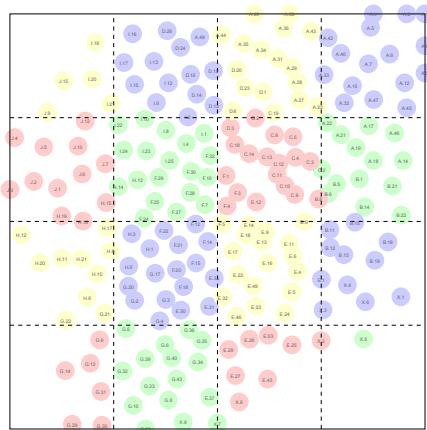
Panel E: iteration 23
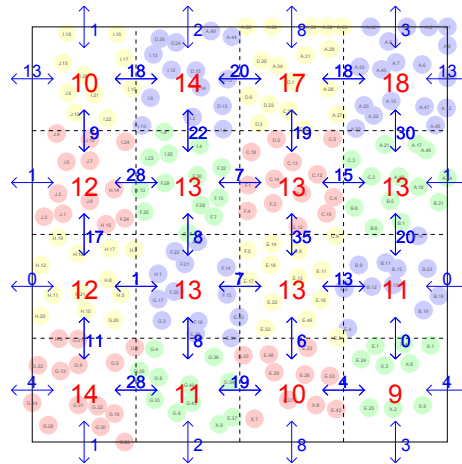
of links crossing the sides of each square.[15] The agreement with our theoretical condition is very good: except for one side of the square in the lower right corner, there are no separating avenues between any two neighboring squares. To better understand what drives the success of this embedding, note that in Figure 5E 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 5A), 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.

Numerical simulations suggest that the Huaraz social network in fact behaves very much like the plane network. We use simulations to 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 is 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 – a result consistent with the empirical findings of Townsend (1994), Ogaki and Zhang (2001), Mazzocco (2007) and others.

2.5   Risk-sharing ability of a group

One commonly used approach to testing full risk-sharing in the data is to regress the consumption of an individual or a group on their own endowment and a community-wide shock. A variant of this regression when there is no aggregate uncertainty is

$$x_F = \alpha + \beta \cdot e_F + \varepsilon$$

where consumption in F is regressed on the endowment shock of F. Equal sharing implies $\beta = 0$; this corresponds to the test of full risk-sharing used in Cochrane (1991), Mace

---

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

(1991), Townsend (1994) and others. When $\beta \neq 0$, equal sharing is rejected; however, small magnitudes of the coefficient can be interpreted to mean that agents in F share their risk with the rest of the community reasonably well. The following result supports this interpretation.

Proposition 3 *We can bound the regression coefficient β as*

$$1 - \frac{c\,[F]}{\sigma_F} \leq \beta.$$

This lower bound is a function of the perimeter $c\,[F]$ relative to the standard deviation of the community-specific shock σ_F. The intuition is familiar: when the perimeter of a set is small, much of the idiosyncratic shock is trapped inside F, resulting in higher correlation. The Proposition is related to Townsend's (1994) finding that there is considerable risk-sharing within, but only limited sharing across villages, as well as Rosenzweig and Stark (1989) who show that Indian households try to create cross-village family links through marriage. Our results are consistent with these facts if cross-village network ties are relatively weaker.

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. As we show below, foundations for these arrangements include both simple rules of thumb and dynamic coalitional bargaining.

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 4 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 4 *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.*

Wilson (1968) establishes a similar equivalence result for risk-sharing in syndicates. His proof builds on the convexity of the set of possible payoff vectors. Since an efficient allocation must lie on the boundary of this set, convexity implies the existence of a tangent hyperplane with some normal vector (λ_i). Maximizing a planner's problem with these λ_i weights will then select the efficient allocation. Adapting this argument to our model requires that the set of coalition-proof payoff vectors be convex. This is straightforward given the perfect substitutes specification: when two transfers satisfy a capacity constraint, their convex combination will also satisfy it.[16] The formal proof is provided in Appendix A.

Observe that 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

---

[16]See section 4 for extending this result to imperfect substitutes.

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 5 *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 6 *[Risk-sharing islands] In any realization (e_i) 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 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) that 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 4, panels B to D. In Figure 4B, 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 4C there are 17 islands on the line and only 4 on the plane; and in Figure 4D 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 4, 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 must 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 as defined above. 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 6 implies that links connecting different islands are all blocked. Except for knife-edge cases when the transfer amount just reaches the capacity over a link but does not "bind" yet, the two definitions are equivalent: $\widehat{W}(i) = W(i)$. It can be shown that these knife-edge cases have zero probability when the distribution of shocks is absolutely continuous, and hence the two definitions can be treated as equivalent for practical purposes.

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

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

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 7 *[Spillovers and local sharing] 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 $x_j(e') = x_j(e)$ for all $j \in W \backslash W(i)$.

(iii) [More sharing with close friends] For any $j \neq i$, there exists a path $i \to j$ 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 $\widehat{W}(i)$ 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.[17] 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.

---

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

Figure 6: Utility cost of shocks to direct and indirect friends
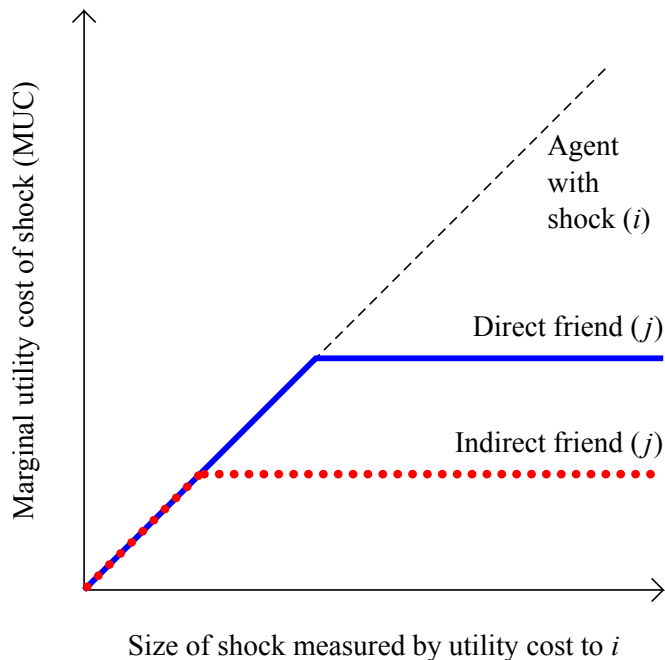
The results of Proposition 7 are summarized in Figure 6, which shows the marginal utility cost of direct and indirect friends in response to a shock to i. The horizontal axis is the marginal utility cost of i himself while the vertical axis measures the MUC of direct and indirect friends. For small shocks, both direct and indirect friends in the same island help out. As the size of the shock grows, some indirect friends hit their capacity constraints (dotted line), but some direct friends continue to help (heavy line). After a point, all direct friends hit their capacity constraints and additional increases in the shock are fully borne by agent i. These implications can be used to test our model against other theories of limited risk-sharing, which do not predict variation in the degree of insurance as a function of network distance.

The results of Proposition 7 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. (2008) 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.

3.3 Foundations for constrained efficiency

We now briefly discuss two intuitive dynamic mechanisms that provide foundations for constrained efficiency. The supplementary appendix contains the corresponding formal results. First consider a decentralized procedure where agents use a simple rule of thumb in helping those who are in need. In every round, agents attempt to equate weighted marginal utilities between neighbors subject to the capacity constraints: intuitively, people help out less fortunate friends. This procedure converges to the constrained efficient allocation corresponding to the welfare weights used. In particular, constrained efficiency can emerge even if in every transaction agents only use local information about the current resources of the parties involved.[18]

A second mechanism is collective dynamic bargaining with renegotiation. Gomes (2000) shows that when agents can propose renegotiable arrangements to subgroups and make side-payments in a dynamic bargaining procedure, ultimately a Pareto-efficient arrangement will be selected.[19] This result can be incorporated in our model by assuming that there is a negotiations phase prior to the endowment realization, and would imply that agents select a constrained efficient risk-sharing arrangement.

We also note that constrained efficient arrangements are particularly stable in that they are also robust to ex ante coalitional deviations, not just ex post ones.

4 Discussion: General preferences

This section discusses how our results extend when goods and friendship are imperfect substitutes. Formal statements and proofs are presented in the supplementary appendix, here we only summarize our findings.

---

[18]Bramoulle and Kranton (2006) use a similar procedure with equal welfare weights and no capacity constraints.

[19]Aghion, Antras and Helpman (2007) establish a similar result in a model involving renegotiating free-trade agreements.

With a general utility function $U(x, c)$, the definition of incentive compatibility (IC) of a transfer arrangement is the following:

Definition 2 *A risk-sharing arrangement t is incentive compatible (IC for short) if*

$$U_i(x_i, c_i) \geq U_i(x_i + t_{ij}, c_i - c(i, j)) \tag{4}$$

for all i and j, for all realizations of uncertainty.

Our key tool is a pair of necessary and sufficient conditions for incentive compatibility with imperfect substitutes. To derive these, define the marginal rate of substitution (MRS) between good and friendship consumption as $MRS_i = (\partial U_i / \partial c_i) / (\partial U_i / \partial x_i)$. We say that the MRS is uniformly bounded if there exist positive constants $m < M$ such that $m \leq MRS_i \leq M$ for all i, x_i and c_i.

When the MRS is uniformly bounded, *(i)* any IC arrangement must satisfy $t_{ij} \leq M \cdot c(i, j)$, and *(ii)* any arrangement that satisfies $t_{ij} \leq m \cdot c(i, j)$ must be IC. The intuition is that the MRS measures the relative price of goods and friendship. If this relative price is always between m and M, then a transfer exceeding $Mc(i, j)$ is always worth more than the link and hence never IC, but a transfer below $mc(i, j)$ is always worth less than the link and hence is IC. With perfect substitutes $MRS_i = 1$, so we can set $m = M = 1$, which yields Theorem 1.

4.1 The limits to risk-sharing with imperfect substitutes

With imperfect substitutes, the results in section 2 extend but the upper and lower bounds on risk-sharing are weakened by constant factors that depend on the degree of substitution. To obtain these extensions, we assume that the marginal rate of substitution (MRS) is uniformly bounded. We continue to find that the first-best can only be achieved in highly expansive graphs where the perimeter-area ratio is bounded from below: we require $a[F] \geq \underline{\sigma}/M$. Our findings about partial risk-sharing are about rates of convergence and hence they extend without modification; in particular, $SDISP$ converges exponentially for geographic networks.

Imperfect substitution also yields additional implications. If the MRS is increasing in consumption, then agents with low consumption value their friends less, reducing the maximum amount they are willing to give up. As a result, if in a society that experiences a negative aggregate shock, the scope for insuring idiosyncratic risk is reduced. We show that reducing the endowments of all agents results in a smaller set of incentive compatible transfer arrangements, and hence an increase in $SDISP$. The aggregate negative shock is thus a double burden: besides its direct negative effect on consumption, it also induces worse sharing of idiosyncratic risks, a finding consistent with Kazianga and Udry (2006), who document limited informal insurance during the severe draught of 1981-85 in rural Burkina Faso.

4.2 Constrained efficient arrangements

The key novelty with imperfect substitutes is that changing the goods consumption of an agent affects his implied link values and hence incentive compatibility. To characterize constrained efficiency, we assume that the marginal rate of substitution MRS_i defined above is concave in x_i. When this holds, we can generalize Proposition 4, establishing the equivalence between constrained efficiency and the planner's problem.

To develop first order conditions, we next analyze the effect of an additional dollar to agent i on the planner's objective. With imperfect substitutes, this marginal welfare gain is no longer equal to λ_i times the marginal utility of i, because increased consumption also softens enforcement constraints. The planner may wish to use these softer constraints and transfer some of the original dollar to neighboring agents. To formalize this, we define the marginal social gain of an additional unit of transfer to i using an iterative procedure, which takes into account the indirect effect of softening constraints.

Using the concept of marginal social gain allows us to extend the characterization of constrained efficient agreements in Proposition 5. Given this result, we can also partition the network into endogenous risk-sharing islands, such that marginal social utility is equalized within islands, and all links connecting the island to the rest of the community are blocked.

Finally, for an agent i who is not on the boundary of his risk-sharing island and hence has no links with binding constraints, the marginal social gain does equal λ_i times his

marginal utility of consumption; hence, for such agents, the results of section 3 hold without modification. For example, weighted marginal utilities are equalized for any two such agents in the same risk-sharing island. Thus if risk-sharing islands are "large", then the results from the perfect substitutes case hold without modification for most agents.

5   Conclusion

This paper showed that the expansiveness of a social network determines the effectiveness of informal risk-sharing. We found that many real-life social networks are likely to be sufficiently expansive to allow for good risk-sharing. We also characterized Pareto-optimal arrangements and found that resources are shared among local groups.

In future work, we would like to develop a dynamic extension of our model, where 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` January 2010.

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, Marcos Rangel, and Imran Rasul**, "Insurance in the Extended Family," *Working paper, University of Arizona*, 2008.

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.*

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.

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), 248Ű264.

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.

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

Ford, Lester Randolph and Delbert Ray Fulkerson, "Maximal Flow Through a Network," *Canadian Journal of Mathematics*, 1956, *8*, 399–404.

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.

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.

Mace, Barbara J, "Full insurance in the presence of aggregate uncertainty.," *Journal of Political Economy*, 1991, *99*, 928Ű956.

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.

Rosenzweig, Mark R. and Oded Stark, "Consumption Smoothing, Migration, and Marriage: Evidence from Rural India," *Journal of Political Economy*, August 1989, *97* (4), 905–26.

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.

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

$$e_F - x_F \leq c^{\text{out}}[F] \tag{5}$$

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

$$g_S = g_{S \cap U} + g_{S \cap D} \geq g_{U_2} + (g_D - g_{D2}) = g_{U_2} + g_{D_1} \tag{6}$$

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 (5), $c^{\text{out}}[S] \geq e_S - x_S = g_S$, and using (6) 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.

Proof of Proposition 4

We prove the following more general result.

Suppose that the $MRS_i = (\partial U_i/\partial c_i)/(\partial U_i/\partial x_i)$ is concave in x_i for every i. Then every constrained efficient arrangement is the solution to a planner's problem with some set of weights (λ_i), and conversely, any solution to the planner's problem is constrained efficient.

Proof. Let $U^* \subseteq \mathbb{R}^W$ be the set of expected utility profiles that can be achieved by IC transfer arrangements: $U^* = \{(v_i)_{i \in W} | \exists$ IC allocation x such that $v_i \leq EU_i(x_i, c_i) \forall i\}$. Our goal is to show that U^* is convex. By concave utility, it suffices to prove that the set of IC arrangements is convex.

To show that the convex combination of IC arrangements is IC, fix an endowment realization e and let x be an IC allocation. Consider an agent i, and for $r \geq 0$ define $y(r, x_i)$ to be the consumption level that makes i indifferent between his current allocation and reducing friendship consumption by r units, that is, $U(x_i, c_i) = U(y(r, x_i), c_i - r)$. For different values of r, the locations $(y(r, x_i), c - r)$ trace out an indifference curve of i. Note that $y(0, x_i) = x_i$ and that the IC constraint for the transfer between i and j can be written as

$$t_{ij} \leq y(c(i, j), x_i) - x_i \tag{7}$$

since $y(c(i, j), x_i) - x_i$ is the dollar gain that makes i accept losing the friendship with j.

Moreover, the implicit function theorem implies that

$$y_r\left(r, x_i\right) = \frac{U_c}{U_x}\left(y, c_i - r\right) \tag{8}$$

which is the marginal rate of substitution MRS_i. This is intuitive: MRS_i measures the dollar value of a marginal change in friendship consumption. Using the concavity of the MRS, we will show that $y\left(r, x_i\right)$ is a concave function in x_i for any $r \geq 0$. When $r = c\left(i, j\right)$, this implies that the convex combination of IC allocations also satisfies the IC constraint (7), and consequently, that the set of IC profiles is convex.

To show that $y\left(r, x_i\right)$ is concave in x_i, let x^1, x^2 be two IC allocations, and let $x_i^3 = \alpha x_i^1 + \left(1 - \alpha\right) x_i^2$ for some $0 \leq \alpha \leq 1$. Define $\overline{y}\left(r\right) = \alpha y\left(r, x_i^1\right) + \left(1 - \alpha\right) y\left(r, x_i^2\right)$, so that $\left(\overline{y}\left(r\right), c_i - r\right)$ traces out the convex combination of the indifference curves passing through $\left(x_i^1, c_i\right)$ and $\left(x_i^2, c_i\right)$, and let $f\left(r\right) = U\left(\overline{y}\left(r\right), c_i - r\right)$, the utility of agent i along this curve. Clearly, $f\left(0\right) = U\left(x_3, c_i\right)$. Moreover, using (8),

$$f'\left(r\right) = U_x\left(\overline{y}\left(r\right), c_i - r\right) \cdot \left[\alpha\frac{U_c}{U_x}\left(y\left(r, x_i^1\right), c_i - r\right) + \left(1 - \alpha\right)\frac{U_c}{U_x}\left(y\left(r, x_i^2\right), c_i - r\right)\right] - U_c\left(\overline{y}\left(r\right), c_i - r\right)$$
$$\leq U_x\left(\overline{y}\left(r\right), c_i - r\right) \cdot \frac{U_c}{U_x}\left(\overline{y}\left(r\right), c_i - r\right) - U_c\left(\overline{y}\left(r\right), c_i - r\right) = 0$$

where we used the assumption that U_c/U_x is concave in the first argument. It follows that f is nonincreasing, and in particular $f\left(r\right) \leq f\left(0\right)$ or equivalently $U\left(\overline{y}\left(r\right), c_i - r\right) \leq U\left(x_i^3, c_i\right)$, which implies that $y\left(x_i^3, r\right) \geq \overline{y}\left(r\right) = \alpha y\left(r, x_i^1\right) + \left(1 - \alpha\right) y\left(r, x_i^2\right)$, and hence that $y\left(x, r\right)$ is concave.

Finally, let $P(U^*)$ denote the Pareto-frontier of U^*. Since U^* is convex, the supporting hyperplane theorem implies that for every $u^0 \in P(U^*)$ there exist positive weights λ_i such that $u^0 \in \arg\max_{U^*} \sum_i \lambda_i u_i$, as desired. The converse statement in the proposition holds for any U^*.

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.

Supplementary Material to: "Consumption Risk-sharing in Social Networks"[1]

Attila Ambrus
Harvard University

Markus Mobius
Harvard University and NBER

Adam Szeidl
UC - Berkeley

January 2010

This material supplements the paper "Consumption Risk-sharing in Social Networks". First of all, we provide missing proofs for results stated in the main paper. Second, we discuss five extensions to the main paper. (1) We provide game theoretic micro-foundations to justify our assumption that links "die" when a promised transfer is not made. (2) We provide background about the mathematical theory of network flows used in the proofs of the paper. (3) We formalize two decentralized mechanisms leading to constrained efficient allocations. (4) We formally develop the extensions of our main results to the case where goods and friendship consumption are imperfect substitutes. (5) We explain the numerical methods used to simulate the model and to construct the geographic network representation of the real world Huaraz network.

A-1   Missing Proofs for Sections 1 to 4

Proof that coalition-proof arrangements are robust to deviating subcoalitions

Our definition of coalition-proofness in the risk-sharing context is equivalent to Bernheim et al.'s (1987) stricter definition of coalition-proofness who only allow for coalitional deviations that are not prone to further deviations by subcoalitions. We establish this result without the perfect substitutes assumption, i.e., for general $U_i(x_i, c_i)$ utility functions.

---

[1]E-mails: ambrus@fas.harvard.edu, mobius@fas.harvard.edu, szeidl@econ.berkeley.edu.

Proposition 8 *Requiring coalitions to be robust to further coalitional deviations does not affect the set of coalition-proof allocations.*

Proof. Let F be a deviating coalition, and let $F' \subseteq F$ be a deviating subcoalition. Then F' is also a deviating coalition in the original set of agents W. To see why, note that the capacities $\widetilde{c}'$ after the subcoalition F' deviates are exactly those associated with links within F', and hence these are also the capacities that remain when F' deviates in W. Moreover, the allocation $\widetilde{x}'$ of the subcoalition F' only uses the resources in F' and hence is also feasible when F' deviates from W. These observations imply that the same allocation is available to all agents in F' if they consider a coalitional deviation from W. Since these agents are better off with this allocation than they were in the coalition F, where in turn they are better off than in the original allocation, it follows that F' is a profitable coalitional deviation in the original network as well. Hence minimal deviating coalitions are robust to further coalitional deviations. Since any allocation that has a deviating coalition also has a minimal one, requiring no deviating coalitions is equivalent to requiring no deviating coalitions that are robust to further group deviations. ∎

Proof of Proposition 1

We denote the supremum of the support of the endowment distribution with M and the infimum with m where $S = M - m$. To show that the perimeter-area inequality implies equal risk-sharing in all states we focus on the worst case scenario where all agents inside F get the maximum endowment M and all agents outside F get the minimum m.[1] In this case, under equal sharing all agents consume $[|F| M + (|W| - |F|)m]/|W|$. This requires agents in set F to give up:

$$|F| M - |F| \left[|F| M + (|W| - |F|)m\right]/|W|$$

This amount has to be less or equal to the group's obligation which equals the perimeter $c[F]$. Some algebra reduces this inequality to $a[F] \geq \left(1 - \frac{|F|}{|W|}\right) S$. Hence the perimeter-area inequality implies that no group will want to deviate even in the worst case scenario. For the

---

[1]If the supremum and infimum do not lie in the support of the endowment distribution, we can assume realizations that are ϵ-close to the supremum and infimum and then take ϵ to 0.

same reason, coalition proofness implies the perimeter-area inequality because the coalitional IC constraint $e_F - x_F \leq c[F]$ has to hold for all states of the world.

Infinite networks in subsection 2.2

Some of our results in subsection 2.2 are stated for infinite networks. We now discuss how to extend our model to these environments. Say that a network is locally finite if W is countable, each agent has a finite number of connections, and every pair of agents is connected through a finite path. A risk-sharing arrangement specifies a consumption allocation $x(e)$ for every realization. Let B_i^r denote the set of agents within network distance r from i. The arrangement x is feasible if with probability one

$$\lim_{r \to \infty} \frac{1}{|B_i^r|} \left| e_{B_i^r} - x_{B_i^r} \right| = 0$$

for all i. This condition is a generalization of the feasibility constraint for finite networks.

We extend the concept of coalition proofness by requiring a consumption allocation x to be coalition-proof in every finite subset. Formally, let $H \subseteq W$ be a finite set of agents, and define the auxiliary network H' by collapsing all agents in $G \backslash H$ into a single node z. In this transformation, all links outside H disappear, all links between $i \in H$ and $j \notin H$ become links between i and z, and all links inside H are preserved. The capacities inherited from G in H' are denoted c_H. Fix realization e; for each $i \in H$ the consumption value x_i is well defined. For z, we let $e_z = 0$ and define x_z such that $e_H - x_H + e_z - x_z = 0$, which guarantees that the resource constraint in H' is satisfied. We also assume that the utility function of z is $c_z + x_z$. With these definitions, we have constructed a feasible allocation x' in H'. If this allocation is coalition-proof for every finite subgraph H, then we say that the original allocation x is coalition-proof in the infinite network G.

Extending Theorem 1. An informal risk-sharing arrangement can be defined in the same way as before. We claim that in this infinite network environment, the statement of Theorem 1 holds word by word. As in the finite case, sufficiency is immediate. To prove necessity, let $H_1 \subseteq H_2 \subseteq ...$ be an increasing sequence of sets such that $\cup_k H_k = W$, and fix a coalition-proof allocation x. For each k, construct the auxiliary network H'_k as above. We can define $g_i = e_i - x_i$ for all $i \in H_k$ as in the proof of Theorem 1, and let $g_z = -\sum_{i \in H_k} g_i$; with

3

this definition, we have constructed a finite implementation problem just like in the proof of Theorem 1. Since we have a coalition-proof allocation in H'_k, Theorem 1 yields an informal risk-sharing arrangement t^k in H'_k. For each (i,j) link we thus obtain a sequence of transfers $t^k_{ij} \in [-c(i,j), c(i,j)]$ for the infinite sequence of k values for which $i,j \in H_k$. Because there are only countably many links, we can select a subsequence that converges to some t^*_{ij} pointwise for each i and j. It is immediate that this transfer arrangement implements consumption allocation x and meets the capacity constraints.

Dispersion. Fix a coalition-proof allocation x in a locally finite network. To define dispersion, fix an agent i, and consider the sequence of ball sets B^r_i around i. We define the dispersion of x as in the infinite network as

$$DISP(x) = \lim_{r \to \infty} \sup DISP^r(x)$$

where $DISP^r(x) = SDISP^r(x)^2$ is just the expected cross-sectional variance of the allocation x restricted to the ball set B^r_i. We then define $SDISP(x)$ to be the square root of $DISP$ in the infinite network. We remark that in general networks, the value of $SDISP$ can depend on the initial agent i used to construct the balls. However, it is easy to see that for the line and plane networks, $SDISP$ is the same for all initial agents.

When the average endowment in the infinite network, $\bar{e} = \lim_{r \to \infty} e_{B^r_i} / |B^r_i|$ is well-defined, it is easy to see that

$$DISP(x) = \lim_{r \to \infty} \frac{1}{|B^r_i|} \sum_{j \in B^r_i} (x_j - \bar{e})^2.$$

In particular, when $\bar{e} = 0$, as in the applications we consider, one can think about $DSIP$ as the limit of the average of Ex^2_j over increasing ball sets. We will use this observation in the proofs below.

Proof of Proposition 2

The following Lemma is used in the proof.

Lemma 1 *Let Z be a random variable such that $|Z| \le c$ almost surely. Then $\sigma_Z \le c$.*

This result appears to be standard; a proof is available upon request.

(i) Dispersion on the line equals the lim sup of $SDISP$ over increasing intervals I_l of length $l = 1, 3,...$ Fix an interval of length l and split it into subintervals of length k. Throughout this argument we ignore integer constraints by assuming that l is large relative to k. For each segment F, $\sigma_F = \sigma\sqrt{k}$ and $c[F] = 2c$. Using Lemma 1, this implies that in any coalition-proof arrangement x, $\text{stdev}(x_F) \geq \sigma\sqrt{k} - 2c$. Even if agents manage to smooth x_F perfectly in F, the standard deviation of per capita consumption is at least $\text{stdev}(x_F)/k$. But this implies that in interval I_l we have $SDISP(x) \geq \text{stdev}(x_F)/k$, i.e.,

$$SDISP(x) \geq \sigma/\sqrt{k} - 2c/k.$$

To obtain the sharpest bound, let $k = 16 (c/\sigma)^2$, which gives $SDISP \geq \sigma^2/(8c)$ as desired.

(ii) We establish a result for more general networks. We fix an initial network with capacities c_0, and explore the behavior of $SDISP$ when capacities are given by $c \cdot c_0$, as $c \to \infty$. Stating the conditions that we impose on the initial network requires some additional notation. Let $G \subseteq F$ be two subsets of W, and define the relative perimeter of G in F, denoted $c_0[G]_F$, as the perimeter of G in the subgraph generated by F. With this definition, $c_0[G]_F$ simply sums the capacities over all links between G and $F\backslash G$. In the subsequent analysis, we continue to use the convention that K, K', K'', as well as K_1, K_2,... denote positive constants, and may represent different values at different occurrences. Our assumptions about the network are the following.

(N1) The network is connected, countably infinite, and all agents have at most K direct friends.

(N2) [Partition] For every $n \geq 1$ integer there exist a collection of sets F_j^i, where $i = 1, ..., n$ and $j = 1, ..., \infty$, such that F_j^i, $j = 1, ..., \infty$ give a partition of W and when $i = 1$, F_j^1 are all singletons.

(N3) [Ascending chain] For all $1 \leq i \leq n - 1$ and all j, j', we have either $F_j^i \cap F_{j'}^{i+1} = \varnothing$ or $F_j^i \subseteq F_{j'}^{i+1}$.

(N4) [Exponential growth.] There exist $1 < \underline{\gamma} < \overline{\gamma}$ constants such that whenever $F_j^i \subseteq$

$F_{j'}^{i+1}$, we have $\underline{\gamma}\left|F_j^i\right| \leq \left|F_j^{i+1}\right| \leq \overline{\gamma}\left|F_j^i\right|$.

(N5) [Relative perimeter] There exists $K > 0$ such that for any $G \subseteq F_j^i$ with $|G| \leq \left|F_j^i\right|/2$ we have $c_0\left[G\right]_{F_j^i} \geq K' \cdot c_0\left[G\right]$.

Note that we define the sets F_j^i separately for each n; we suppress the dependence on n in notation for simplicity. (N2) implies that for each i, the i-level sets partition the entire network. (N3) requires that each $i+1$-level set is a disjoint union of some i-level sets, so i-level sets partition the $i+1$-level sets. (N4) requires that the size of these sets grows exponentially; this implies in particular that the number of i-level sets in an $i+1$ level set is bounded by some constant K for all n and i. (N5) requires that the partitioning sets F_j^i are "representative" in the sense that the relative perimeter of sets inside F_j^i is the same order of magnitude as their true perimeter in G.

A specific example where (N1)-(N5) are satisfied is the plane network, where the sets can be chosen to be squares. Specifically, define F_j^n for $j = 1, 2, \ldots$ to be a partition of the plane by 2^n by 2^n sized squares. Split each of these squares in four 2^{n-1} by 2^{n-1} subsquares, and index these smaller squares by F_j^{n-1} for $j = 1, 2, \ldots$. Split these squares again and again to define F_j^i for lower values of i, until we arrive at singleton sets when $i = 1$. In this construction, conditions (N1)-(N4) are satisfied: we can set $K = 4$ for (N1) and $\underline{\gamma} = \overline{\gamma} = 4$ for (N4). It is also easy to see that (N5) is satisfied with $K' = 1/3$; equality can be achieved only when the side length of F_j^i is even, in which case G can be chosen as a rectangle-shaped half-square such that three sides of G lie on the sides of F_j^i.

To obtain a result about risk-sharing, we need to connect the network structure with the distribution of shocks. We require the following key perimeter/area condition, which can be viewed as an extension of the conditions used in Proposition 1:

(K) There exists $K > 0$ such that for all G finite, $\sigma_G \leq K \cdot c_0\left[G\right]$.

For the plane network, this condition essentially requires that for all squares F, the standard deviation σ_F is at most proportional to the side length of F, which in turn is a consequence of assumption (P3). We now state and prove the following result.

Proposition 9 *Under conditions (P1)-(P5), (N1)-(N5) and (K), there exist positive constants K' and K'' and a coalition-proof allocation $x(c)$ such that for every agent i, $Ex_i^2(c) \leq K' \exp\left[-K'' \cdot c^{2/3}\right]$.*

Proposition 2 (ii) is an immediate consequence of this result. This is because (1) the plane network satisfies conditions (N1)-(N5) and (K); and (2) $DISP$ is defined as the limit of averages of $\mathrm{E}x_i^2(c)$ over increasing sets of agents, and in consequence also satisfies the exponential bound that each $\mathrm{E}x_i^2(c)$ satisfies.

Proof. Note that (N5) and (K) together imply that here exists $K > 0$ such that for all $G \subseteq F_j^i$ with $|G| \le |F_j^i|/2$, we have $\sigma_G \le K \cdot c_0 [G]_{F_j^i}$. Since our goal is to obtain a result about the rate of convergence, we can re-scale the initial capacity c_0 by a positive constant without loss of generality. Hence we can assume that the following condition is satisfied:

(K') For all $G \subseteq F_j^i$ with $|G| \le |F_j^i|/2$, we have $\sigma_G \le c_0 [G]_{F_j^i}$.

Roadmap. Our proof constructs an incentive-compatible risk-sharing arrangement in several steps. Fix n, and consider the decomposition described above. We begin by constructing an "unconstrained" risk-sharing arrangement that implements equal sharing in each set F_j^n, $j = 1, ..., \infty$, but does not necessarily satisfy the capacity constraints. We compute the implied typical link use of this transfer arrangement for each link, and choose n and c such that capacity constraints are satisfied most of the time. This arrangement results in exponentially small $SDISP$. We then bound the contribution of non-typical shocks to $SDISP$ and combine these terms to obtain the result stated in the proposition.

Iterative logic. The unconstrained arrangement is constructed by first smoothing consumption within each F_j^1 set; then smoothing consumption within each F_j^2 set; and so on. When $i = 1$, all sets are singletons, so there is no need to smooth within a set. Now consider the step when we move from i to $i + 1$. As we have seen, by (N4) the number of i level sets in F_j^{i+1} is bounded by a positive constant K. To simplify notation, denote F_j^{i+1} by F, and denote the i-level sets $F_{j'}^i$ that are subsets of F by $F_1, ..., F_k$ where $k \le K$. We know from (N2) and (N3) that $F_1, ..., F_k$ partition F. We smooth consumption in F_j^{i+1} by first smoothing the total amount of resources currently present in F_1 through the entire set F; then smoothing the total amount currently in F_2 through the set F, and so on until F_k. Note that the total consumption in F_1 at this round is the same as the total endowment e_{F_1}, because in each round i, we smooth all endowments within an i-level set. Having completely smoothed resources in F_1 in the previous round, all agents in F_1 are currently allocated $e_{F_1}/|F_1|$ units of consumption.

Auxiliary network flow. To smooth consumption over F, we define an auxiliary network flow. This is a key step in the proof. For this flow, focus on the subgraph generated by F together with capacities c_0, and assume for the moment that each agent in F_1 has $\sigma_{F_1}/|F_1|$ units of the consumption good (so the total in F_1 is exactly σ_{F_1}), while each agent in $F \backslash F_1$ has zero. We will show that a flow respecting capacities c_0 can achieve equal sharing in F from this endowment profile; and then use this flow to construct an unconstrained flow implementing the desired sharing over F for arbitrary shock realizations.

To verify that equal sharing can be implemented in the above endowment profile, note that equal sharing can be implemented through some IC transfer if for each set $G \subseteq F$ the excess demand for goods does not exceed the perimeter relative to F (this is where the key perimeter/area condition (K) plays it's role). What is this excess demand? Since we want equal sharing, we should allocate $\sigma_{F_1}/|F|$ to every agent in G. But those agents in G who are also in F_1 each have $\sigma_{F_1}/|F_1|$. So the excess demand for goods in the set G is

$$ed(G) = |G| \cdot \frac{\sigma_{F_1}}{|F|} - |G \cap F_1| \cdot \frac{\sigma_{F_1}}{|F_1|}. \tag{9}$$

Applying Theorem 1 to the finite network F, there is a feasible flow if for every G, we have $|ed(g)| \leq c_0 [G]_F$. To check that this holds, first assume that $|G|/|F| \geq |G \cap F_1|/|F_1|$; then the first term in (9) is larger, and hence $|ed(G)| \leq \sigma_{F_1} \cdot |G|/|F|$. From (P4) we have $\sigma_{F_1} \leq \sigma_F$, implying $|ed(G)| \leq \sigma_F \cdot |G|/|F|$. Now (P5) implies $\sigma_F/|F| \leq \sigma_G/|G|$, and hence $|ed(G)| \leq \sigma_G$. Now recall the key condition (K') that $\sigma_G \leq c_0 [G]_F$; it follows that $|ed(G)| \leq c_0 [G]_F$ as desired. We now check that (9) also holds when $|G|/|F| < |G \cap F_1|/|F_1|$. In this case, the second term in (9) dominates, and hence $|ed(G)| \leq \sigma_{F_1} \cdot |G \cap F_1|/|F_1|$. Since $\sigma_{F_1}/|F_1| \leq \sigma_{G \cap F_1}/|G \cap F_1|$ by (P3), we can bound the right hand side by $\sigma_{G \cap F_1}$, which satisfies $\sigma_{G \cap F_1} \leq \sigma_G \leq c_0 [G]_F$ again verifying that $|ed(G)| \leq c_0 [G]_F$. This shows that the auxiliary flow can be implemented.

Smoothing with auxiliary flow. Denote the transfers associated with the auxiliary flow by t_1. To smooth the consumption of F_1 over F for arbitrary shocks, we just use the transfers $t_1 \cdot (e_{F_1}/\sigma_{F_1})$; that is, we scale up the above flow with the actual size of the shock in F_1. This works, because t_1 was constructed to smooth a shock of exactly one standard deviation

σ_{F_1}. Extending this logic, to smooth the endowment of each other F_j through the set F, we construct auxiliary flows t_2, ..., t_k analogously, and implement the total transfer given by $t_1 \cdot e_{F_1}/\sigma_{F_1} + ... + t_k \cdot e_{F_k}/\sigma_{F_k}$. This construction results in an unconstrained flow which smooths consumption in the entire set F.

Note that while we used the capacities to construct the flow (this is how we got $t_1,...,$ t_k), the actual flow is a stochastic object that may violate some capacity constraints, both because it is scaled by e_{F_1}/σ_{F_1} and because it is summed over all j.

Iteration. We do the above step for all $i + 1$-level sets F_j^{i+1}; this concludes round $i + 1$ of the algorithm. Then we go on to round $i + 2$, and so on, until finally we implement equal sharing in each of the highest-level sets F_j^n, $j = 1,...,\infty$. Denote the unconstrained arrangement obtained in this way by t^U.

How low is $SDISP$ in the arrangement t^U? To answer, recall that (N4) implies $\left| F_j^n \right| \geq \underline{\gamma}^n$, and (P3) implies $\sigma_F / |F| \leq K \cdot |F|^{-1/2}$, so that $SDISP \leq K \cdot \underline{\gamma}^{-m/2} = K_1 \cdot \exp\left[-K_2 m \right]$. This $SDISP$, however, is implemented with an unconstrained flow; and now we want to assess how often the flow violates capacity constraints once we choose c and m. To do this, we need to compute the distribution of the flow over each link in the network.

Link usage. Consider the step where we smooth the consumption of F_1 over the entire set F using the flow $t_1 \cdot e_{F_1}/\sigma_{F_1}$. Fix some (u, v) link; then the use of this link in the flow at this round is $t_1(u, v) \cdot e_{F_1}/\sigma_{F_1}$. This is a random variable with mean zero and standard deviation $t_1(u, v)$, since e_{F_1}/σ_{F_1} has unit standard deviation. Moreover, we know that $t_1(u, v) \leq c_0(u, v)$ because this is how t_1 was constructed (this is why it was important to construct t_1 such that it satisfies the capacity constraints c_0.) It follows from Lemma 1 that the standard deviation of link use at this step is at most $c_0(u, v)$.

Now consider link use as we smooth the consumption of all sets F_1, ..., F_k over the set F. As we have seen, smoothing for each of these sets implies adding a flow over the (u, v) link that has standard deviation of at most $c_0(u, v)$. Given that $k \leq K$ for some constant, the total standard deviation of the flow over (u, v) in each round of the algorithm is at the most $K \cdot c_0(u, v)$. Adding up these flows over all n rounds shows that the total standard deviation of the unconstrained arrangement over the (u, v) link is at most $nK \cdot c_0(u, v)$.

Constrained arrangement. We construct an arrangement which satisfies the capacity

9

constraints in a simple way. We fix c and n, and for each agent u, try to implement his inflows and outflows according to the unconstrained flow we just constructed. If this is not possible, then we just implement as much of the prescribed flows as possible. This approach ensures that binding capacity constraints do not propagate down the network.

Bounding exceptional event. Denote $F_j^n = F$, and consider some agent $u \in F$. We begin bounding the exceptional event by looking at those realizations where the capacity constraint binds on exactly one of u's links: $t^U(u, v) > c \cdot c_0(u, v)$. We explore the effect of multiple binding constraints later. We focus on the contribution of these realizations to $\mathrm{E}x_u^2$, recalling that $SDISP$ is the square root of the average of this quantity over all agents u. The contribution of realizations where $t^U(u, v) > c \cdot c_0(u, v)$ but the other constraints of u do not bind to $\mathrm{E}x_u^2$ is at most

$$\int_{t^U(u,v)>c\cdot c_0(u,v)} \left[\bar{e}_F + t(u, v) - c(u, v)\right]^2 \, dP$$

where $\bar{e}_F = e_F / |F|$, the integral is taken over the probability space on which all random variables are defined and P is the associated probability measure. Noting that $(x + y)^2 \leq 3(x^2 + y^2)$, we can bound this from above by

$$3 \int \bar{e}_F^2 \, dP + 3 \int_{t^U(u,v)>c\cdot c_0(u,v)} \left[t(u, v) - c \cdot c_0(u, v)\right]^2 \, dP. \tag{10}$$

Here the first term is proportional to the variance of the unconstrained flow, which, as we have seen, is exponentially small. Thus we have to bound the contribution of the second term.

Large deviations. Let $z = \sum_j \alpha_j y_j$ for some α_j satisfying $\sum \alpha_j^2 < \infty$. Then, for any $c > 0$ and $\theta > 0$,

$$\Pr[z > c] \leq \mathrm{E}\exp\left[\theta(z - c)\right] = e^{-\theta c}\mathrm{E}\exp\left[\theta \sum \alpha_j y_j\right] = e^{-\theta c}\prod_j \mathrm{E}\exp\left[\theta \alpha_j y_j\right].$$

Now we can bound the last term using (P1) to obtain

$$\Pr[z > c] \leq e^{-\theta c}\prod_j \mathrm{E}\exp\left[K\alpha_j^2\theta^2/2\right] = e^{-\theta c}\mathrm{E}\exp\left[K\theta^2/2 \cdot \sum \alpha_j^2\right].$$

This holds for any θ, in particular, for $\theta = c/\left(K\sum\alpha_j^2\right)$, resulting in the bound $\Pr\left[z > c\right] \leq \exp\left[-c^2/\left(2K\sigma_z^2\right)\right]$, where we used the fact that the variance of z is $\sigma_z^2 = \sum\alpha_j^2$. This shows that the tail probabilities of z can be bounded by a term exponentially small in $(c/\sigma_z)^2$, just like in the case when z is normally distributed.

Bound on remaining variance. Using the bound on the tail probability, we can estimate the final term in (10). Let $z = t^U(u, v)$ which is a weighted sum of the y_j shocks by construction. Denoting the c.d.f. of z by $H(z)$ we have

$$\int_{t^U(u,v)>c\cdot c_0(u,v)} [t(u,v) - c\cdot c_0(u,v)]^2\ dP = \int_{z=c\cdot c_0(u,v)}^{\infty} (z - c\cdot c_0(u,v))^2\, dH(z)$$

$$= -\int_{z=c\cdot c_0(u,v)}^{\infty} (z - c\cdot c_0(u,v))^2\, d[1 - H(z)] =$$

$$= -\left[(z - c\cdot c_0(u,v))^2(1 - H(z))\right]_{c(u,v)}^{\infty} + \int_{z=c\cdot c_0(u,v)}^{\infty} 2(z - c\cdot c_0(u,v))[1 - H(z)]\, dz$$

where we integrated by parts. The above argument with large deviations proves $1 - H(z) \leq \exp\left[-z^2/2K\sigma_z^2\right]$. This implies that the first term is zero, and combining it with the second term, direct integration shows that

$$\int_{t^U(u,v)>c\cdot c_0(u,v)} [t(u,v) - c\cdot c_0(u,v)]^2\ dP \leq K'c\cdot c_0(u,v)\exp\left[-c^2\cdot c_0(u,v)^2/2K\sigma_z^2\right]$$

for appropriate constants K and K'.

Since $\sigma_z \leq nKc_0(u,v)$, the last term is bounded by $K\cdot\exp\left[-K'\cdot(c/n)^2\right]$, where the values of the constants are now different.

Combine bounds. We have obtained a bound on the exceptional event where the capacity constrained on a single link is binding. We must similarly bound the contribution to $\mathrm{E}x_u^2$ of binding capacity on all other single links of u; all possible pairs of links; all possible sets of three links; and so on. Since u has a bounded number of links, doing this just increases the bound we just obtained by a constant factor. In total, all exceptional events thus contribute to $\mathrm{E}x_u^2$ at most $K\cdot\exp\left[-K'\cdot(c/n)^2\right]$.

To obtain a bound on $SDISP$, we first bound $DISP = SDISP^2$, which is just the

average of $\mathrm{E}x_u^2$ over the entire network. We have seen that for each u,

$$\mathrm{E}x_u^2 \leq K_1 \cdot \exp\left[-K_2 n\right] + K_3 \cdot \exp\left[-K_4 \cdot (c/n)^2\right]$$

where the first term is the variance of the unconstrained flow and the second term is the bound coming from exceptional events. Setting $n = c^{2/3}$ yields $\mathrm{E}x_u^2 \leq K_5 \cdot \exp\left[-K_6 \cdot c^{2/3}\right]$, as desired.

Proof of Corollary 1

For this proof we also construct an informal risk-sharing arrangement step by step. The logic of the proof is to fix a grid associated with the geographic embedding, show that inside grid squares risk-sharing is good because the embedding is local and there are only a bounded number of people, and use the result for the plane to show that insurance is good across squares.

Fix the geographic embedding, and consider the grid with step size A for which the no separating avenues condition holds: for this grid, there is at least capacity $K > 0$ between any pair of adjacent squares under c_0. Since capacities are bounded away from zero, after re-scaling we can assume that all link capacities are at least 1; in this case all neighboring squares have connecting flow of at least 1 as well in c_0. Index the squares in the grid by $j = 1, ..., \infty$ and denote the set of agents in square j by G_j.

We have to accomplish good risk-sharing inside each square as well as across the squares. We will do this by using a share of the capacity of each link for within square sharing, and the remaining capacity for cross-square sharing. By locality of the embedding, any two agents in a given square are connected through a path lies within a bounded distance from the square. Assign, for each pair of agents inside a square one such path. By evenness, any link in the network is used by at most a bounded number of such paths. Let K^* be large enough such that all links are used by no more than K^* paths (K^* will denote this fixed quantity for the rest of the proof.)

Now fix $c > 0$, and use a share $1/(10K^*)$ of capacities to implement between-squares risk-sharing using Proposition 2, taking e_{G_j} as the "endowment shocks" of the squares. The conditions of the proposition are easily seen to be satisfied, and hence we obtain between-

squares dispersion which is exponentially small in $c^{2/3}$.

Second, we have to smooth the incoming and outgoing transfers for each square. Use a share 4/10 of capacities to smooth all incoming and outgoing transfers of each square. To do this, we need to use the paths connecting agents. Since the perimeter of each square used for incoming and outgoing transfers is $4c/(10K^*)$, and each link is used for at most K^* connecting paths, a total capacity of $4c/(10K^*) \cdot K^* = 4c/10$ will be sufficient to completely share the incoming and outgoing transfers among agents inside each square.

Third, we also have to smooth the total endowment shock realized in each square. To do this, first note that for any network of bounded size where capacities are bounded below and endowment shocks satisfy (P1) and (P2), the large deviations argument of the previous proof imply that $SDISP$ can be bounded by $K \exp[K' \cdot c^2/2]$. Since the number of agents in a square are bounded and shocks satisfy (P1) and (P2), and all pairs of agents are connected by (potentially external) paths of remaining capacity $5c/(10K^*)$ or more, it follows that we can achieve within-square dispersion on the order of $\exp[-K' \cdot c^2/2]$ This is of smaller order than the main $\exp[-K''c^{2/3}]$ term; hence the proof is complete.

Proof of Proposition 3

By definition

$$\beta = \frac{\text{cov}[x_F, e_F]}{\text{var}[e_F]} = \frac{\text{cov}[e_F - t_F, e_F]}{\text{var}[e_F]} = 1 - \frac{\text{cov}[t_F, e_F]}{\text{var}[e_F]}$$

where t_F denotes the total transfer leaving F. Moreover, $|\text{cov}[t_F, e_F]| = \sigma(t_F) \cdot \sigma(e_F) \cdot |corr[t_F, e_F]| \leq c[F] \cdot \sigma(e_F)$, using Lemma 1, and the claim of the Proposition follows.

Proof of Proposition 5

Fix realization e, and let t denote the vector of transfers over all links in a given IC arrangement. Denote the planner's objective with a given set of weights λ_i by $V(t) = \sum_i \lambda_i U_i \left(e_i - \sum_j t_{ij}, c_i\right)$. Then the planner's maximization problem can be written as $\max_t V(t)$ subject to $t_{ij} \leq c(i,j)$ and $t_{ij} = -t_{ji}$ for all i and j. It is easy to see that Karush-Kuhn-Tucker first order conditions associated with this problem are those given in the Proposition. Since we have a concave maximization problem where the inequality constraints are linear,

the Karush-Kuhn-Tucker conditions are both necessary and sufficient for characterizing a global maximum. For uniqueness, rewrite the planner's objective as a function of the consumption profile x, $\overline{V}(x) = V(t)$. This function is strictly concave in x and maximized over a convex domain, and hence the maximizing consumption allocation is unique, although the transfer profile supporting it need not be.

Proof of Proposition 6

For each i and j, say that i and j are in the same equivalence class if there is an $i \to j$ path such that for all agents l on this path, including j, we have $\lambda_i U_i' = \lambda_l U_l'$. The partition generated by these equivalence classes is the set of risk-sharing islands W_k. If $i \in W_k$ and $j \notin W_k$, then either $c(i,j) = 0$, in which case $t_{ij} = c(i,j)$ by definition, or $c(i,j) > 0$, which implies that $\lambda_i U_i' \neq \lambda_l U_l'$ by construction of the equivalence classes. But then Proposition 5 implies that $|t_{ij}| = c(i,j)$, as desired.

Proof of Proposition 7

In this proof we focus on transfer arrangements that are acyclical, i.e., have the property that after any endowment realization there is no path of linked agents $i_1 \to i_k$ such that $i_1 = i_k$, and $t_{i_l i_{l+1}} > 0 \; \forall \, l \in \{1, ..., k-1\}$. This is without loss of generality, as it is easy to show that for any IC arrangement there is an outcome equivalent acyclical IC arrangement that achieves the same consumption vector after any endowment realization.

(i): We begin with the weak inequalities of the claim $(x_j(e') \leq x_j(e) \; \forall \, j)$, which we establish in a slightly more general setup. Say that a transfer arrangement is monotone over all sets if for any $F \subseteq W$ and any two endowment realizations (e) and (e') such that $e_i' \leq e_i$ for all $i \in F$ and $t_{ji}' \leq t_{ji}$ for all $i \in F$ and $j \notin F$, we have $x_i' \leq x_i$ for all $i \in F$. Monotonicity over all sets means that for any set of agents F, reducing their endowments and/or their incoming transfers weakly reduces everybody's consumption. Note that this property indeed implies monotonicity in the sense of the Proposition, by taking $F = W$.

Fix a constrained efficient arrangement, and suppose it is not monotone over all sets. Let F be a set where this property fails, and fix a connected component of the subgraph spanned by F that contains an agent i such that $x_i' > x_i$. Let S be the set of agents for whom $x_i' \leq x_i$, and T be the set of agents for whom $x_i' > x_i$ in this component. S is

non-empty, because the total endowment available in any connected component of F has decreased, and T is non-empty by assumption. In addition, there exist $s \in S$ and $t \in T$ such that $t'_{st} > t_{st}$, because consumption in T is higher under e' than under e. But $t'_{st} > t_{st}$ implies $c(s,t) > t_{st}$ and $c(t,s) > t'_{ts}$, and hence, by Proposition 5, $\lambda_s U'_s(x_s) \geq \lambda_t U'_t(x_t)$ in e, and also $\lambda_s U'_s(x'_s) \leq \lambda_t U'_t(x'_t)$ in e'. Since $x'_t > x_t$ by assumption, strict concavity implies $\lambda_t U'_t(x'_t) < \lambda_t U'_t(x_t)$, which, combined with the previous two inequalities, yields $\lambda_s U'_s(x'_s) < \lambda_s U'_s(x_s)$. But this implies $x_s < x'_s$, which is a contradiction.

Finally, the claim that $x'_j < x_j$ for all $j \in \widehat{W}(i)$ follows directly from this monotonicity condition combined with (ii) which is proved below.

(ii): Let $\widehat{L}_i$ denote the set of links connecting agents in $\widehat{W}(i)$. Let L_i denote the set of links connecting agents in $W(i)$. Let t be a transfer arrangement respecting the capacity constraints and achieving $x(e)$ at endowment realization e, such that $t_{kl} < c(k,l) \ \forall \ (k,l) \in \widehat{L}_i$. In words, in transfer arrangement t, the capacity constraints for all links in $\widehat{L}_i$ are slack. Such a t exists by the definition of $\widehat{W}(i)$. Let b be the minimum amount of slackness on a link in $\widehat{L}_i$: $b = \min_{(k,l) \in \widehat{L}_i} (c(k,l) - |t_{kl}|)$.

Let L'_i denote the set of links connecting agents in $W(i)$ with agents in $W \backslash W(i)$. For every $(k,l) \in L'_i$, let t'_{kl} be such that $\lambda_k U'_k(x_k(e) - t'_{kl}) = \lambda_l U'_l(x_l(e) + t'_{kl})$. In words, t'_{kl} is the amount of transfer between k and l that would equate the weighted marginal utilities of k and l. By Proposition 5 and by the definition of $W(i)$, $t'_{kl} \neq 0 \ \forall \ (k,l) \in L'_i$. Let b' be the minimum amount of transfer that would equate the weighted marginal utilities of an agent in $W(i)$ and a neighboring agent outside $W(i)$: $b' = \min_{(k,l) \in L'_i} |t'_{kl}|$.

We claim that the result holds for $\Delta = \min(b, b')$, that is whenever $|e_i - e'_i| < \min(b, b')$, we have $\lambda_j U'_j(x_j(e')) = \lambda_i U'_i(x_i(e')) \ \forall \ j \in \widehat{W}(i)$, and $U_j(x_j(e')) = U_i(x_i(e)) \ \forall \ j \notin W(i)$. To see this, consider the restricted set of agents $W(i)$, and endowments $x_i(e) + e'_i - e_i$ for agent i, and $x_j(e)$ for $j \in W(i)/\{i\}$ (where $x_i(e)$ still refers to the constrained efficient allocation given set of agents W and endowment realization e). Let $x^{e,e'}$ denote this endowment vector on $W(i)$. Consider now the consumption arrangement over $W(i)$ that maximizes $\sum_{j \in W(i)} \lambda_j U_j(x_j)$ subject to x being achievable from $x^{e,e'}$ by transfer scheme t' (over $W(i)$) for which $|t_{jj'} + t'_{jj'}| \leq c(j, j') \ \forall \ j, j' \in W(i)$. Let this arrangement be denoted by $x^{W(i)}$. Because $\lambda_j U'_j(x_j)$ is decreasing in x_j for all j, $|x^{W(i)} - x_i(e)| \leq |e_i - e'_i|$. Then

15

there is a transfer scheme t' over $W(i)$ that achieves $x^{W(i)}$ from endowments $x^{e,e'}$, for which $|t'_{jj'}| \leq |e_i - e'_i| < \Delta$. Since $\Delta < b$, all the capacity constraints in $\widehat{L}_i$ are still slack. By Proposition 5 this means that $\lambda_j U'_j(x_j^{W(i)}) = \lambda_i U'_i(x_i^{W(i)})$. Moreover, since $\Delta < b'$, all the capacity constraints in L'_i are still binding, in the same direction. Extend now $x^{W(i)}$ to W such that $x_j^{W(i)} = x_j(e)$ for $j \in W \backslash W(i)$. Similarly, extend transfer scheme t' to W such that $t'_{jj'} = 0$ whenever at least one of j an j' are not in $W(i)$. Note that $t + t'$ is a direct transfer arrangement on W which meets the capacity constraints, and that $x^{W(i)}$ satisfies the conditions of Proposition 5. Hence $x^{W(i)}$ is the constrained efficient allocation given endowment realization e', and as shown above, satisfies the claims in (ii).

(iii): Let t' be an acyclical transfer arrangement achieving $x(e')$ after endowment realization e'. Then we can decompose t' as the sum of acyclical transfer arrangements t and t'' such that t achieves $x(e)$ after endowment realization e. By part (i) above, $x_{j'}(e') \leq x_{j'}(e)$ $\forall j' \in W$, implying that $MUC_{j'} \geq 1$ $\forall j' \in W$. Therefore if $x_j(e') = x_j(e)$, hence $MUC_j = 1$, then the statement in the claim holds. Assume now that $x_j(e') < x_j(e)$. Since $x_{j'}(e') \leq x_{j'}(e)$ $\forall j' \in W$ by part (i), for any $j' \in W \backslash \{i\}$ it must hold that the sum of transfers received by j' in transfer arrangement t'' is non-positive: $\sum_{l \in W \backslash \{j'\}} t''_{lj'} \leq 0$. Hence, only i can be a net recipient in the transfer arrangement t''. This, together with $x_j(e') < x_j(e)$ implies that there is a $j \to i$ path such that $t''_{i_m i_{m+1}} > 0$ along the path. Hence, in transfer scheme t no link (i_m, i_{m+1}) along the above $j \to i$ path is blocked, implying $\lambda_{i_{m+1}} U'_{i_{m+1}}(x_{i_{m+1}}(e)) \leq \lambda_{i_m} U'_{i_m}(x_{i_m}(e))$, and that no link (i_{m+1}, i_m) along the reverse $i \to j$ path is blocked, implying $\lambda_{i_{m+1}} U'_{i_{m+1}}(x_{i_{m+1}}(e')) \geq \lambda_{i_m} U'_{i_m}(x_{i_m}(e'))$. Dividing these inequalities yields the result.

A-2   Microfoundations for link-level punishment

Consider the following multi-stage game.

Stage 1. An endowment vector e is drawn from a commonly known prior distribution.

Stage 2. Each agent i makes a transfer t^e_{ij} to every neighbor j. Transfer t^e_{ij} is only observed by players i and j.

Stage 3. Agents play friendship games over links. The game over the (i, j) link is

	C	D
C	$c(i,j)$ $c(i,j)$	-1 $c(i,j)/2$
D	$c(i,j)/2$ -1	0 0

which is a coordination game with two pure strategy equilibria, (C, C) and (D, D). Denote the payoff of i from the game with j by $c'(i, j)$.

Stage 4. The realized utility of agent i is $U_i(x'_i, c'_i)$.

Proposition 10 *An allocation $x(e)$ is the outcome of a pure-strategy subgame-perfect equilibrium of this game if and only if it can implemented through an incentive-compatible informal risk-sharing arrangement.*

Proof. Fix an incentive-compatible informal risk-sharing arrangement and consider the following strategy profile σ. In Stage 2, each agent is supposed to make the transfer according to the above arrangement. In Stage 3, the neighbors across links where transfers were made as prescribed coordinate on the high equilibrium (C, C) and otherwise they coordinate on the low equilibrium (D, D). It is easy to see that making the promised transfers is an SPE. Conversely, consider a pure strategy SPE, and the corresponding risk-sharing arrangement it induces. Note that in any such profile, in stage 3 any two neighbors should either play (D, D), resulting in a payoff of $(0, 0)$, or play (C, C), resulting in a payoff of $(c(i, j), c(i, j))$. But then all transfers in Stage 2 have to satisfy the IC constraints because the actual transfer from i to j can only influence the continuation strategy of j, not agents in $W/\{i, j\}$ (since they do not observe the actual transfer). Therefore the actual transfer from i to j can only influence the payoff i gets from the friendship game with j, not the payoff from other friendship games he is involved at in Stage 3. Hence the maximum loss in Stage 3 payoffs in a pure SPE when not delivering a promised transfer t_{ij}^e is $c(i, j)$, the difference between the best Nash equilibrium payoffs of the friendship game $(c(i, j))$ and the payoff that a player can guarantee in the friendship game (0). This implies that the transfer scheme has to be IC. ∎

A-3 Background on the theory of network flows

The following concepts from the theory of network flows are useful for many of the proofs in the paper. Cormen, Leiserson, Rivest and Stein (2001) provides a more careful treatment. Fix a finite graph G two nodes s and t (for "source" and "target") and a capacity c.

Definition 3 *An $s \rightarrow t$ flow with respect to capacity c is a function $f : G \times G \rightarrow \mathbb{R}$ which satisfies*

- *(i) Skew symmetry: $f(u, v) = -f(v, u)$.*
- *(ii) Capacity constraints: $f(u, v) \leq c(u, v)$.*
- *(iii) Flow conservation: $\sum_w f(u, w) = 0$ unless $u = s$ or $u = t$.*

A useful physical analogue is to think about a flow as some liquid flowing through the network from s to t, which must respect the capacity constraints on all links. The value of a flow is the amount that leaves s, given by $|f| = \sum_w f(s, w)$. The maximum flow is the highest feasible flow value in G. Flows are particularly useful in our setting, because the capacity constraints associated with our direct transfer representation are exactly the constraints (ii) in the above definition. In particular, a direct transfer representation that meets the capacity constraints is called a circulation in the computer science literature.

Definition 4 *A cut in G is a disjoint partition of the nodes into two sets $G = S \cup T$ such that $s \in S$ and $t \in T$. The value of the cut is the sum of $c(u, v)$ for all links such that $u \in S$ and $v \in T$.*

It is easy to see that the maximum flow is always less than or equal to the minimum cut value. The following well-known result establishes that these two quantities are equal.

Theorem 2 *[Ford and Fulkerson, 1958] The maximum flow value equals the minimum cut value.*

We rely both on the concept of network flows and the maximum flow - minimum cut theorem in the proofs of the paper.

A-4   Formal results for subsection 3.3 of the paper

A decentralized exchange implementing any constrained efficient arrangement.
We show that for any constrained efficient allocation, there exists a simple iterative procedure
that only uses local information in each round of the iteration, and converges to the allocation
as the number of iterations grow. A simpler version of this procedure, with equal welfare
weights and no capacity constraints, was proposed by Bramoulle and Kranton (2006). The
basic idea is to equalize, subject to the capacity constraints, the marginal utility of every
pair of connected agents at each round of iteration. This procedure can be interpreted as a
set of rules of thumb for behavior that implements constrained efficiency in a decentralized
way.

Fix an endowment realization e, and denote the efficient allocation corresponding to
welfare weights λ_i by x^*. Fix an order of all links in the network: $l_1,...,l_L$, and let i_k and
j_k denote the agents connected by l_k. To initialize the procedure, set $x_i = e_i$ and $t_{ij} = 0$
for all i and j. Then, in every round $m = 1, 2, ...$, go through the links $l_1, ..., l_L$ in this
order, and for every l_k, given the current values x_{i_k}, x_{j_k}, and $t_{i_k j_k}$, define the new values x'_{i_k}
and x'_{j_k} and $t'_{i_k j_k} = t_{i_k j_k} + x'_{j_k} - x_{j_k}$ such that they satisfy the following two properties: (1)
$x'_{i_k} + x'_{j_k} = x_{i_k} + x_{j_k}$. (2) Either $\lambda_{i_k} U'_{i_k}(x'_{i_k}) = \lambda_{j_k} U'(x'_{j_k})$, or $\lambda_{i_k} U'_{i_k}(x'_{i_k}) > \lambda_{j_k} U'_{j_k}(x'_{j_k})$ and
$t'_{i_k j_k} = -c(i,j)$, or $\lambda_{i_k} U'_{i_k}(x'_{i_k}) < \lambda_{j_k} U'_{j_k}(x'_{j_k})$ and $t'_{i_k j_k} = c(i,j)$. This amounts to the agent
with lower marginal utility helping out his friend up to the point where either their marginal
utility is equalized, or the capacity constraint starts to bind. Once this step is completed
for link k, we set $x = x'$ and $t = t'$ before moving on to link $k + 1$. For $m = 1, 2, ...$ let x_i^m
denote the value of x_i, and let t_{ij}^m denote the value of t_{ij}, at the end of round m. Note that
x_m meets the capacity constraints by design for every m.

Proposition 11   *If consumption and friendship are perfect substitutes, then $x^m \to x^*$ as*
$m \to \infty$.

Proof. Let $V(x)$ denote the value of the planner's objective in allocation x. The above
procedure weakly increases $V(x)$ in every round and for every link l_k. Hence $V(x_1) \leq$
$V(x_2) \leq ..$, and since $V(x) \leq V(x^*)$ for all x that are IC, we have $\lim_{m \to \infty} V(x_m) = V \leq$
$V(x^*)$. Since the set of IC allocations is compact, and x_m is IC for every m, there exists a

convergent subsequence of x_m, with limit x and associated transfers t. Clearly, $V(x) = V$. If $V = V^*$ then $x = x^*$ since the optimum is unique. If $V < V^*$, then x is not optimal, and hence does not satisfy the first order condition over all links. Let l_k be the first link in the above order for which the first order condition fails in x and t. Then there is a transfer meeting the capacity constraints at x that increases the planner's objective by a strictly positive amount δ. But this means that for every x_m far along the convergent subsequence, the planner's objective increases by at least $\delta/2$ at that round, which implies that $V(x_m)$ is divergent, a contradiction. Hence $\lim x_m = x^*$ along all convergent subsequences, which implies that x_m itself converges to x^*. ∎

Ex ante coalition-proofness of constrained efficiency. If we require stability with respect to both ex ante and ex post coalitional deviations we obtain a subset of the constrained efficient agreements. In the case of perfect substitutes this subset is exactly the set of constrained efficient agreements. We say that a coalition-proof agreement x admits no ex ante coalitional deviations if there is no coalition S and coalition-proof risk-sharing agreement x'_S within S such that all agents in S weakly prefer losing all their links to agents in W/S and having agreement x'_S to keeping all their links and having agreement x, and at least one agent in S strictly prefer the former. Intuitively, an ex ante coalitional deviation implies that the agents of the deviating coalition leave the community (cut their ties with the rest of the community) and agree upon a new risk-sharing agreement among each other (using only their own resources).

Proposition 12 *A coalition-proof agreement that admits no profitable ex ante coalitional deviations is constrained efficient. If goods and friendship are perfect substitutes then the set of coalition-proof agreements that admit no profitable ex ante deviations is equal to the set of constrained efficient agreements.*

Proof. Consider first a coalition-proof agreement x that is not constrained efficient. Then there is another coalition-proof agreement x' that ex ante Pareto-dominates x. But then x' is a profitable ex ante coalitional deviation for coalition W. This concludes the first part of the statement.

Assume now that goods and friendship are perfect substitutes and consider a coalition-proof agreement x that is constrained efficient. Suppose there is coalition S and a profitable ex ante deviation x'_S by S. Theorem 1 implies that x can be achieved by a direct-transfer agreement t that respects all capacity constraints. Similarly, x'_S can be achieved by a direct transfer agreement t'_S within S that respects all capacity constraints (within S). Consider now a combined direct transfer agreement (t'_S, t_{-S}) that is equal to t'_S for links within S, and it is equal to t otherwise. Since both t and t'_S respect capacity constraints, so does (t'_S, t_{-S}), hence the resulting consumption profile x'' is coalition-proof. By construction x is equivalent to x'' for agents in $W \backslash S$. Agents in S are at least weakly better off with consumption profile x'' and not losing any of their links than with consumption profile x'_S and losing their links with agents in $W \backslash S$, since x'' is coalition-proof. But this, combined with x'_S being a profitable ex ante coalitional deviation, implies that coalition-proof agreement x'' Pareto-dominates x, which contradicts that x is constrained efficient. ∎

A-5 Analysis with imperfect substitutes

A-5.1 Formal results for subsection 4.1 of the paper

The extension of the results in Section 2 follows directly from the discussion in subsection 4.1. We now also show that with an increasing MRS, the set of IC arrangements contracts after a negative aggregate shock.

Proposition 13 *Assume that MRS_i is increasing in x_i for all i. Then for any pair of endowment realizations $\underline{e}$ and $\overline{e}$ such that $\underline{e}_i \leq \overline{e}_i$ for all i, an incentive compatible set of transfers in $\underline{e}$ is also incentive compatible given $\overline{e}$.*

Proof. Let $V(y_i, c_i; s_i) = U_i(y_i + s_i, c_i)$, then $(V_x/V_c)(y_i, c_i; s_i) = (U_x/U_c)(y_i + s_i, c_i)$, and hence the condition that $MRS_i = (U_x/U_c)(x_i, c_i)$ is increasing in x_i implies that $(V_x/V_c)(y_i, c_i; s)$ is increasing in s for any fixed (y_i, c_i), i.e., that $V(y_i, c_i; s)$ satisfies the Spence-Mirrlees single-crossing condition. Since U_i is continuously differentiable and U_x, $U_c > 0$, Theorem 3 in Milgrom and Shannon (1994) implies that V has the single crossing property. In particular, $V(y_i, c_i; 0) \geq V(y'_i, c'_i; 0)$ implies $V(y_i, c_i; s_i) \geq V(y'_i, c'_i; s_i)$ for any

$s_i \geq 0$, or equivalently, $U_i(x_i, c_i) \geq U_i(x_i', c_i')$ implies $U_i(x_i + s_i, c_i) \geq U_i(x_i' + s_i, c_i')$. It follows that for any $s_i \geq 0$, the compensating variation satisfies

$$CV_i(x_i, c_i, c_i') \leq CV_i(x_i + s, c_i, c_i')$$

and hence for any set F, we have $c^x[F] \leq c^{x+s}[F]$. Now denote $\overline{e} - \underline{e} = s \geq 0$; it follows immediately that any IC transfer scheme given $\underline{e}$ is IC given $\overline{e}$ as well. ∎

A-5.2 Formal results for subsection 4.2 of the paper

The equivalence between the planner's problem and constrained efficiency with general preferences and a concave MRS was established in Appendix A to the paper. To present our characterization result building on this equivalence, first we define a measure of marginal social welfare gain of transfers to agents. Fix an IC arrangement x, and recalling the definition of acyclical transfer arrangements from the proof of Corollary 7, let t be an acyclical implementation of x in endowment realization e. Consider the following iterative construction. We say that the IC constraint from i to j binds if $U_i(x_i, c_i) = U_i(x_i + t_{ij}, \widehat{c}_{i,j})$. Let $W^1 \subseteq W$ denote the set of agents i for whom (i) there is no j such that $c(i,j) > 0$; and (ii) the IC constraint from i to j binds. Since t is acyclical, W^1 is nonempty. For any $i \in W^1$, let $\Delta_i = \lambda_i U_{i,x}(x_i, c_i)$ be the marginal benefit of an additional dollar to i. This is both the private and social marginal welfare gain, because no IC constraint binds for transfers from i.

Suppose now that we have defined the sets $W^1, ..., W^{k-1}$ and the corresponding Δ_i for any $i \in \cup_{l \leq k-1} W^l$. Let W^k denote the set of agents i such that $i \notin \cup_{l \leq k-1} W^l$ but whenever $c(i,j) > 0$ and the IC constraint from i to j binds, $j \in \cup_{l \leq k-1} W^l$. To define Δ_i, first denote, for every j such that the IC constraint from i to j binds, $\widehat{x}_{i,j} = x_i + t_{ij}$, and $\widehat{c}_{i,j} = c_i - c(i,j)$, and let

$$\delta_{ij} = \lambda_i U_{i,x}(x_i, c_i) \cdot \frac{U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})}{U_{i,x}(x_i, c_i)} + \Delta_j \cdot \left[1 - \frac{U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})}{U_{i,x}(x_i, c_i)}\right].$$

As we will show below, δ_{ij} measures the marginal social gain of an additional dollar to i, under the assumption that i optimally transfers some of the dollar to j. Intuitively, to transfer

to j, i has to increase his own consumption somewhat to maintain incentive compatibility. More formally, we show below that a share $U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})/U_{i,x}(x_i, c_i)$ of the marginal dollar must be kept by i, and only the remaining share can be transferred to j, where it has a welfare impact of Δ_j. Denote $\delta_{ii} = \lambda_i U_{i,x}(x_i, c_i)$, and to account for the softening of the IC constraint over all links, let

$$\Delta_i = \max\{\delta_{ij} \mid j : \text{the IC constraint from } i \text{ to } j \text{ binds or } j = i\}.$$

With this recursive definition, the marginal social welfare of an additional dollar takes into account both the marginal increase in i's consumption, and the softening of the IC constraints which allow transfers of resources through a chain of agents.

Proposition 14 *[Constrained efficiency with imperfect substitutes] Assume that MRS_i is concave in x_i for every i. A transfer arrangement t is constrained efficient iff there exist positive $(\lambda_i)_{i \in W}$ such that for every $i, j \in W$ one of the following conditions holds:*
1) $\Delta_j = \Delta_i$
2) $\Delta_j > \Delta_i$ and the IC constraint binds for t_{ij}
3) $\Delta_j < \Delta_i$ and the IC constraint binds for t_{ji}.

Proof. We begin with some preliminary observations. Suppose that the IC constraint from i to j binds, and i receives an additional dollar. Suppose that i keeps a share α of the dollar and transfers the remaining $1 - \alpha$ such that the IC constraint continues to bind. Then it must be that $\alpha U_{i,x}(x_i, c_i) = U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})$, or equivalently, $\alpha = U_{i,x}(x_i, c_i)/U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})$. To maintain incentive compatibility, this share of the dollar has to be consumed by i, and only the remainder can be transferred to j.

Now we establish the necessity part of the proposition. Fix a constrained efficient arrangement, and let λ_i be the associated planner weights. Consider realization e. We first show that the marginal value to the planner of an additional dollar to an agent i is Δ_i. Let $i \in W^1$, then the marginal value to the planner of endowing i with an additional dollar is at least Δ_i. It cannot be larger, since that would imply that transferring a dollar away

23

from i increases social welfare in the original allocation, contradicting constrained efficiency. Hence, the marginal social value of a dollar to i is exactly Δ_i. Suppose we established for all $j \in \cup_{l \leq k-1} W^l$ that the marginal social value of a dollar to j is Δ_j. Let $i \in W^k$. For any j such that the IC constraint from i to j is binding, Δ_j is at least as large as the marginal social value of an additional dollar to i, because otherwise optimality requires reducing t_{ij}. Hence the marginal social value of a dollar to i is obtained when i transfers as much of the dollar as possible under incentive compatibility to some agent j. Given our above argument, i can transfer at most $1 - U_{i,x}(x_i, c_i)/U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})$ to j, hence the marginal welfare gain if he chooses to transfer to j will be δ_{ij}. Since i will choose to transfer the dollar to the agent where it is most productive, the marginal social gain will be the maximum of δ_{ij} over j, which is Δ_i.

It follows easily that if $\Delta_j > \Delta_i$ for some i, j, then the IC constraint for t_{ij} has to bind: otherwise social welfare could be improved by marginally increasing t_{ij}. This establishes that in a constrained efficient allocation, for any endowment realization and any pair of agents one of conditions (1)-(3) from the theorem have to hold.

For sufficiency, let now x denote the unique welfare maximizing consumption, let t be an IC transfer scheme achieving this allocation, and let $\widehat{\Delta}_i = \Delta_i(x, t)$, for every $i \in W$. Assume now that there exists another consumption vector $x' \neq x$ achieved by IC transfer scheme t' such that (x', t') satisfy conditions (1)-(3), and let $\Delta'_i = \Delta_i(x', t')$, for every $i \in N$. Then there exists an acyclical nonzero transfer scheme t^d that achieves x from x', and which is such that $t' + t^d$ is IC. By definition of x, t^d from x' improves social welfare. Let now $W^d = \{i \in W | \exists \, j \text{ such that } t^d_{ij} \neq 0\}$, and partition W^d into sets $W^d_0, ..., W^d_K$ the following way. Let $W^d_0 = \{i \in W^d | -\exists \, j \in W^d \text{ st. } t^d_{ij} > 0\}$. Given $W^d_0, ..., W^d_k$ for some $k \geq 0$, let $W^d_{k+1} = \{i \in W^d \backslash (\underset{l=0,...,k}{\cup} W^d_l) | -\exists \, j \in W^d \backslash (\underset{l=0,...,k}{\cup} W^d_l) \text{ st. } t^d_{ij} > 0\}$. Note that $x'_i > x_i \, \forall \, i \in W^d_0$, which together with there being no agent j such that $t^d_{i\,j} > 0$ implies that $\Delta'_i < \widehat{\Delta}_i$. Now we iteratively establish that $\Delta'_i < \widehat{\Delta}_i \, \forall \, i \in W^d$. Suppose that $\Delta'_i < \widehat{\Delta}_i \, \forall \, i \in \underset{l=0,...,k}{\cup} W^d_l$ for some $k \geq 0$. Let $i \in W^d_{k+1}$. Note that by definition there is $j \in \underset{l=0,...,k}{\cup} W^d_l$ such that $t^d_{i\,j} > 0$, and there is no $j' \in W^d \backslash (\underset{l=0,...,k}{\cup} W^d_l)$ such that $t^d_{i\,j'} > 0$. Suppose $\Delta'_i \geq \widehat{\Delta}_i$. This can only be compatible with $t^d_{i\,j} > 0$, $\Delta'_j < \widehat{\Delta}_j$, and (1)-(3) holding for both (x', t') and $(x, t' + t^d)$ if $x_i > x'_i$. But $x_i > x'_i$, and $\Delta'_{i'} < \widehat{\Delta}_{i'} \, \forall \, i' \in W$ such that $t^d_{ii'} > 0$ implies $\Delta'_i < \widehat{\Delta}_i$,

a contradiction. Hence $\Delta'_i < \widehat{\Delta}_i \; \forall \; i \in W^d_{k+1}$, and then by induction $\Delta'_i < \widehat{\Delta}_i \; \forall \; i \in W^d$. But note that for any $i \in W^d_K$ it holds that $x_i < x'_i$ and there is no $j \in W$ such that $t^d_{ji} > 0$, and hence $\Delta'_i > \widehat{\Delta}_i$. This contradicts $\Delta'_i < \widehat{\Delta}_i \; \forall \; i \in W^d$, hence there cannot be (x', t') satisfying (1)-(3) such that t' is IC and $x' \neq x$. ∎

Corollary 7 can also be extended to the imperfect substitutes case. Fix a constrained efficient arrangement, and let e and e' be two endowment realizations such that $e_i > e'_i$ for some $i \in W$, and $e_j = e'_j \; \forall \; j \in W \backslash \{i\}$. Let $x^*(e)$ be the consumption in the constrained efficient allocation after e. Analogously to the perfect substitutes case, let $\widehat{W}(i)$ the largest set of connected agents containing i such that all IC constraints within the set are slack given some transfer arrangement achieving the constrained efficient allocation after e_i. For any endowment realization e, let $\Delta_j(e)$ be Δ_j, as defined above, given any transfer scheme with the maximal number of links on which the IC constraints are slack, among the ones that attain the constrained efficient allocation. It is straightforward to show that there is a transfer scheme with a maximal number of links on which the IC constraints are slack, among the ones achieving the constrained efficient allocation, and that for all such transfer arrangements Δ_j is the same.

Corollary 2 *[Spillovers with imperfect substitutes] Assume that MRS_i is concave, then*

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

(ii) [Local sharing] There exists $\delta > 0$ such that $|e_i - e'_i| < \delta$ implies $\Delta_i(e') = \Delta_j(e')$ for all $j \in \widehat{W}(i)$.

(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, $\Delta_l(e') \geq \Delta_j(e')$.

The proof of this result is analogous to the perfect substitutes case and hence omitted. Note that (ii) is weaker than in Corollary 7, because even small shocks can spill over the boundaries of the risk-sharing islands of agent hit by the shocks. Also note that since $\Delta_i = \lambda_i U_{i,x}$ for any agent not on the boundary of an island, (i) implies that consumption is monotonic in the endowment realization for such agents.

A-6 Numerical methods

Risk-sharing simulations. We use the following numerical approach for the simulations underlying Figure 4. We assume throughout that endowment shocks are uniformly distributed with support $[-1, 1]$. We build on Theorem 1 and express a SDISP-minimizing incentive-compatible risk-sharing arrangement as a cost-minimizing flow as follows. (1) Create two artificial nodes s and t as in the proof of Theorem 1. (2) Divide the shock support into K equal intervals. For each agent i, denote the subinterval into which i's endowment falls by k_i (treating $[-1, -1 + 2/K]$ as the first interval and $[1 - 2/K, 1]$ as the Kth interval). Create k_i links between s and i such that each link has capacity $2/K$ in the direction from s to i and zero in reverse direction. Define the "cost" of a flow going from s to i across any of these links to be j for the jth link of out k_i links. Similarly, create $K - k_i$ links between t and i. such that each link has capacity $2/K$ in the direction from i to t and zero in reverse direction. Define the cost of a flow going from i to t across any of these links to be j for the jth link of out k_i links. (3) Use Edmonds and Karp's (1972) algorithm to calculate a cost-minimizing flow in this augmented network. This solution induces an incentive-compatible risk-sharing arrangement that maximizes a piecewise linear approximation to the quadratic utility function assumed in the definition of $SDISP$, where the marginal utility of consumption for any agent is constant within each of the K intervals. Simulations (not reported) show that this approximation generates highly accurate predictions for $K = 20$. For the results presented in the text we set $K = 100$.

Geographic network representation. The algorithm used in the geographic representation constructed in Figure 5 is the following. For each household i, we first construct vectors v_j to every other households j in the unit square using households' initial (re-scaled) geographic coordinates. We also calculate the length d_i of each of these vectors. Note, that the maximum distance between two households is $\sqrt{2}$. We then calculate a shift vector as the weighted sum $-\sum(\sqrt{2} - d_i)v_j/\|v_j\|$ and move each household in the direction of this shift vector. Shifts are larger if a household is closely surrounded by other households and the shift will push the household away from its neighbors. This procedure is repeated 23 times to obtain the representation in Figure 5E.