

Consumption Risk-sharing in Social Networks*

Attila Ambrus Markus Mobius Adam Szeidl
Harvard University Harvard University UC - Berkeley

November 2007

Abstract

We build a model of informal risk-sharing among agents organized in a social network. A connection between individuals serves as collateral that can be used to enforce insurance payments. We characterize incentive compatible risk-sharing arrangements for any network structure, and develop two main results. (1) Expansive networks, where every group of agents have a large number of links with the rest of the community relative to the size of the group, facilitate better risk-sharing. In particular, “two-dimensional” village networks organized by geography are sufficiently expansive to allow very good risk-sharing. (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, risk-sharing in second-best arrangements is local: socially closer agents insure each other more. In an application of the model, we explore the spillover effect of development aid on the consumption of non-treated individuals.

*E-mails: ambrus@fas.harvard.edu, mobius@fas.harvard.edu, szeidl@econ.berkeley.edu. We thank In Koo Cho, Erica Field, Drew Fudenberg, Eric Maskin, Stephen Morris, Gabor Pete and seminar participants for helpful comments and suggestions.

Households in developing countries are often exposed to substantial risk. Obtaining formal insurance against this risk can be difficult: due to weaknesses in the legal system, financial and insurance markets are often underdeveloped. To cope with this problem, households sometimes rely on informal risk-sharing arrangements, such as exchanging gifts or providing transfers and services to those in need. Evidence suggests that these informal arrangements frequently take place in the social network. For instance, Townsend (1994) emphasizes the importance of informal insurance networks in Indian villages; similarly, Udry (1994) documents that the majority of transfers take place between neighbors and relatives in Northern Nigeria. The prevalence of transfers in the developing world is also illustrated by Figure 1, which depicts the web of financial and asset transfers in a shantytown in Peru.

In this paper, we develop a model of informal risk-sharing that formalizes the role of the social network in providing contract enforcement. In our economy, agents organized in an exogenously given social network face endowment risk. To obtain insurance, these agents engage in an informal risk-sharing contract, which specifies a set of transfers contingent on the realization of uncertainty. This contract involves moral hazard, because ex post, individuals who are required to make transfers may prefer to deviate and withhold payment. Informal contract enforcement comes from the fact that failure to make a promised payment to a friend leads to losing the associated link. In turn, losing a link is costly, because agents derive utility from their remaining network connections. This utility value of links represents either the present value of future interaction or more direct “social” benefits, and can be used as *social collateral* to provide informal contract enforcement.

Our goal is to understand the degree and structure of informal insurance in social networks using this model. Our first main result is a characterization of the degree of risk-sharing that can obtain in a given network. We begin with identifying a property of networks – expansiveness, measured with the number of links that sets of agents have with others in the community, relative to the number of agents in the set – that facilitates good risk-sharing. To gain intuition about this property, consider the three example networks depicted in Figure 2. Among these networks, the infinite line in Figure 2a is the least expansive, because even large sets of consecutive agents have only two links with the rest of the community. Higher expansion is obtained in the infinite “plane” network of Figure 2b, where the “perimeter” of square shaped sets of agents grows with size, and yet more expansive is the infinite binary tree, where the perimeter of all sets grows at least proportionally with size. The connection between expansiveness and risk-sharing is intuitive:

having more links with the rest of the network allows for higher transfers, which makes it easier for every set of agents to dispose of their set-specific idiosyncratic risk.

To quantify the implications of expansiveness for insurance, we first show that perfect risk-sharing only obtains in highly expansive networks like the infinite binary tree. In these networks, the perimeter of every set is at least proportional to its size, which makes full sharing feasible even in the improbable event when all agents in a large set receive a negative income shock. However, this high level of expansiveness is unlikely to obtain in real-world networks: as Figure 1 illustrates, in practice networks of transfers and gifts are organized partly on the basis of geographic closeness, and hence are more similar to the plane. Motivated by this observation, we turn to explore partial risk-sharing in less expansive networks. We begin by showing that when shocks are not too correlated, the plane network allows for reasonably good risk-sharing. For an intuition, assume that endowment shocks are i.i.d. Independence implies that the standard deviation of the total endowment in any set of agents is proportional to the square root of set size. But on the plane, the perimeter of sets is also at least proportional to the square root of size, implying that “typical” shocks can pass through the perimeter. Thus most sets of agents can dispose of their idiosyncratic shocks, which results in reasonably good insurance. This logic also shows that risk-sharing is necessarily poor on the line, because the perimeter of interval sets is uniformly bounded.

What do these results imply about insurance in real-world networks like Figure 1? To address this question, we next consider a class of “geographic networks,” that have a map representation where agents tend to have connections in multiple directions. We show that the expansion properties of these networks are similar to the plane, and hence they allow for very good but imperfect risk-sharing. In particular, we predict reasonably good informal insurance for real-world villages like the one depicted in Figure 1, because two-dimensional village networks are likely to be “geographic.” This theoretical result is consistent with the empirical findings of Townsend (1994), Ogaki and Zhang (2001) and Mazzocco (2007), who document very good and in some cases perfect risk-sharing in Indian villages.

The above results constitute a quantitative analysis of the degree of informal risk-sharing. Our second main contribution is a qualitative analysis of insurance behavior in constrained efficient “second-best” arrangements. We show that in these arrangements, for every realization of uncertainty, the network can be partitioned into a set of endogenously organized connected components called “risk-sharing islands.” This partition has the property that shocks are completely shared within, but only imperfectly across islands. For an intuition, note that in each realization, island

boundaries are defined by links where agents are paying the maximum incentive compatible transfer amount. Higher transfers over these links are not incentive compatible, and hence insurance across island boundaries is limited; but links inside an island do allow for marginally higher transfers, explaining complete sharing within boundaries. In this partition the size and location of islands, and hence the set of agents who fully share each others' shocks, is endogenous to the endowment realization. This result differentiates our model from group-based theories of risk-sharing, where insurance groups are determined exogenously and do not vary with the realization of uncertainty.

One implication of the islands result is that risk-sharing in networks is local. The intuition is straightforward: because risk-sharing islands are connected subgraphs, agents who are socially closer are more likely to belong to the same island, and hence insure each other more. This observation helps characterize the mechanics of informal insurance as a function of shock size. Risk-sharing works well for relatively small shocks, because both direct and indirect friends help out. As the size of the shock increases, only a selected group of close friends help shoulder the additional burden; and risk-sharing completely breaks down for large shocks. This prediction can be used to test our model against other theories of limited risk-sharing, which do not imply differences in partial risk-sharing as a function of network distance.

In current work, we are exploring the interaction between government policies and network-based insurance by simulating the effects of a hypothetical development aid program using network data from Peru. Aid programs where some agents receive government transfers are common in the developing world (e.g., Progresa in Mexico). In our model, part of this aid will be transferred by informal arrangements through the network and therefore also affects the consumption of the non-treated, consistent with the empirical findings in Angelucci and De Giorgi (2007). Simulations allow us to better understand the mechanics of aid spillovers. We expect that the identities of the treated can matter for the overall impact of the program: well-connected individuals are better at allocating resources to those who need them most. Identifying observable demographic correlates such as gender or education that help targeting aid to these individuals can have implications for the optimal design of development aid.

Our work builds on recent theories of informal contracting where the network structure is explicitly modelled. The paper most closely related to ours is Bloch, Genicot, and Ray (2005), who build a model of informal insurance in social networks where agents face both informational and commitment constraints. Their main result is a characterization of network structures that are stable under certain exogenously specified risk-sharing arrangements. We conduct the opposite

investigation: taking the network as exogenous, we study the degree and structure of informal risk-sharing. Bramouille and Kranton (2006) and Bramouille and Kranton (2007) also study insurance arrangements in networks, but in their models there are no enforcement constraints. Mobius and Szeidl (2007) explore informal borrowing in networks with a model related to ours, and Dixit (2003) analyses the trade-off between relational and formal governance when agents are organized in a circle network.¹

Empirical research on insurance in networks includes Dercon and Weerdt (2006), Fafchamps and Lund (2003) and Fafchamps and Gubert (2007), who use data on village networks, while Mazzocco (2007) emphasizes the role of within caste-transfers. We also build on the influential early work of Mace (1991), Cochrane (1991), Townsend (1994) and Udry (1994) on limited risk-sharing.

The rest of this paper is organized as follows. The next Section presents our model of informal insurance in networks. Section 2 characterizes the limits to risk-sharing, and Section 3 analyses constrained efficient arrangements. We discuss briefly our current work on the indirect effects of development aid in Section 4, and conclude in Section 5.

1 A model of risk-sharing in the network

1.1 Setup

The basic logic of our model is the following. We consider an economy where agents face endowment risk and have no access to formal insurance markets. To reduce their risk exposure, agents agree on an informal risk-sharing arrangement, which is a set of state-contingent transfers to be paid after the realization of uncertainty. These transfer payments are used to implement risk-sharing by compensating those who experience bad shocks. The informal contract is enforced by the threat of social sanctions: agents keep their promises because failure to make a transfer payment leads to losing a valuable friend in the network.

More formally, consider a social network $G = (W, L)$ where W is the set of agents (vertices) and L is the set of links (edges) between agents. Each link in the network represents a friendship or business relationship. The strength of an (i, j) relationship is exogenous, and is denoted by

¹More broadly, our paper is related to the literature on informal contracting in repeated interactions. Ligon (1998), Coate and Ravallion (1993), Kocherlakota (1996) and Ligon, Thomas, and Worrall (2002) develop models of consumption insurance with limited commitment, but do not study the effects of network structure. In earlier work, Kandori (1992), Greif (1993) and Ellison (1994) study community enforcement, and Kranton (1996) analyses the interaction between relational and formal markets.

$c(i, j) \geq 0$, where we call c the capacity function. For ease of presentation, we assume that friendship is symmetric, so that $c(i, j) = c(j, i)$ for all i and j agents.² We think about the capacity $c(i, j)$ as a measure of the benefit that i derives from his relationship with j . These benefits can represent the direct utility that agents enjoy when they are in a social relationship, or the utility or monetary value of economic interaction in the present or in future periods.

Prior to the realization of uncertainty, agents agree on a transfer arrangement, which is a set of state-contingent transfer payments t_{ij} . Here t_{ij} is the net transfer from agent i to agent j , to be delivered after the endowment shocks are realized. For now, we do not model how the particular transfer arrangement is chosen. Next, nature moves, and each agent i receives an endowment realization e_i , where the vector of endowments (e_i) is drawn from a commonly known joint distribution. We assume that agents fully observe the endowments of others, effectively ruling out information-based reasons for limited risk-sharing.

After observing the endowments, agents can send transfers to each other. Let \tilde{t}_{ij} be the net transfer sent by agent i to agent j ; by definition, $\tilde{t}_{ij} = -\tilde{t}_{ji}$. Whenever an agent i chooses to transfer a different amount from what he had promised to pay, $\tilde{t}_{ij} \neq t_{ij}$, he loses his friendship link with j .³ Loss of a friendship is costly, because friends generate utility value: it is this social sanction that provides incentives to agents to keep their promises. Formally, at the end of the game, agents derive utility from two sources: goods consumption and friendship. Denoting $x_i = e_i - \sum_j \tilde{t}_{ij}$ the total goods consumption of i and $c_i = \sum_j \tilde{c}(i, j)$ his total remaining friendships, his realized utility is $U_i(x_i, c_i)$, where U_i is a smooth, increasing and concave utility function. The ex ante expected payoff of i is then $EU_i(x_i, c_i)$, where the expectation is taken over the realizations of endowment shocks.

1.2 Discussion of modelling assumptions

The two main ingredients in the model are the concept of transfer arrangements and the specification of social sanctions. Transfer arrangements can represent social customs and norms that have developed in a given community, as well as more explicit informal agreements. While the promised transfer payments t_{ij} are measured in dollar terms, they may also stand for transfers in kind, such

²We emphasize that this assumption is made for presentational purposes only. All our results extend to the case with asymmetric capacities.

³Such punishments may be less realistic when the value of the link to agent j is high, and the difference between the promised and actual transfer is small. However, the analysis in the paper is unaffected if we assume that making a lower than promised transfer results in a partial loss of friendship, high enough to make the deviation suboptimal. That is, the ‘‘punishment’’ can be in proportion with the ‘‘crime’’, without affecting our results.

as gifts of goods and services.

In this model, we formalize social sanctions by assuming that when an agent fails to make a promised transfer, the associated link is automatically lost. This loss of friendship captures the idea that friendly feelings may cease to exist if a promise is broken. It is possible to provide precise microfoundations for such behavior: Failure to make a transfer might signal that an agent no longer values a particular friendship, in which case the former friends might find it optimal not to interact with each other in the future. Mobius and Szeidl (2007) develop this idea formally. It is also possible that people break a link because of emotional or instinctive reasons in response to cheating: Fehr and Gächter (2000) provide evidence for such behavior.

One can imagine other social sanctions as well: for example, a deviating agent could be punished by all his friends, or by all agents in the community. Because these sanctions are stronger, we expect that they implement better risk-sharing than what can be obtained in our model. By modelling sanctions at the level of network connections, we essentially assume that in the event when a relationship goes bad, outsiders cannot assign the blame: they do not learn who broke a promise. This assumption is particularly realistic if relationships can go bad for reasons not connected to risk-sharing arrangements as well, such as personal conflicts.

Two important aspects of relational contracting in practice are repeated interactions and asymmetric information about endowments. While our model setup is a static one, we emphasize that it can also be interpreted as a “snapshot” of a more fully dynamic model, where the value of a network connection derives in part from the ability to conduct transactions through the link in the future. In any fixed period, conditional on the endogenous link values the dynamic model reduces to our current setup with quasi-linear utility: $U_i(x_i, c_i) = u_i(x_i) + c_i$, where c_i is the continuation value from future relationships.⁴ Since our static analysis applies for any set of capacities, it follows that our results about risk-sharing in networks extend to the dynamic model as well. We plan to analyze the dynamics of risk-sharing more explicitly in future work.

Agents in our model perfectly observe each other’s endowment shocks, and hence we have no information-based limits to insurance. Thus our results can be interpreted as a benchmark about the importance of limited commitment in explaining imperfect risk-sharing. Moreover, our full

⁴In dynamic settings it may be unrealistic to expect that $c_i = \sum_{j \neq i} c(i, j)$, i.e., the continuation value from keeping all links is equal to the sum of continuation values of individual links. However, our analysis remains valid as long as in the underlying dynamic model it is sufficient to consider deviations that involve withholding only one transfer. To see this, note that if c_i is the continuation value of i if all his links survive the current round and c'_i is his continuation value if he loses the link with j , then $c(i, j)$ can be defined to be $c_i - c'_i$. In this case, while $c_i = \sum_{j \neq i} c(i, j)$ is not necessarily true, we can still verify incentive compatibility based on link-specific capacity values $c(i, j)$. This one-link deviation property holds for all settings where the value of a link for an agent increases if he loses other links.

information assumption seems reasonable in the village environments that we are most interested in, where individuals can easily observe important economic attributes like the state of livestock or crops. This view is also supported by Udry (1994), who shows that asymmetric information between borrowers and lenders is relatively unimportant in villages in Northern Nigeria. This is not to say that asymmetric information is irrelevant for consumption insurance in general: the costs of observability rapidly increase with distance, and hence asymmetric information may be an important limitation to cross-village risk-sharing.

1.3 Incentive compatible risk-sharing arrangements

We focus on allocations that can be implemented using arrangements where agents always find it optimal to keep their promises ex post. This leads to the following definition:

Definition 1 *A risk-sharing arrangement t is incentive compatible (IC for short) if*

$$U_i(x_i, c_i) \geq U_i(x_i + t_{ij}, c_i - c(i, j)) \quad (1)$$

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

Intuitively, agent i must prefer to keep his friendship with j to defaulting on the promised transfer payment of t_{ij} . It is easy to see that if i does not find it optimal to default on a single transfer, he will also prefer not to default on a set of transfers: this follows from the quasi-concavity of the isoquants of U_i . Restricting our attention to IC arrangements does not reduce the set of feasible payoffs: For every non-IC arrangement t , we can construct an alternative IC arrangement t' by replacing t with a zero promised transfer in whenever it is optimal to default. This arrangement implements the same transfer payments as t , but no agent ever defaults, and hence utility is weakly improved.

Perfect substitutes. Our expected utility specification admits a useful special case when goods and friends are *perfect substitutes*, i.e., when $U_i(x_i, c_i) = u_i(x_i + c_i)$. Here a unit increase in the total capacity of friends is equivalent to a unit increase in goods consumption, and hence the value of friends is fixed in dollar terms. This case arises naturally when link values come from contemporaneous economic interactions that have an associated surplus measured in dollars. When goods and friendship are perfect substitutes, the incentive compatibility condition (1) simplifies

considerably: a transfer arrangement is IC if and only if

$$t_{ij} \leq c(i, j) \tag{2}$$

holds for all i and j . This condition means that the transfer amount can never exceed the capacity of a link: agent i cannot credibly commit to paying more to j than the value of their friendship. The simplicity and transparency of (2) makes this setting highly tractable. Because of this, we pay particular attention to the perfect substitutes case in the subsequent analysis, while highlighting how our results extend to general utility functions.

A key tool in extending our results to the imperfect substitutes case is a pair of necessary and sufficient conditions for incentive compatibility with general utility functions. To derive these conditions, 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 constants $k < K$ such that $k \leq MRS_i \leq K$ for all i , x_i and c_i . It is easy to see that when the MRS is uniformly bounded, then (i) any IC arrangement must satisfy $t_{ij} \leq K \cdot c(i, j)$, and (ii) any arrangement that satisfies $t_{ij} \leq k \cdot c(i, j)$ must be IC. The intuition can be seen by noting that the MRS measures the relative price of goods and friendship. If this relative price is always between k and K , then a transfer exceeding $Kc(i, j)$ is always worth more than the link and hence never IC, but a transfer below $kc(i, j)$ is always worth less than the link and hence is IC. In the perfect substitutes case, $MRS_i = 1$, so we can set $k = K = 1$, which yields (2).

2 The limits to risk-sharing

Our goal in this section is to characterize the degree of risk-sharing that obtains in a given social network. The central theme in the analysis is that good risk-sharing requires networks to have good expansion properties; that is, all groups of agents should have enough connections with the rest of the network, relative to group size. The intuition is simple: these connections allow every subset of agents to off-load their idiosyncratic shocks to the rest of the community. For most of the analysis, we assume that goods and friendship are perfect substitutes; we discuss how to extend the results to general utility functions at the end of the section.

2.1 An implementation result

We begin by looking at the problem of implementing a consumption profile in a fixed endowment realization. This will be helpful when we study the implementation problem under uncertainty later. To gain some intuition, consider the infinite line network depicted in Figure 2a, where all link capacities are equal to some fixed number c , and let F be a group of four consecutive agents. Suppose that the total endowment of these four agents is $e_F = \sum_{i \in F} e_i$, and that our goal is to implement a profile (x_i) where their total consumption is $x_F = \sum_{i \in F} x_i$. This implies that the group F as a whole must receive a transfer of $x_F - e_F$ from the rest of the network. This transfer can only flow through the two links at the endpoints of the interval F ; hence incentive compatibility requires that the capacity of these two links, $2c$, is greater than or equal to the total demand for resources, i.e., $2c \geq x_F - e_F$. To extend this logic for arbitrary sets of agents, we define the “perimeter” of a set of agents $F \subseteq W$ to be $c[F] = \sum_{i \in F, j \notin F} c(i, j)$.

Theorem 1 *Given endowment realization (e_i) , consumption profile (x_i) can be implemented with an IC arrangement if and only if (i) the resource constraint $x_W = e_W$ holds, and (ii) for all sets $|F| \leq N/2$,*

$$x_F - e_F \leq c[F]. \tag{3}$$

Since W is the set of all agents, the constraint $x_W = e_W$ simply means that the target allocation uses all available endowment. The key part of the Theorem is the set of bounds (3). We have already established that these bounds are necessary: the excess demand $x_F - e_F$ must flow through the perimeter $c[F]$. The surprising part of the theorem is that they are also sufficient. This result makes use of the mathematical theory of network flows, and in particular a corollary of the Ford and Fulkerson (1956) maximum flow - minimum cut theorem. To understand the basic idea, note that the maximum flow between vertices s and t in a graph is defined as the highest amount that can flow from s to t along the edges respecting the capacity constraints given by link values. Ford and Fulkerson show that the value of the maximum flow equals the value of the minimum cut, i.e., the smallest capacity that has to be deleted such that s and t end up in different components. To apply this result here, we add two hypothetical agents s and t to the network G and transform the implementation problem into a flow problem such that implementing profile (x) is equivalent to finding a large enough flow between s and t . Every cut in this transformed problem corresponds to the perimeter of some set F in the original network, and hence the desired flow exists if all cuts are large enough, which is exactly condition (3).

2.2 The limits to full risk-sharing

Theorem 1 can be used to characterize the networks that allow Pareto-optimal full risk-sharing. To understand the logic, consider the three infinite networks depicted in Figure 2, where all link capacities are equal to some number c . Let the endowment shocks be independent binary random variables, so that for each agent i , e_i is either σ or $-\sigma$ with equal probabilities.

Can equal sharing be implemented in these examples? The law of large numbers implies that the average endowment is almost surely zero, hence equal sharing implies all agents consuming zero almost surely. Consider first the “line” network in Figure 2a. Select an “interval set” F : since endowment shocks are i.i.d., with positive probability all agents in F receive a negative income shock of $-\sigma$. Because all these agents must consume zero, Theorem 1 implies that $2c \geq |F| \cdot \sigma$ has to hold for every F . But for any fixed value of c , we can find a long enough interval that violates this inequality; as a result, full risk-sharing cannot be implemented on the line. A similar negative result holds for the 2-dimensional “plane” network in Figure 2b. The perimeter of a square-shaped set F in this network is $4c \cdot |F|^{1/2}$, which is smaller than $|F|$ for a large square; hence in the event where all agents in F get a negative shock, equal sharing must fail. However, this argument does not rule out equal sharing for the infinite “binary tree” network in Figure 2c. Here, the perimeter of any set F is at least $\sigma \cdot |F|$, and so for $c > \sigma$, the transfers required for equal sharing can flow into any set F in any endowment realization.

The above examples suggest that the perimeter relative to the size of certain sets F governs whether full risk-sharing can be implemented. To formalize this intuition, we first introduce some notation. Let $a[F] = c[F] / |F|$ be the “perimeter-to-area ratio” of F , where the “area” is simply the number of agents in F , and let $\underline{\sigma} = \min_i \sigma_i$ denote the minimum standard deviation of endowment shocks in the network. We say that endowments have a product support if for all i , the support of e_i given any realization of (e_{-i}) is the same as its unconditional support. This is a weak assumption, ensuring that there is some idiosyncratic component in each agent’s endowment shock.⁵

Proposition 1 *[Limits to full risk-sharing]*

(i) *Suppose endowments have a product support. If $a[F] < \underline{\sigma}$ for some F with $|F| \leq N/2$, then no IC allocation implements equal sharing.*

(ii) *Suppose shocks are symmetric binary, with $e_W = 0$. If $a[F] \geq \sigma$ for all F with $|F| \leq N/2$, then there exists an IC allocation implementing equal sharing.*

⁵Bloch, Genicot and Ray (2006) impose a similar condition on endowment shocks in their Assumption 1.

Part (i) states that when the perimeter/area ratio of at least one set is smaller than the measure of endowment risk $\underline{\sigma}$, then full risk-sharing cannot be implemented. The intuition is simple: by the product support assumption, there are realizations where the cumulative idiosyncratic shock inside F is larger than the perimeter. This shock cannot be completely transferred away, and hence equal sharing must fail. Part (ii) is a partial converse for symmetric binary shocks and no aggregate uncertainty. This result follows directly from Theorem 1. When $e_W = 0$, equal sharing means that all agents must consume zero. Given that shocks are binary, any set F has an excess demand of at most $|F| \cdot \sigma$ relative to the target of zero consumption. This demand is less than or equal to the perimeter $c[F]$ because $a[F] \geq \sigma$, and hence can be satisfied for all sets F .

The proposition shows that full risk-sharing imposes very strong expansion conditions on the network structure, which are unlikely to be satisfied in real-world networks. In fact, since social networks in practice are often organized on the basis of geographic location, we expect that their structure more closely resembles the 2-dimensional grid on the plane. For these networks, full risk-sharing fails by part (i) of the Proposition, and hence we need to explore the degree of partial risk-sharing that might obtain.

2.3 Partial risk-sharing

We begin our analysis of partial risk-sharing with an intuitive argument. Assume that our goal is to implement a profile where all agents consume zero. We know from Theorem 1 that the cumulative shock in a set F can leave the set in a realization if $e_F \leq c[F]$. This result suggests that risk-sharing should be reasonably good when $e_F \leq c[F]$ holds “most of the time.” This will be the case for example when $\sigma_F = \text{stdev}[e_F]$ is small relative to $c[F]$, because the standard deviation is a measure of the “typical magnitude” of e_F . This logic suggests that we might expect good but imperfect risk-sharing if $c[F]$ is sufficiently larger than σ_F for most sets F .

Note, there is an important difference between this argument and the results in Proposition 1. In the Proposition, the perimeter/area ratio is compared to the standard deviation of *individual* endowment shocks, and hence the correlation structure across agents is not exploited. Here, we make use of the correlation structure by computing the standard deviation over *sets* of agents. To see why this matters, note that full risk-sharing as characterized by Proposition 1 requires $c[F]$ to be proportional to $|F|$. But if endowment shocks are i.i.d., then the standard deviation of e_F is only of order $|F|^{1/2}$, and hence our argument suggests that good risk-sharing can obtain even if the perimeter is of order $|F|^{1/2}$, which can be much smaller than $|F|$. In particular, for the plane

network, where the perimeter of square shape sets is $c[F] \sim |F|^{1/2}$, this logic suggests reasonably good risk-sharing. In contrast, for the line we still expect risk-sharing to be poor, because the perimeter of long interval sets will be smaller than $|F|^{1/2}$.

To formalize this intuition, we need to develop a measure of partial risk-sharing. Using the equal sharing profile where all agents consume $\bar{e} = (1/N) \sum_i e_i$ as the full risk-sharing benchmark, assuming that agents have identical preferences over goods consumption, a natural measure of risk-sharing is

$$UDISP(x) = \mathbb{E} \frac{1}{N} \sum_i \{U(\bar{e}) - U(x_i)\}$$

where we ignored the dependence of utility on links to simplify notation. *UDISP*, or “utility-based dispersion,” is simply the difference between average expected utility in the allocation (x) relative to the first best of equal sharing, and hence lower values correspond to better risk-sharing. In particular, $UDISP(x) \geq 0$, with equality if and only if x implements equal sharing. If all agents have the same quadratic utility function over x , then we can express *UDISP* as an increasing function of

$$SDISP(x) = \left[\mathbb{E} \frac{1}{N} \sum x_u^2 \right]^{1/2}, \quad (4)$$

which is the square-root of the expected cross-sectional variance of x . For non-quadratic utilities, *SDISP*(x) can be interpreted as a second order approximation of the utility based measure. *SDISP* is highly tractable, and inherits the intuitive properties of *UDISP*: it is zero only in equal sharing and positive otherwise, and its magnitude measures the departure from equal sharing: e.g., if e_u are symmetric binary, then in autarky $SDISP(e) = \sigma$. For these reasons, we concentrate on *SDISP* in the analysis below.

Before stating the formal results, we make some assumptions about the distribution of shocks. We assume that the source of uncertainty is a collection of independent random variables y_j , $j = 1, \dots, \infty$, which can represent idiosyncratic shocks like illness as well as aggregate shocks like weather. Different agents may have different exposure to these shocks, so that $e_i = \sum_j \alpha_{ij} y_j$, where α_{ij} measures the extent to which agent i is exposed to shock j . We assume that $\sum_j \alpha_{ij}^2$ is uniformly bounded for all agents, and that y_j have uniformly bounded support.⁶ One natural special case is when $e_i = y_i$ are i.i.d. We also require that shocks are not too correlated, so that aggregate uncertainty disappears at a rate $\sigma_F \leq K_1 \cdot |F|^{1/2}$ with some $K_1 > 0$. On the line or the plane, this

⁶This assumption can be relaxed: we only require bounds on the moment-generating function of y_j . Normally distributed random variables also satisfy these bounds, and hence all our results extend to jointly normal shocks.

property holds for example when the correlation between endowment shocks decays geometrically with network distance. Finally, we assume that a larger group of people face more risk, so that for all $G \subseteq F$, we have $\sigma_G \leq \sigma_F$; and that sharing risk with more people is always good, i.e., that for all $G \subseteq F$, we have $\sigma_F/|F| \leq \sigma_G/|G|$.

Proposition 2 *Under the above conditions, there exist K and K' constants such that*

(i) *On the line with capacities c and i.i.d. shocks, we have $SDISP(x) \geq K/c$ for all IC arrangements (x).*

(ii) *On the plane with capacities c , we have $SDISP(x) \leq \exp[-K'c^{2/3}]$ for some IC arrangement (x).*

This proposition formalizes our earlier intuition by comparing the rate of convergence to full risk-sharing as we increase capacities, between the line and plane networks.⁷ Intuitively, this criterion examines the effectiveness of strengthening the existing links in the network. For the line we expect poor risk-sharing, because the size of shocks grows faster than the perimeter of sets. Formally, this means that $SDISP$ goes to zero at a slow rate of $1/c$ as c goes to infinity. But for the plane, we expect very good risk-sharing because the perimeter has the same order of magnitude as the standard deviation of the shock; as a result, $SDISP$ should go to zero at a fast, exponential rate. Risk-sharing on the plane is thus qualitatively different from risk-sharing on the line: convergence to equal sharing is exponential as opposed to polynomial.

Numerical simulations suggest that the asymptotic results of the Proposition are good descriptions of behavior for finite c as well. To take one example, consider Figure 3, which shows constrained optimal allocations for finite line and plane networks with unit capacities, for a given realization of binary shocks with $\sigma = 1$.⁸ For both the line and the plane, the black-and-white panel shows the endowment realization, while the grey panel shows the optimal allocation that can be implemented with IC transfers. The figures represent typical endowment realizations: they were randomly drawn, and we have played around with many realizations. The contrast between the line and the plane is remarkable: for the line, we see substantial color variation in the grey panel, reflecting imperfect risk-sharing ($SDISP = 31\%$); but for the plane, equal sharing can be implemented in this particular realization ($SDISP = 0$).

The proof of the proposition works the following way. For part (i), we split the line into equal

⁷It can be shown that under mild conditions on the distribution of endowments, for any connected network $SDISP$ converges to zero as the capacities of links go to infinity.

⁸In these simulations, the line network is a segment, and the plane network is a square.

interval sets, and show that for long intervals, much of the shock over each interval must remain inside the set. We choose intervals of approximate length $16c^2$, which implies that σ_F for each interval is on the order of $(16c^2)^{1/2} = 4c$. Since the perimeter of each interval is $2c$, a standard deviation of $4c - 2c = 2c$ must remain in each of these sets. If agents in a set manage to smooth this residual shock perfectly, then the per capita residual standard deviation will be $2c/|F| \sim 1/c$, establishing the desired lower bound.

The result for the plane is more difficult. Here, we have to construct an IC arrangement that implements very good risk-sharing. We do this in two steps: first we construct an “unconstrained” arrangement that implements equal sharing on a 2^m by 2^m sized square for some m chosen as a function of c ; second, we show that this unconstrained arrangement violates capacity constraints infrequently, and hence there exists an IC arrangement that implements almost as good risk-sharing as the unconstrained one. For the unconstrained arrangement, we make use of a partition of the network where we split the square of size 2^m by 2^m into four equal squares, split each of these into four smaller squares, and repeat this procedure m times. Then we build the unconstrained flow from the bottom up: first we smooth consumption in the smallest squares, then we smooth consumption in the squares at the next level, and so on. There are m steps in this procedure, and hence each link is used only m times. Assuming that $\sigma = 1$, every time a link is used, the required capacity is of order 1, because the standard deviation in a square F is $\sigma_F = |F|^{1/2}$, and this must be distributed over the $4 \cdot |F|^{1/2}$ links on the perimeter. Thus with m iterations, we implement equal sharing in the big square while only using a link capacity of order m on average. Equal sharing in the big square corresponds to an *SDISP* of order e^{-m} , thus a choice of $m = c$ would implement exponentially good risk-sharing. However, we have to worry about the exceptional events when the capacity constraints are violated. Using the theory of large deviations we prove that these exceptional events can be taken care of with a choice of $m = c^{2/3}$, resulting in the bound of the proposition.

2.4 Geographic networks

The results for the plane are interesting because real-world networks are likely to be similarly two-dimensional. To formalize this idea, suppose that the social network can be represented by a map on the plane. If the correlation between agents’ endowment shocks falls fast enough with distance on the map, then we expect that σ_F grows at the rate of $|F|^{1/2}$, just like in the plane network. Moreover, if agents tend to have friends at close geographic distance in multiple directions, then

it is plausible that the perimeter of sets F also grows at a rate of $|F|^{1/2}$. These observations suggest that the key properties of the plane are preserved for a wider class of *geographic networks*, and hence we expect good risk-sharing for them.

To make this argument precise, consider an infinite network, and let $\pi : W \rightarrow \mathbb{R}^2$ map agents in this network to locations in the plane such that different agents are assigned different locations. To ensure that agents have friends in multiple directions, consider two a by a squares on the plane sharing a common side, with sides parallel to the axes, and define the crossing density as the total capacity of all links connecting agents in one square with agents in the other, normalized by a . We say that the embedding exhibits *no separating avenues* if the crossing density of all large enough squares is bounded away from zero. If this holds, then agents in a large square always have friends in neighboring squares. We also assume that for large squares, the part of the network that falls into a square is connected; and that the population density in all large enough squares is bounded from above and below. Regarding the endowment shocks, we require that the correlation between individual endowments falls geometrically with distance, i.e., that $\text{corr}[e_i, e_j] \leq K_2 \exp[-d(i, j)/K_2]$ for some $K_2 > 0$, where $d(u, v)$ is the Euclidean distance between $\pi(u)$ and $\pi(v)$. A network is called a geographic network if these conditions are satisfied.

Corollary 1 *In geographic networks, we have $SDISP(x) \leq \exp[-K'c^{2/3}]$ for some IC arrangement (x) .*

The proof combines Proposition 2 with a renormalization argument. We take a geographic network, and superimpose on its planar representation a grid with large squares. Then we merge all people within each square to create a new network. Because of the no separating avenues condition, this new network is essentially a plane, and hence Proposition 2 (ii) can be applied to yield a bound for $SDISP$ of the new network. We then lift this bound back to the old network using the assumptions of bounded population density and connectedness inside the squares.

The Corollary is useful because it can help explain stylized facts in development economics. Real-world village networks are likely to be organized partly on a geographic basis, and hence are likely to satisfy the properties of a geographic network. As a result, our model predicts very good informal risk-sharing in these villages, which is 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 test full risk-sharing in the data is to regress the consumption of an individual or a group on their own endowment shock as well as 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 specific to F . It is easy to see that with full risk-sharing, we get $\beta_F = 0$; this corresponds to the test of full risk-sharing used in Cochrane (1991), Mace (1991), Townsend (1994) and others. When $\beta \neq 0$, full risk-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 have*

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

The regression coefficient β has a lower bound which 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, there is insufficient capacity for the shock to exit, which yields high correlation between consumption and shocks. The proposition is related to the empirical findings in Townsend (1994), who shows that there is considerable risk-sharing among households within an Indian village, but only limited sharing of village-specific shocks across villages. The proposition is consistent with these findings if cross-village network ties are weak relative to the size of the villages.

2.6 The limits to risk-sharing with imperfect substitutes

We now discuss briefly how the results in this section extend to the imperfect substitutes case. We find that all results extend, but the upper and lower bounds on risk-sharing are weakened by constant factors that depend on the degree of substitutability between goods and friendship. Since the results about partial risk-sharing characterize limit behavior, they remain unaffected by these constant factors. To obtain our extensions, we assume that the marginal rate of substitution

(MRS) is uniformly bounded: $k < MRS_i < K$ for all agents i in the relevant range of endowment realizations.

We begin with extending Theorem 1. When the MRS is bounded, the necessary and sufficient condition in the Theorem is replaced by the following two conditions, one being necessary, the other sufficient for IC implementation. (i) Any IC profile must satisfy $x_F - e_F \leq K \cdot c[F]$ for all sets F with $|F| \leq N/2$. (ii) A profile that satisfies $x_F - e_F \leq k \cdot c[F]$ for all F with $|F| \leq N/2$ is IC. The extension follows directly from the logic of Theorem 1, noting that bounded MRS implies that the relative price of friendship and goods is always between k and K . This extension is particularly informative for environments where k and K are close to each other, e.g., when endowment shocks have a small support: then, by continuity, the MRS does not vary much in the relevant range of realizations.

The uniform bound on the MRS can also be used to extend the characterization of environments where full risk-sharing can be implemented. We continue to find that the first-best can only be achieved in expander graphs where the perimeter/area ratio is bounded from below: we require $a[F] \geq \underline{\sigma}/K$. We also find that in the binary shock case, full risk-sharing can be implemented when $a[F] \geq \sigma/k$. These results imply that full risk-sharing fails for most plausible networks even with imperfect substitutes. Our findings about partial risk-sharing characterize convergence rates, and hence they extend without modification to the imperfect substitutes case. In particular, *SDISP* continues to converge exponentially for geographic networks, and therefore our argument about good risk-sharing in real-world networks is unaffected.

The setup where goods and friendship are imperfect substitutes yields some additional implications as well. If the marginal rate of substitution between goods and friendship is increasing in consumption, then agents with low consumption value their friends less, reducing the maximum amount they are willing to transfer to them. As a result, if in a society that experiences a negative aggregate shock, the scope for insuring idiosyncratic risk is reduced because of the drop in the dollar value of links. We formalize this intuition in the appendix by showing that when the *MRS* is increasing, reducing the endowments of all agents results in a smaller set of IC transfer arrangements. In particular, *SDISP* is larger after a negative aggregate shock, because agents are more constrained in insuring idiosyncratic risk. These results are consistent with the findings of Kazianga and Udry (2006), who document that during the severe draught of 1981-85 in rural Burkina Faso, risk-sharing between households was quite limited.

3 Constrained efficient risk-sharing

In this section, we study allocations that are optimal subject to the enforcement constraints imposed by the network. A risk-sharing arrangement is *constrained efficient* or second-best if it is Pareto-optimal among the set of IC arrangements. Constrained efficient arrangements provide a natural benchmark, because they achieve the highest possible level of risk-sharing in a given social network. In addition, we show below that constrained efficiency can arise both when agents follow simple rules of thumb for helping each other, and also as a result of dynamic coalitional bargaining. Once reached, constrained efficient arrangements are likely to remain stable, because they are robust to both individual and coalitional deviations.

As in the previous section, we start out by assuming that goods and friendship are perfect substitutes, and extend the results to general preferences later.

3.1 Characterizing constrained efficiency

In this subsection we maintain the assumption that goods and friendship are perfect substitutes. The study of second-best arrangements is facilitated by the fact that they can be characterized using a planner's problem. Formally, let $(\lambda_i)_{i \in W}$ be a set of positive weights, and define the *planner's problem* as the constrained optimization problem

$$\max \sum_{i \in W} \lambda_i \cdot EU_i(x_i, c_i) \tag{5}$$

subject to the IC-constraint (1). We then have the following result.

Proposition 4 *Every constrained efficient 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 idea that the set of possible payoff vectors is convex. 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 select the efficient allocation by design. Adapting this argument to our model requires that the set of IC payoff vectors be convex. In the perfect substitutes case, this follows easily: when two transfers satisfy a capacity constraint, their convex combination will also satisfy it. As we detail in the Appendix, the result

can also be extended to the imperfect substitutes case under an additional condition about the curvature of the marginal rate of substitution.

Proposition 4 implies that maximizing the planner's expected utility $E\sum \lambda_i U_i$ is equivalent to maximizing realized utility $\sum \lambda_i U_i$ independently for each state, because conditional on the planner weights, the states are not connected in the maximization problem. This separation of states simplifies the characterization of second-best arrangements, and makes it easier to solve for them. In particular, we can derive a set of first-order conditions for the planner's problem separately for each state, which allows for a simple characterization of second-best arrangements. We say that a link from i to j is blocked in a given realization, if $t_{ij} = c(i, j)$, that is, if i is sending the maximum IC amount.

Proposition 5 *A transfer arrangement t is constrained efficient iff there exist positive welfare weights $(\lambda_i)_{i \in W}$ such that for every $i, j \in N$ one of the following conditions holds:*

- 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 states the set of first order necessary and sufficient conditions for the planner's problem. To understand the intuition, recall that as Wilson (1968) has shown, unconstrained Pareto optimal risk-sharing implies that $\lambda_i U'_i(x_i) = \lambda_j U'_j(x_j)$ for all i and j . If this condition is violated, e.g., $\lambda_i U'_i(x_i) < \lambda_j U'_j(x_j)$, then the planner's objective can be improved by transferring a small amount from i to j . In a second best arrangement, this transfer must violate the incentive compatibility constraint; as a result, the maximum possible amount must already be transferred from i to j . This logic establishes the necessity of the above first order conditions; sufficiency follows because the planner's objective function is concave and the domain of IC consumption profiles is convex. These conditions also guarantee uniqueness.

The Proposition also implies that for any pair of agents i and j , if $\lambda_i U'_i < \lambda_j U'_j$, then all paths connecting i and j have to be blocked in the sense that at least one link along each path is used at maximal capacity. This observation uncovers an important feature of constrained efficient agreements, namely that 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 for $i, j \in W_k$ we have $\lambda_i U'_i = \lambda_j U'_j$ and for $i \in W_k$, and $j \notin W_k$ either $t_{ij} = c(i, j)$ or $t_{ji} = c(j, i)$.

The sharing island W_k of i is the maximal connected set containing i with the property that $\lambda_i U'_i = \lambda_j U'_j$ for all $j \in W_k$. For each realization, these sharing islands provide a partition of the network, and have the property that shocks are fully shared within an island but there is imperfect insurance across islands. This island structure is illustrated in the line network in Figure 3, where the grey panel depicts the constrained efficient allocation corresponding to equal planner weights in one endowment realization. The dashed lines in the figure indicate the boundaries of the islands; marginal utility and hence consumption is equalized within an island, but differs across islands.

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 endowment realization. This result differentiates our model from group-based theories of risk-sharing, where insurance groups are determined exogenously and do not vary with the realization of uncertainty.

3.2 Spillover effects and local sharing

The characterization of constrained efficient allocations in terms of risk-sharing islands can be used to explore the degree of partial insurance as a function of network distance. This analysis sheds light on the spillover effects of shocks in networks, and yields new testable implications about local risk-sharing.

We begin by introducing 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, i.e., the maximal connected set with the property that $\lambda_i U'_i = \lambda_j U'_j$ for all $j \in W(i)$. We now define $\widehat{W}(i)$ to be the maximal connected set of agents j such that there exists a path between i and j along which no links are blocked in either direction. With this definition, $\widehat{W}(i) \subset W(i)$, because the first order condition of Proposition 5 implies $\lambda_i U'_i = \lambda_j U'_j$ for all $j \in \widehat{W}(i)$. Moreover, except for knife-edge cases when the transfer amount just reaches the capacity over a link but does not “bind” yet, these two island definitions are equivalent, and $\widehat{W}(i) = W(i)$. It can be shown that these knife-edge cases happen with zero probability when the distribution of endowment shocks is absolutely continuous; as a result, the two definitions can be treated as equivalent for practical purposes.

To understand the connection between partial risk-sharing and network distance, we 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 endowment realizations $e = (e_i)$ and $e' = (e'_i)$, where $e'_j = e_j$ for all $j \neq i$, and $e'_i < e_i$. These two realizations can be viewed either as agent i experiencing an idiosyncratic negative shock in e' relative to e , or as agent i experiencing an idiosyncratic positive shock (aid) in e relative to e' . For ease of exposition, in the discussions below we focus on the first interpretation: that agent i receives a negative shock in e' . We can measure the impact of this shock on agent j by computing the ratio of marginal utilities

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

Here MUC_j 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.⁹

Corollary 2 [*Spillovers and local sharing*]

- (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 $j \in W \setminus \widehat{W}(i)$.
- (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, $MUC_l \geq MUC_j$.

Part (i) shows that efficient arrangements 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 who are in the same risk-sharing island, who all reduce their consumption by a positive amount. As a result, unless i is in a singleton island, he has access to at least some insurance. The intuition follows from the definition of $\widehat{W}(i)$: links within the risk-sharing island of i are not blocked, and hence a small transfer can flow through them to help out i . As part (ii) shows, for small shocks, the set of agents who insure i is exactly his sharing island $\widehat{W}(i)$. All these agents share an equal burden of the shock, and hence experience the same marginal utility cost. Agents outside of $W(i)$ do not reduce their consumption at all; and in the knife edge case where $\widehat{W}(i) \neq W(i)$, agents in $W(i) \setminus \widehat{W}(i)$ may or may not share. Finally, part (iii) shows how the utility cost of agents in response to an idiosyncratic shock to i varies by social

⁹Analysing the impact of a positive iniosyncratic shock to i yields symmetric results.

distance. The result states that indirect friends provide less insurance to i than direct friends: for any agent $j \neq i$, there exists some direct friend of i , denoted l , who shares at least as much of the burden of the shock as j does.

The results of the Corollary are also illustrated in Figure 4, which shows the marginal utility cost of direct and indirect friends in response to an endowment shock to i . The horizontal axis is the marginal utility cost of agent i himself, and the vertical axis measures the marginal utility cost of some direct and indirect friends. For small shocks, both direct and indirect friends who are in the same risk-sharing island as i help out. As the size of the shock grows, some indirect friends hit their capacity constraints and stop reducing consumption, but some direct friends continue to help. After a point, all direct friends hit their capacity constraints; as a result, 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 partial risk-sharing as a function of network distance.

3.3 Foundations for constrained efficiency

One reason for analyzing constrained efficient allocations is that they naturally emerge from intuitive dynamic mechanisms among agents in the network. Here we briefly discuss two such mechanisms. First, constrained efficient allocations can be obtained in a decentralized procedure where agents use a simple rule of thumb in helping those who are in need. In every round of this dynamic procedure, agents attempt to equate weighted marginal utilities between neighbors subject to the capacity constraints: intuitively, people help out those friends and relatives who are in need. The appendix shows that this procedure converges to the constrained efficient allocation corresponding to the welfare weights used. As a result, constrained efficient allocations can emerge even if agents only use local information: in every round of the procedure, agents engage in binary transactions that only knowledge about the current resources of the two parties involved.¹⁰

A second mechanism that leads to constrained efficient arrangements is collective dynamic bargaining with renegotiation. Gomes (2005) shows that when agents can propose renegotiable arrangements to subgroups and make side-payments in a dynamic bargaining procedure, as the community become infinitely patient, a Pareto-efficient arrangement will be selected.¹¹ This result can be incorporated in our model by assuming that there is a negotiations phase prior to the

¹⁰Bramouille and Kranton (2007) use a similar procedure with equal welfare weights and no capacity constraints.

¹¹Aghion, Antras, and Helpman (2007) establish a similar result in a model involving renegotiating free-trade agreements.

endowment realization, and would imply that agents select a constrained efficient risk-sharing arrangement.

Finally, constrained efficient arrangements are quite stable in our model, because they are robust to all possible coalitional deviations. The appendix shows that after any endowment realization, no group of agents can have a credible and profitable deviation that involves IC transfers among members of the group, and possibly renegeing on some of the transfers to agents outside the group.

3.4 General preferences

We now turn to discuss how our results about constrained efficiency extend to the imperfect substitutes case. We find that all our conceptual results generalize. The formal results are provided in the Appendix; here we present an intuitive summary.

The key novelty with imperfect substitutes is that changing the goods consumption of an agent affects the agent's link values, and hence his incentive compatibility constraints over transfers. To characterize constrained efficiency in this environment, we assume that the marginal rate of substitution MRS_i , i.e., the relative price of friendship in terms of goods consumption, is concave in x_i . Intuitively, this means that increases in goods consumption have a diminishing effect on the value of friendship. When this condition 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 his marginal utility of i 's consumption, because increased consumption also softens i 's IC constraints over transfers to neighboring agents. As a result, it may be optimal from the planner's perspective to transfer some of the original dollar to such neighboring agents with whom the IC constraint of i was previously binding. Due to this difference between private and social marginal utility, we can have constrained efficient arrangements where i is transferring a positive amount to i' even though i' has lower weighted marginal utility, because this transfer, by keeping the consumption of i' high, softens the IC constraint of i' on transfers to another agent i'' , who has high marginal utility. To get around this issue, in the Appendix we define the marginal social gain of an additional unit of transfer to i , denoted Δ_i , for each agent i using an iterative procedure, which takes into account the indirect effect of softening IC constraints and transferring further some of the additional resources.

Using Δ_i instead of the private marginal utilities allows us to extend all the results in this

section. We obtain first order conditions that are analogous to Proposition 5: in a constrained efficient arrangement, either $\Delta_i = \Delta_j$ of the link between i and j has to be blocked. Using this result, we can also partition the network into endogenous risk-sharing islands, such that Δ_i is equalized within islands while links are blocked across islands. The results of Corollary 2 on monotonicity, local sharing, and more sharing with close friends also have analogous extensions to this environment, which are formally presented in the Appendix.

Finally, for an agent i who is not on the boundary of his risk-sharing island and hence has no binding IC constraints, the marginal social gain does equal λ_i times his marginal utility of consumption; hence, for such agents, the results presented in this section hold without modification. For example, weighted marginal utilities are equalized for two agents inside the same risk-sharing island and away from the island boundary. This argument establishes that if risk-sharing islands are “large”, then the results from the perfect substitute case hold without modification for most agents.

4 Indirect effects of an aid program

We plan to simulate our model using network data from Peru, to evaluate the indirect effects of a hypothetical development aid program. Because of network spillovers, in our model aid will also benefit the non-treated, as shown in Corollary 1 in the previous section. This is consistent with the empirical findings of Angelucci and De Giorgi (2007). We plan to identify individuals who should be targeted to maximize both the direct and indirect effect of aid.

5 Conclusion

This paper has developed a theory of informal risk-sharing in social networks. We have shown that expansive networks facilitate informal insurance, and argued that many real-life social networks are likely to be sufficiently expansive to allow for good risk-sharing. We also characterized second-best arrangements and found that they exhibit local risk-sharing. In current work, we are exploring the implications of our model for the indirect effect of development aid. In future work, we would like to develop other empirical applications.

References

- AGHION, P., P. ANTRAS, AND E. HELPMAN (2007): “Negotiating free trade,” *Journal of International Economics*, forthcoming.
- ANGELUCCI, M., AND G. DE GIORGI (2007): “Indirect Effects of an Aid Program: How do Cash Injections Affect Non-Eligibles’ Consumption?,” Working paper.
- BLOCH, F., G. GENICOT, AND D. RAY (2005): “Informal Insurance in Social Networks,” Working paper, New York University.
- BRAMOULLE, Y., AND R. E. KRANTON (2006): “Risk-Sharing Networks,” *Journal of Economic Behavior and Organization*, forthcoming.
- (2007): “Risk Sharing Across Communities,” *American Economic Review, Papers and Proceedings*.
- COATE, S., AND M. RAVAILLON (1993): “Reciprocity Without Commitment: Characterization and Performance of Informal Insurance Arrangements,” *Journal of Development Economics*, 40, 1–24.
- COCHRANE, J. H. (1991): “A Simple Test of Consumption Insurance,” *Journal of Political Economy*, 99(5), 957–976.
- DERCON, S., AND J. D. WEERDT (2006): “Risk-sharing networks and insurance against illness,” *Journal of Development Economics*, 81, 337–356.
- DIXIT, A. (2003): “Trade Expansion and Contract Enforcement,” *Journal of Political Economy*, 111, 1293–1317.
- ELLISON, G. (1994): “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching,” *Review of Economic Studies*, 61, 567–588.
- FAFCHAMPS, M., AND F. GUBERT (2007): “The formation of risk-sharing networks,” *Journal of Development Economics*, forthcoming.
- FAFCHAMPS, M., AND S. LUND (2003): “Risk sharing networks in rural Philippines,” *Journal of Development Economics*, 71, 261–287.

- FEHR, E., AND S. GACHTER (2000): “Cooperation and punishment in public goods experiments,” *American Economic Review*, 90, 980–994.
- FORD, L. R. J., AND D. R. FULKERSON (1956): “Maximal Flow Through a Network,” *Canadian Journal of Mathematics*, pp. 399–404.
- GOMES, A. (2005): “Multilateral contracting with externalities,” *Econometrica*, 73, 1329–1350.
- GREIF, A. (1993): “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition,” *American Economic Review*, 83, 525–548.
- KANDORI, M. (1992): “Social Norms and Community Enforcement,” *Review of Economic Studies*, 59, 63–80.
- KAZIANGA, H., AND C. UDRY (2006): “Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso,” *Journal of Development Economics*, 79, 413–446.
- KOCHERLAKOTA, N. (1996): “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 63(4), 595–609.
- KRANTON, R. E. (1996): “Reciprocal Exchange: A Self-Sustaining System,” *American Economic Review*, 86, 830–851.
- LIGON, E. (1998): “Risk-sharing and information in village economies,” *Review of Economic Studies*, 65, 847–864.
- LIGON, E., J. P. THOMAS, AND T. WORRALL (2002): “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies,” *Review of Economic Studies*, 69(1), 209–244.
- MACE, B. (1991): “Consumption volatility: borrowing constraints or full insurance,” *Journal of Political Economy*, 99, 928–956.
- MAZZOCCO, M. (2007): “Household Intertemporal Behavior: a Collective Characterization and a Test of Commitment,” *Review of Economic Studies*, 74(3), 857–895.
- MOBIUS, M., AND A. SZEIDL (2007): “Trust and Social Collateral,” Working paper no. 13126, NBER.

- OGAKI, M., AND Q. ZHANG (2001): “Decreasing Relative Risk Aversion and Tests of Risk Sharing,” *Econometrica*, 69(2), 515–526.
- TOWNSEND, R. (1994): “Risk and Insurance in Village India,” *Econometrica*, 62, 539–591.
- UDRY, C. (1994): “Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria,” *Review of Economic Studies*, 61(3), 495–526.
- WILSON, R. (1968): “The theory of syndicates,” *Econometrica*, 36, 119–132.

Appendix: Proofs

Proof of Theorem 1

We prove the more general version of the theorem allowing for directed links, so that $c(i, j)$ and $c(j, i)$ may differ. Necessity is immediate. To prove sufficiency, 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 S be the set of agents for whom $g_i \geq 0$ and let $T = W \setminus S$. Define the auxiliary graph G' which has two additional vertices, s and t , and additional edges connecting s with all agents in S , and additional edges connecting t with all agents in T . For any $i \in S$, define the capacity $c(s, i) = g_i$ and $c(i, s) = 0$. Similarly, for any $j \in T$, let $c(j, t) = -g_j$ and $c(t, j) = 0$.

The auxiliary graph is useful, because implementing the desired consumption allocation is equivalent to finding an $s \rightarrow t$ flow in G' that has value $g_S = \sum_{e_i \geq 0} g_i$. To see why, note that in the desired allocation, exactly g_i must leave each agent $i \in S$. The capacities on the new links ensure that in any $s \rightarrow t$ flow, at most g_i can leave agent i . Similarly, to implement the target, exactly $-g_j$ must flow to each agent $j \in T$, 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 $\sum_{e_i \geq 0} g_i$ 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 $s \rightarrow t$ flow, we instead compute the value of the minimum cut. Fix a minimum cut. In this cut, some links of s and t are cut. Let $S_1 \subseteq S$ denote those agents whose link with s is cut in the minimum cut, and let $T_1 \subseteq T$ denote those agents in T whose link with t is cut. Clearly, the total value of the links cut that connected S_1 with s and T_1 with t is $g_{S_1} - g_{T_1}$. Let $S_2 = S \setminus S_1$ and $T_2 = T \setminus T_1$. We claim that if we consider the restriction of the cut to the original graph G , then there will be no $S_2 \rightarrow T_2$ paths that survive. Suppose not; then there is some $s_2 \rightarrow t_2$ path in G after the cut. But this is also an $s_2 \rightarrow t_2$ path in the auxiliary graph G' , and since s_2 is connected to s and t_2 to t after the cut, it generates an $s \rightarrow t$ path after the cut, which is a contradiction.

Let H be the set of agents h who can be accessed by $s \rightarrow h$ paths in G after the cut. By the above argument, $S_2 \subseteq H$ and $T_2 \subseteq T \setminus H$. By construction of H , the value of the cut in G must be $c^{\text{out}}[H] = c^{\text{in}}[W \setminus H]$, and therefore the value of the cut in G' is $g_{S_1} - g_{T_1} + c^{\text{out}}[H]$. Suppose that $|H| \leq N/2$. Then we know from (3) that $c^{\text{out}}[H] \geq g_H$. Thus the value of the cut in G' can be bounded from below as

$$g_{S_1} - g_{T_1} + g_H = g_{S_1} - g_{T_1} + (g_{S_2} + g_{H \cap T_1} + g_{H \cap S_1}) \geq g_{S_1} + g_{S_2} = g_S$$

where we used that H can be decomposed as a disjoint union of S_2 , $H \cap T_1$ and $H \cap S_1$ and that $-g_{T_1} \geq -g_{T_1 \cap H}$ because g_j is negative for all $j \in T_1$. It follows that the value of the maximum flow is at least g_S , as desired. Note that the maximum flow cannot exceed g_S , because deleting all links between s and S is a valid cut with value g_S . Thus the value of the maximum flow is exactly g_S . When $|H| > N/2$, an analogous argument can be used with $W \setminus H$ instead of H .

The following Lemma will be useful.

Lemma 1 *Let Z be a random variable such that $|Z| \leq c$ almost surely. Then $\sigma_Z \leq c$, and this bound is sharp.*

Proof. Let $G(z_0)$ be the family of probability distributions of random variables that satisfy $|Z| \leq c$ and $\mathbb{E}Z = z_0$. This family of measures is tight, and by Prokhorov's theorem, it is relatively compact in the weak topology. The variance of a random variable with distribution $G \in G(z_0)$ is $\int (z - z_0)^2 G(dz)$. Here the integrand is a bounded continuous function because G is concentrated on the $[-c, c]$ interval, implying that variance is a continuous function with respect to the weak topology on $G(z_0)$, and then compactness implies that there exists $G^* \in G(z_0)$ that has the highest variance.

Let Z^* be a random variable with distribution G^* . Suppose that $|Z^*| < c$ with positive probability; then there is some $\varepsilon > 0$ such that $\Pr(|Z^*| < c - \varepsilon) > \varepsilon$. Let Y be a random variable independent of Z^* that assigns equal probabilities to $+1$ and -1 , and consider $Z' = Z^* + \chi\{|Z^*| < c - \varepsilon\} \cdot \varepsilon Y$ where $\chi\{|Z^*| < c - \varepsilon\}$ is an indicator function. It is easy to see that $\mathbb{E}Z' = \mathbb{E}Z = z_0$, and that $|Z'| \leq c$; thus G' , the probability distribution of Z' , is in $G(z_0)$. The variance of Z' can be written as

$$\Pr(|Z^*| \geq c - \varepsilon) \cdot \mathbb{E} \left[(Z^* - z_0)^2 \mid |Z^*| \geq c - \varepsilon \right] + \Pr(|Z^*| < c - \varepsilon) \cdot \mathbb{E} \left[(Z^* - z_0 + \varepsilon Y)^2 \mid |Z^*| < c - \varepsilon \right].$$

The second term on the right hand side is

$$\Pr(|Z^*| < c - \varepsilon) \cdot \left\{ \mathbb{E} \left[(Z^* - z_0)^2 \mid |Z^*| < c - \varepsilon \right] + \varepsilon^2 \right\} \geq \Pr(|Z^*| < c - \varepsilon) \cdot \mathbb{E} \left[(Z^* - z_0)^2 \mid |Z^*| < c - \varepsilon \right] + \varepsilon^3$$

where we used that Y has mean zero, unit variance, and is independent of Z^* . Combining this bound with the previous expression yields $\text{var}[Z'] \geq \mathbb{E}(Z^* - z_0)^2 + \varepsilon^3$ which contradicts the optimality of Z^* . It follows that the optimal Z^* must satisfy $|Z^*| = c$ with probability one. Given that $\mathbb{E}Z^* = z_0$

must also hold this implies that $Z^* = c$ with probability $1/2 + z_0/2c$ and $Z^* = -c$ with remaining probability. The variance of Z^* is then $(c - z_0)^2$, which is maximal when $z_0 = 0$, and the highest possible variance of Z that satisfies $|Z| \leq c$ is c^2 .

Proof of Proposition 1

(i) We prove the more general result that when the MRS is bounded from above by K , no Pareto-optimal allocation can be implemented when $a[F] < \underline{\sigma}/K$ for some F . Let $(e_i)_{i \in W}$ and $(e'_i)_{i \in W}$ be two endowment realizations, and i and j two agents. Wilson (1968) shows that any Pareto-efficient arrangement satisfies the first order condition

$$\frac{\partial U_i / \partial x_i(x_i, c_i)}{\partial U_i / \partial x_i(x'_i, c_i)} = \frac{\partial U_j / \partial x_j(x_j, c_j)}{\partial U_j / \partial x_j(x'_j, c_j)}$$

where x_i and x'_i are the goods consumption of agent i in the two endowment realizations. This first order condition simply states that in any Pareto-optimal arrangement, the marginal rates of substitution across different states are equalized for all agents. This equation implies that in a Pareto-optimal arrangement, agents have the same cardinal ranking for all states of the world: if agent i prefers $(e_i)_{i \in W}$ to $(e'_i)_{i \in W}$, then so does agent j .

Let $\underline{a} = \min_{|F| \leq N/2} a[F]$. By assumption, $\sigma_i \geq K\underline{a}$ for all agents i . By Lemma 1, this implies that for every i there exists m_i such that e_i assumes both values below m_i and values above $M_i = m_i + K\underline{a}$ with positive probability. Now suppose that, contrary to the assertion in the Proposition, a Pareto-optimal incentive compatible arrangement exists. Let F be a set with $a[F] < \underline{\sigma}/K$. By Assumption 1, the set of realizations $(e_i)_{i \in W}$ where for all $i \in F$, $e_i > M_i$, and for all $j \notin F$, $e_j < m_j$ form a positive probability event. Note that for any such realization, the total goods consumption of agents in F is at least $\sum_{i \in F} M_i - Kc[F]$, where the second term is the maximum amount that can leave the set F . Similarly, the total goods consumption of agents outside F is at most $\sum_{i \notin F} m_i + Kc[F]$.

Now consider a second set of realizations (e'_i) , where for all $i \in F$ we have $e'_i < m_i$, and for all $j \notin F$ we have $e'_j > M_j$. By assumption, this set of realizations also has positive probability. For each such realization, the total consumption of agents in F is at most $\sum_{i \in F} m_i + Kc[F]$, and the total consumption of agents in $W \setminus F$ is at least $\sum_{i \notin F} M_i - Kc[F]$.

These results imply that the total consumption of F is higher in (e_i) than in (e'_i) , since the

lower bound in (e_i) equals the upper bound in (e'_i) :

$$\sum_{i \in F} M_i - Kc[F] = \sum_{i \in F} m_i + Kc[F]$$

holds because $M_i - m_i = K\underline{a} = Kc[F]/|F|$ by definition. In contrast, the total consumption of agents in $W \setminus F$ is higher in (e'_i) than in (e_i) , because the upper bound in (e_i) does not exceed the lower bound in (e'_i) :

$$\sum_{i \notin F} m_i + Kc[F] \leq \sum_{i \notin F} M_i - Kc[F]$$

since $\sum_{i \notin F} M_i - m_i = (N - [F])K\underline{a} = Kc[F] \cdot (N - [F])/|F| \geq Kc[F]$. Taken together, these findings violate the property of Pareto optimality that agents have a common ranking over states of the world: there must be some agent in F who prefers (e_i) to (e'_i) , and at the same time there must be some agent in $W \setminus F$ who prefers (e'_i) to (e_i) . This is a contradiction.

To show that there is an agent in F who is not fully insured, note that there is some $i \in F$ who prefers a positive probability of realizations $(e_i)_{i \in W}$ to a positive probability of realizations $(e'_i)_{i \in W}$; and within these events, there is $j \notin F$ who prefers a positive probability of (e'_i) to (e_i) . But this means that we could improve the expected utility of i and j by transferring a small amount from i to j in (e_i) and transferring a small amount back from j to i in (e'_i) .

(ii) Theorem 1 provides necessary and sufficient conditions for implementing the profile when all agents consume zero. For any set F , the excess demand is at most $|F| \cdot \sigma$, which must be less than or equal to $c[F]$. This is equivalent to $\sigma \leq a[F]$ for all F with $|F| \leq N/2$.

Proof of Proposition 2

(i) Consider an N long segment on the line, and split it into intervals of length k . For each segment F , $\sigma_F = \sigma\sqrt{k}$ and $c[F] = 2c$. Using Lemma 1, this implies for any IC arrangement

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

To obtain the sharpest bound, let $k = 16(c/\sigma)^2$, which gives

$$SDISP \geq \sigma \cdot \frac{1}{8} \frac{\sigma}{c}$$

as desired.

(ii) We establish a considerably more general result. We begin by listing our assumptions:

(D1) [Thin tails] The y_j variables are independent, have zero mean and unit variance, and satisfy $\log \text{Eexp}[\theta y_j] \leq K\theta^2/2$ for some K with all $\theta \geq 0$.

(D2) [Correlated shocks] The endowment shock of agent i is $e_i = \sum_j \alpha_{ij} y_j$ where $\sum_j \alpha_{ij}^2 < \infty$.

Here (D1) is a uniform bound on the moment-generating function of y_j , and allows us to use the theory of large deviations to bound the tails of weighted sums of y_j . This condition is satisfied if the y_j random variables are i.i.d. normal, or if they have a common compact support, and in many other cases.

(E1) [No aggregate risk] Endowments satisfy $\sigma_F/|F| \leq K_1 \cdot |F|^{-K_2}$ for some $K_1, K_2 > 0$.

(E2) [More people have more risk] For all $G \subseteq F$, we have $\sigma_G \leq \sigma_F$.

(E3) [Sharing with more people is always good.] For all $G \subseteq F$, we have $\sigma_F/|F| \leq \sigma_G/|G|$.

(N1) The network is connected, countably infinite, and there exists a constant K_3 such that all agents have at most K_3 direct friends.

We also impose a set of conditions on the network that allows for a decomposition similar to the square structure on the plane. Specifically, we require that for all $m \geq 1$ integers there exist a collection of sets F_j^i , where $i = 1, \dots, m$ and $j = 1, \dots, \infty$ that satisfy the following properties.

(S1) [Partition] For all $1 \leq i \leq m$, the sets $F_j^i, j = 1, \dots, \infty$ give a partition of the set of agents; and when $i = 1$, all sets F_j^1 are singletons.

(S2) [Ascending chain] For all $1 \leq i \leq m - 1$ and all j, j' , we have either $F_j^i \cap F_{j'}^{i+1} = \emptyset$ or $F_j^i \subseteq F_{j'}^{i+1}$.

(S3) [Exponential growth.] There exist $1 < \underline{\gamma} < \bar{\gamma}$ constants such that whenever $F_j^i \subseteq F_{j'}^{i+1}$, we have $\underline{\gamma} |F_j^i| \leq |F_{j'}^{i+1}| \leq \bar{\gamma} |F_j^i|$.

Let $G \subseteq F$ be two sets, and define the relative perimeter of G in F , denoted $c_0[G]_F$, as the perimeter of G in the subgraph generated by the set of vertices F .

(S4) [Relative perimeter] There exists $K > 0$ such that for any $G \subseteq F_j^i$ with $|G| \leq |F_j^i|/2$ we have $c_0[G]_{F_j^i} \geq K \cdot c_0[G]$.

Here (S1) means that for each i , the i -level sets partition the entire network. (S2) 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. (S3) requires that the size of these sets grows exponentially, and (S4) means that the F_j^i sets are good “snapshots” of the network: the perimeter of sets inside F_j^i is proportional to their total perimeter. For the plane a decomposition with squares generates a partition that satisfies these properties; since there are 4 identical squares in each larger square, we can set $\underline{\gamma} = \bar{\gamma} = 4$.

The following is the key substantive condition.

(K) [Key perimeter/area condition] There exists $c > 0$ such that for all i, j , and for all $G \subseteq F_j^i$ with $|G| \leq |F_j^i|/2$, we have $\sigma_G \leq c \cdot c_0 [G]_{F_j^i}$.

This is just a version of the usual perimeter/area ratio condition. Together with (S4), it implies that there exists $c' > 0$ such that for all i, j , and for all $G \subseteq F_j^i$ with $|G| \leq |F_j^i|/2$, we have $\sigma_G \leq c' \cdot c_0 [G]_{F_j^i}$. To simplify notation, we redefine c_0 to be $c' \cdot c_0$. Then we have the following result.

Proposition 7 *Under the above conditions, there exist constants K_4 and K_5 such that $SDISP(c) \leq K_4 \exp[-K_5 \cdot c^{2/3}]$.*

Proposition 2 (i) is an immediate consequence of this result. The proof is the following.

Intuition. Fix m , and consider the decomposition described above. Our strategy is to construct an unconstrained flow that implements equal sharing within each set F_j^m , $j = 1, \dots, \infty$ (these are the “biggest squares” in the plane). We then compute the implied capacity use of this flow for each link. Then we choose m and c simultaneously in such a way that the unconstrained flow actually satisfies all capacity constraints “most of the time.” This flow will then implement essentially equal sharing for the F^m sets, and hence by (E1) implements exponentially small $SDISP$. But we also have to bound the impact of the exceptional events when the flow does hit some capacity constraints. This results in an additional term that is sub-exponential in c ; and combining these two terms leads to the bound of the proposition.

Induction logic. The unconstrained flow is constructed iteratively, 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 $i + 1$. Let $n(F_j^{i+1}) = \left| \left\{ j' \mid F_{j'}^i \subseteq F_j^{i+1} \right\} \right|$ denote the number of i level sets in F_j^{i+1} . By (S3), there exists some K_6 such that for all i and j we have $2 \leq n(F_j^{i+1}) \leq K_6$. 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, \dots, F_k where $k \leq K_6$. We know from (S1) and (S2) that F_1, \dots, F_k partition F . We will now 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 .

Induction step. To smooth the total consumption of F_1 in F , first note that this quantity is the same as the total endowment in F_1 , because in each round i , we are smoothing all endowments within an i -level set. Second, having completely smoothed resources in F_1 in the previous round,

currently all agents in F_1 are allocated $e_{F_1}/|F_1|$ units of consumption (for a total of $|F_1| \cdot e_{F_1}/|F_1| = e_{F_1}$.)

Flow construction. To smooth this over F , we now define a flow. This is a key step in the proof. For this flow, focus on the subgraph generated by F together with the original 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 $F \setminus F_1$ 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, we use Theorem 1; this is where the key perimeter/area condition (K) plays its role. According to the theorem, we can implement equal sharing if for each set $G \subseteq F$ with $|G| \leq |F|/2$, the excess demand for goods does not exceed the perimeter (relative to F). What is this excess demand? Since we want equal sharing, we should allocate $\sigma_{F_1}/|F|$ to every agent in G . But those guys in G who are also in F_1 each have $\sigma_{F_1}/|F_1|$. So the excess demand for goods in G is

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

If there is a feasible flow, then for every G , the absolute value of this excess demand $ed(G)$ should be less than $c_0[G]_F$. To check that this holds, first assume that $|G|/|F| \geq |G \cap F_1|/|F_1|$; then the above formula implies $|ed(G)| \leq \sigma_{F_1} \cdot |G|/|F|$. But from (E2) we have $\sigma_{F_1} \leq \sigma_F$, implying $|ed(G)| \leq \sigma_F \cdot |G|/|F|$; and from (E3), $\sigma_F/|F| \leq \sigma_G/|G|$, which then implies that $|ed(G)| \leq \sigma_G$. Now recall our 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 this inequality also holds when $|G|/|F| < |G \cap F_1|/|F_1|$. In this case, the formula for $ed(G)$ displayed above implies $|ed(G)| \leq \sigma_{F_1} \cdot |G \cap F_1|/|F_1|$. Since $\sigma_{F_1}/|F_1| \leq \sigma_{G \cap F_1}/|G \cap F_1|$ by (E3), we can bound the right hand side from above by $\sigma_{G \cap F_1}$, which satisfies $\sigma_{G \cap F_1} \leq \sigma_G \leq c_0[G]_F$ and we are done.

So the proposed flow can indeed be implemented. Let the associated transfers be denoted by t_1 . To get a flow smoothing 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 . Extending this logic, to smooth the endowment of each F_j through the set F , we just construct t_2, \dots, t_k analogously, and implement the transfers $t_1 \cdot e_{F_1}/\sigma_{F_1} + \dots + t_k \cdot e_{F_k}/\sigma_{F_k}$. It is important

to note that this construction results in an unconstrained flow. While we used the capacities to construct the flow (this is how we got t_1, \dots, 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 .

Repeat. 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^m , $j = 1, \dots, \infty$. How low is SDISP at the end of this procedure? To answer, recall that (S3) implies $|F_j^m| \geq \underline{\gamma}^m$, and (E1) implies $\sigma_F/|F| \leq K_1 \cdot |F|^{-K_2}$, so that $SDISP \leq K_1 \cdot \underline{\gamma}^{-K_2 m} = K_1 \cdot \exp[-K_7 m]$ for some $K_7 > 0$. 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 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 that the standard deviation of the link use in this flow is at most $c_0(u, v)$.

Now consider link use as we smooth the consumption of all sets F_1, \dots, F_k over the set F . By the argument of the previous paragraph, as we smooth each of these sets, we add a flow over the (u, v) link that is normally distributed, and has a standard deviation of at most $c_0(u, v)$. So the total standard deviation of the flow over (u, v) generated in one round of the algorithm is at the most $K_6 \cdot c_0(u, v)$. Finally, we have to add up these flow demands over all m rounds; thus the total standard deviation of the flow demand over a link is at most $mK_6 \cdot c_0(u, v)$.

Bounding exceptional event. To bound the contribution of the exceptional events to SDISP, we first need to specify what is the constrained flow. We do the following: Fix some c and m , and for each agent u , try to implement his inflows and outflows according to the unconstrained flow corresponding to m we just constructed. If this is not possible because some of his constraints are hit, we implement as much of the prescribed flows as possible. This procedure assumes that binding constraints do not propagate down the network.

Consider some agent $u \in F_j^m = F$, and suppose that the constraint binds the unconstrained flow over an (u, v) link, but on no other link of u . The contribution of this event to the variance of

u 's consumption is bounded by

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

where $\bar{e}_F = e_F/|F|$ and P is the probability measure. This can be bounded from above by

$$6 \int_{t(v,u) > c(v,u)} \bar{e}_F^2 + [t(v,u) - c(v,u)]^2 dP \leq 6 \int \bar{e}_F^2 dP + 6 \int_{t(v,u) > c(v,u)} [t(v,u) - c(v,u)]^2 dP. \quad (6)$$

Here the first term is six times the unconstrained DISP, and the second term is six times the integral of the square of the random variable $(t(u,v))$ on a tail event. We now bound this latter term using (D1), using the theory of large deviations.

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 $\theta > 0$,

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

Now we can bound the last term using (D1) to obtain

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

This holds for any θ , in particular, for $\theta = c / \left[K \sum \alpha_j^2 \right]$, resulting in the bound

$$\Pr [z > c] \leq \exp \left[-\frac{c^2}{2K \sum \alpha_j^2} \right]$$

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 (6). Let $z = t(u,v)$ which is a weighted sum of the y_j shocks by construction; denoting the c.d.f. of z by $H(z)$ we have

$$\begin{aligned} \int_{t(v,u) > c(v,u)} [t(v,u) - c(v,u)]^2 dP &= \int_{z=c(u,v)}^{\infty} (z - c(u,v))^2 dH(z) = - \int_{z=c(u,v)}^{\infty} (z - c(u,v))^2 d[1 - H(z)] = \\ &= - \left[(z - c(u,v))^2 (1 - H(z)) \right]_{c(u,v)}^{\infty} + \int_{z=c(u,v)}^{\infty} 2(z - c(u,v)) [1 - H(z)] dz \end{aligned}$$

where we integrated by parts. The above argument with large deviations implies $1 - H(z) \leq \exp[-z^2/2K\sigma_z^2]$, and hence

$$\int_{t(v,u) > c(v,u)} [t(v,u) - c(v,u)]^2 dP \leq K_8 c(u,v) \exp[-c(u,v)^2/2K\sigma_z^2].$$

Since the standard deviation of $z = t(u,v)$ is at most $mK_6c_0(u,v)$ and $c(u,v) = c \cdot c_0(u,v)$, the last term is bounded by $K_8 \cdot \exp[-K_9 \cdot (c/m)^2]$. We need to similarly bound the contribution of the exceptional event for every other link of u where the constraint may be violated, and for every pair of links, and so on. Since u has a bounded number of links, this will just increase the above bound by a constant factor. So overall, the exceptional event contributes a constant times the unconstrained SDISP plus $\exp[-K_9 \cdot (c/m)^2]$ to the SDISP of the constrained allocation. So in total, we have

$$SDISP \leq K_7 \cdot \left[\exp[-K_7m] + \exp[-K_9 \cdot (c/m)^2] \right].$$

Now let $m = c^{2/3}$, then we obtain

$$SDISP \leq K_4 \cdot \exp[-K_5c^{2/3}]$$

as desired.

Proof of Corollary 1

We will pick a large enough grid that does not intersect with any of the agents. Inside each of the squares in the grid, we can get good risk-sharing because there are only a bounded number of people. To share across squares, we make use of the result that we have good risk-sharing on the plane, and the fact that the squares in the grid approximate a plane.

Pick a grid size $g > K_1$, and place a grid with this stepsize on the plane that does not overlap with any agent (this is possible, because there are only a countable number of agents). By (P1), under capacities c_0 there is a capacity of at least 1 between any pair of adjacent squares on the grid. If some pairs of adjacent squares have capacity exceeding 1, then delete some links or reduce capacities such that in the resulting network the capacity between any pair of adjacent squares is exactly 1. Index the squares in the grid by $j = 1, \dots, \infty$ and denote the set of agents in square j by G_j .

We know from (P2) that the subgraph spanned by G_j is connected, and that $K_1 \leq |G_j| \leq K_1^3$. The upper bound here implies that we have $\sigma_{G_j} \leq K_7$ for $K_7 = K_1^3 \bar{\sigma}$ by assumptions (E4) and

(E5). We now do the following. Pick c , and use capacity $c/10$ to implement between-squares risk-sharing using the previous Proposition, taking e_{G_j} as the “endowment shocks” of the squares. To see that the previous result is indeed applicable, we only need to check the key condition (K). But that holds, because for any union of grid-squares G , we have $\sigma_G \leq K_8 |G|^{1/2}$ (here $|G|$ refers to the number of grid-squares in G) by condition (E5). This allocation generates between-squares dispersion which is exponentially small in $c^{2/3}$.

Now we just have to smooth the incoming and outgoing transfers as well as the endowments within each square. Use capacity $4c/10$ within the squares to smooth all incoming and outgoing transfers across all agents in the square. Since the total perimeter of the square is $4c/10$, and G_j is connected, this can be done with probability one. Now it remains to smooth the total endowment shock realized in the square, and we still have capacity $c/2$ to do this. Since the square is a finite network, the number of agents and the variances of all shocks are bounded, and all pairs of agents are connected by a path of at least a constant capacity, we know that we can achieve within-square dispersion on the order of $\exp[-K(c/2)^2]$ with shocks satisfying (D1) and (D2). This is smaller than the main $\exp[-c^{2/3}]$ term; hence the proof is complete.

Proof of Proposition 3

We have

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

where t_F denotes the total transfer leaving F . Using Lemma 1, $|\text{cov}[t_F, e_F]| \leq c[F] \cdot \sigma(e_F)$, which implies the claim of the proposition.

Proofs for Section 2.6

The following result is discussed in the text in Section 2.6:

Proposition 8 *Assume that MRS_i is increasing in x_i for all i , then for any pair of endowment realizations \underline{e} and \bar{e} such that $\underline{e}_i \leq \bar{e}_i$ for all i , the set of IC transfer arrangements in \underline{e} is a subset of the set of IC transfer arrangements in \bar{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 $U_i(x_i + s_i, c_i) \geq U_i(x'_i + s_i, c'_i)$.

Now let t be an IC transfer arrangement under \underline{e} and let \underline{x} be the associated consumption profile. Incentive compatibility implies $U_i(\underline{x}_i, c_i) \geq U_i(\underline{x}_i + t_{ij}, c - c(i, j))$. Denoting $\bar{e} - \underline{e} = s \geq 0$, by the single crossing property we have $U_i(\underline{x}_i + s_i, c_i) \geq U_i(\underline{x}_i + s_i + t_{ij}, c - c(i, j))$. But the consumption profile resulting from the transfer arrangement t under realization \bar{e} is exactly $\bar{x} = \underline{x} + s$, and hence the last inequality shows that t is IC under \bar{e} as well.

Proof of Proposition 4

We prove the following more general result.

Proposition 9 *Suppose that the MRS_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: $U^* = \{(v_i)_{i \in W} \mid \exists \text{ IC allocation } x \text{ such that } v_i \leq EU_i(x_i, c_i) \forall i\}$. 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 $(y(r, x_i), c_i - r)$ 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

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

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 that concavity of the MRS, we will

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

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

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 $f(r) \leq f(0)$ or equivalently $U(\bar{y}(r), c_i - r) \leq U(x_i^3, c_i)$, which implies that $y(x_i^3, r) \geq \bar{y}(r) = \alpha y(r, x_i^1) + (1 - \alpha)y(r, x_i^2)$, and hence that $y(x, r)$ 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^* .

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 $V(t) = \sum_i \lambda_i U_i(e_i - \sum_j t_{ij}, c_i)$. 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 , $\bar{V}(x) = V(t)$. 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 $i \rightarrow j$ 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 Corollary 2

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 $i_1 \rightarrow i_k$ such that $i_1 = i_k$, and $t_{i_l i_{l+1}} > 0 \forall l \in \{1, \dots, k-1\}$. 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 establish a stronger monotonicity property. Say that a transfer arrangement is strongly monotone 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 $i \in F$. Strong monotonicity means that for any set of agents F , reducing their endowments and/or their incoming transfers weakly reduces everybody's consumption.

Fix a constrained efficient arrangement, and suppose that is not strongly monotone. Let F be a set where this property fails, and fix a connected component of the subgraph spanned 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 an IC transfer arrangement achieving $x(e)$ at endowment realization e , such that $t_{kl} < c(k, l) \forall (k, l) \in L_i$. Let $b = \min_{(k, l) \in \widehat{L}_i} (c(k, l) - |t_{kl}|)$. Let L'_i denote

the set of links connecting agents in $W(i)$ with agents in $N/W(i)$. For every $(k, l) \in L'_i$, let t'_{kl} be such that $\lambda_k U'_k(x_k(e) - t'_{kl}) = \lambda_l U'_l(x_l(e) + t'_{kl})$. By Proposition 5, $t'_{kl} \neq 0 \forall (k, l) \in L'_i$. Let $b' = \min_{(k, l) \in L'_i} |t'_{kl}|$. Let $b^0 = \min(b, b')$. Recall that $\lambda_j U'_j(x_j(e)) = \lambda_i U'_i(x_i(e)) \forall j \in W(i)$. Then for any $|e_i - e'_i| < b^0$ there is transfer arrangement t'' such that (i) $t + t''$ is IC; (ii) $t''_{ij} < b^0 \forall (i, j) \in L$; (iii) $t''_{ij} = 0$ whenever $i \notin W(i)$ or $j \notin W(i)$; (iv) the first-order conditions of Proposition 5 hold for any $(i, j) \in L_i$. Then (ii), (iii), and $b^0 \leq b'$ imply that the first-order conditions of Proposition 5 hold for any $(i, j) \in L'$. Moreover, (iii) implies that the first-order conditions of Proposition 5 hold for any $(i, j) \in L$ such that $i, j \notin W(i)$. Therefore the first-order conditions of Proposition 5 hold for every link at consumption vector $e' + t + t''$ after endowment realization e' . Proposition 5 then implies that $e' + t + t''$ is the constrained efficient consumption vector after e' . Note that $b^0 \leq b$ and (ii) above imply $t_{ij} + t''_{ij} < c(i, j) \forall (i, j) \in \widehat{L}_i$, hence by Proposition 5 $\lambda_j U'_j(x_j(e')) = \lambda_i U'_i(x_i(e')) \forall j \in \widehat{W}(i)$. Finally, (iii) above implies $x_j(e') = x_j(e) \forall j \notin W(i)$.

(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, $x_{j'}(e') \leq x_{j'}(e) \forall j' \in W$. Therefore if $x_j(e') = x_j(e)$ then the statement in the claim holds. Assume now that $x_j(e') < x_j(e)$. Since $x_l(e') \leq x_l(e) \forall l \in W$, for any $l \in W \setminus \{i\}$ it must hold that $\sum_{l' \in W \setminus \{l\}} t''_{l'l} \leq 0$. This, together with $x_j(e') < x_j(e)$ implies that there is a $j \rightarrow i$ 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 $j \rightarrow i$ 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 $i \rightarrow j$ 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.

Formal results for Section 3.3

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, \dots, 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 $m = 1, 2, \dots$, go through the links l_1, \dots, 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 $\lambda_{i_k} U'_{i_k}(x'_{i_k}) > \lambda_{j_k} U'(x'_{j_k})$ and $t'_{i_k j_k} = -c(i, j)$, or $\lambda_{i_k} U'_{i_k}(x'_{i_k}) < \lambda_{j_k} U'(x'_{j_k})$ 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 $m = 1, 2, \dots$ 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 is IC by design for every m .

Proposition 10 *If consumption and friendship are perfect substitutes, then $x^m \rightarrow x^*$ as $m \rightarrow \infty$.*

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 $V(x_1) \leq V(x_2) \leq \dots$, and since $V(x) \leq V(x^*)$ for all x that are IC, we have $\lim_{m \rightarrow \infty} V(x_m) = V \leq V(x^*)$. 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 an IC transfer 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^* .

Constrained efficient arrangements are robust to coalitional deviations

Given an IC arrangement t , an ex ante coalitional deviation by a set of agents F is a transfer arrangement t' among agents in F that takes on values conditional on the endowment realizations of all agents in W . An ex ante coalitional deviation is profitable if t' is IC, and given transfer arrangement that specifies transfers t' for agents in F and transfers according to t otherwise, all agents in F are weakly better off and one of them is strictly better off than given transfer

arrangement t . Note that we do not require the transfer arrangement that specifies transfers t' for agents in F and transfers according to t otherwise to be IC. That is, a coalitional deviation might involve withholding transfers specified by t to agents outside F , for some state realizations.¹²

Proposition 11 *If goods and friendship are perfect substitutes, then an IC allocation is constrained efficient if and only if it admits no profitable ex ante coalitional deviations.*

Proof: Fix an IC arrangement x , and consider an ex ante profitable coalitional deviation t' by agents in a coalition F . With perfect substitutes, incentive compatibility of a transfer over a link connecting F with $W \setminus F$ depends only on the capacity of the link. As a result, following the coalitional deviation, agents in F will not default on their promised transfers with agents outside F . This means that the allocation x' achieved by the coalitional deviation is IC, weakly improves all agents' utility, and strictly improves some agent's utility, showing that x is not constrained efficient. Conversely, for any IC x that is Pareto dominated by some IC allocation x' , there exists a coalitional deviation in x where the coalition of all agents $F = W$ chooses x' .

It is easy to see that nonexistence of profitable ex ante coalitional deviations implies nonexistence of ex post coalitional deviations, too. In fact, any IC transfer arrangement is immune to ex post coalitional deviations. The proofs of these claims are straightforward, hence omitted.

First-order conditions for general preferences

To present our characterization result for general preferences, 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 2, let t be an acyclical implementation of x in endowment realization e . Consider the following iterative construction. Let $W^1 \subseteq W$ denote the set of agents i for whom there is no j such that $c(i, j) > 0$ and the IC constraint from i to j binds. Since t is acyclical, W^1 is nonempty. For any $i \in W^1$, let $\Delta_i = \lambda_i U_{i,x}(x_i, c_i)$, 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, \dots, 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

¹²This definition of a coalitional deviation is closely related to the concept of side-deal proof equilibrium in Mobius and Szeidl (2007).

that the IC constraint from i to j binds, $\hat{x}_{i,j} = x_i + t_{ij}$, and $\hat{c}_{i,j} = c_i - c(i, j)$, and let

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

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 $U_{i,x}(\hat{x}_{i,j}, \hat{c}_{i,j})/U_{i,x}(x_i, c_i)$ 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 12 *[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, i.e., $U_i(x_i, c_i) = U_i(x_i + t_{ij}, \hat{c}_{i,j})$, 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 $\alpha U_{i,x}(x_i, c_i) = U_{i,x}(\hat{x}_{i,j}, \hat{c}_{i,j})$, or equivalently, $\alpha = U_{i,x}(x_i, c_i)/U_{i,x}(\hat{x}_{i,j}, \hat{c}_{i,j})$. 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 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 $1 - U_{i,x}(x_i, c_i)/U_{i,x}(\hat{x}_{i,j}, \hat{c}_{i,j})$ 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 $\hat{\Delta}_i = \Delta_i(x, t)$, 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 $W^d = \{i \in W \mid \exists j \text{ such that } t_{ij}^d \neq 0\}$, and partition W^d into sets W_0^d, \dots, W_K^d the following way. Let $W_0^d = \{i \in W^d \mid \exists j \in W^d \text{ st } t_{ij}^d > 0\}$. Given W_0^d, \dots, W_k^d for some $k \geq 0$, let $W_{k+1}^d = \{i \in W^d \mid \exists j \in W^d \setminus (\cup_{l=0, \dots, k} W_l^d) \text{ st } t_{ij}^d > 0\}$. Note that $x'_i > x_i \forall i \in W_0^d$, which together with there being no agent j such that $t_{ij}^d > 0$ implies that $\Delta'_i < \hat{\Delta}_i$. Now we iteratively establish that $\Delta'_i < \hat{\Delta}_i \forall i \in W^d$. Suppose that $\Delta'_i < \hat{\Delta}_i \forall i \in \cup_{l=0, \dots, k} W_l^d$ for some $k \geq 0$. Let $i \in W_{k+1}^d$. Note that by definition there is $j \in \cup_{l=0, \dots, k} W_l^d$ such that $t_{ij}^d > 0$, and there is no $j' \in W^d \setminus (\cup_{l=0, \dots, k} W_l^d)$ such that $t_{ij'}^d > 0$. Suppose $\Delta'_i \geq \hat{\Delta}_i$. This can only be compatible with $t_{ij}^d > 0$, $\Delta'_j < \hat{\Delta}_j$, and (1)-(3) holding for both (x', t') and $(x, t' + t^d)$ if $x_i > x'_i$. But $x_i > x'_i$, and $\Delta'_{i'} < \hat{\Delta}_{i'} \forall i' \in W$ such that $t_{ii'}^d > 0$ implies $\Delta'_i < \hat{\Delta}_i$, a contradiction. Hence $\Delta'_i < \hat{\Delta}_i \forall i \in W_{k+1}^d$, and then by induction $\Delta'_i < \hat{\Delta}_i \forall i \in W^d$. But note that for any $i \in W_K^d$ it holds that $x_i < x'_i$ and there is no $j \in W$ such that $t_{ji}^d > 0$, and hence $\Delta'_i > \hat{\Delta}_i$. This contradicts $\Delta'_i < \hat{\Delta}_i \forall i \in W^d$, hence there cannot be (x', t') satisfying (1)-(3) such that t' is IC

and $x' \neq x$.

Corollary 2 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 $e_j = e'_j \forall j \in W \setminus \{i\}$. 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.

Corollary 3 *[Spillovers with imperfect substitutes] Assume that MRS_i is concave, then*

- (i) *[Monotonicity] $\Delta_j(e') \geq \Delta_j(e)$ 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 $\Delta_i = \Delta_j$ 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, $\Delta_l \leq \Delta_j$.*

The proof of this result is analogous to the perfect substitutes case and hence omitted. Note that (ii) is weaker than in Corollary 2, 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.

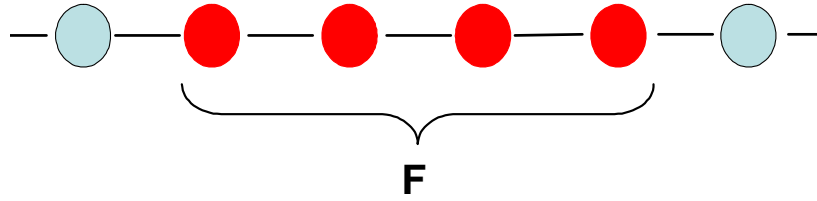
FIGURE 1: TRANSFERS IN A PERU SHANTYTOWN



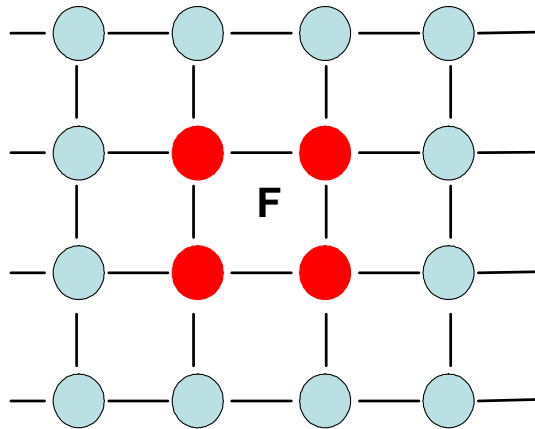
NOTE—Figure shows network of transactions in the map of a shantytown in Peru. Red links represent financial transactions, and thickness indicates value measured in soles. Blue respectively green links represent objects of high and low value. Figure is constructed using data collected by Markus Mobius.

FIGURE 2: RISKSHARING IN SIMPLE NETWORKS

A. Line network



B. Plane network



C. Binary tree network

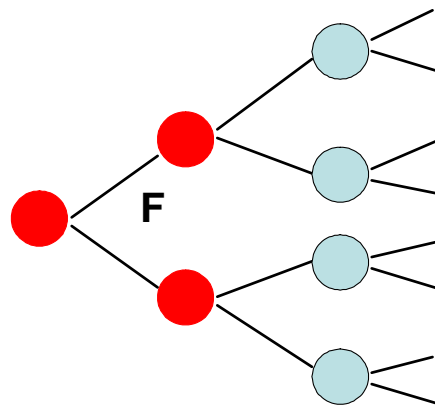
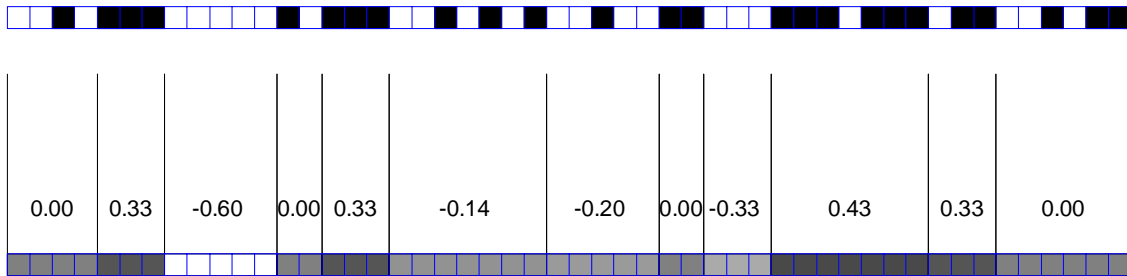


FIGURE 3: RISKSHARING SIMULATIONS

A. Line



B. Plane



NOTE—Figure shows constrained efficient allocations in the line and plane networks with binary endowment shocks. For both networks, the black and white panel represents the endowment realizations and the grey panel represents the “second best” risksharing arrangement with equal planner weights. The plane allows for better risksharing: in this realization, $SDISP = 31\%$ for the line and $SDISP = 0\%$ for the plane. The line network also illustrates the emergence of risksharing islands: there is perfect smoothing and equal consumption within islands, but imperfect smoothing across boundaries.

FIGURE 4: UTILITY COST OF SHOCKS TO DIRECT AND INDIRECT FRIENDS

