Consumption Risk-sharing in Social Networks
Online Appendix

By ATTILA AMBRUS, MARKUS MOBIUS AND ADAM SZEIDL*

This material contains proofs and extensions which supplement the paper "Consumption Risk-sharing in Social Networks". We provide the following analysis. (1) We provide missing proofs for results stated in the main paper. (2) We provide game theoretic micro-foundations to justify our assumption that links "die" when a promised transfer is not made. (3) We provide background about the mathematical theory of network flows used in the proofs of the paper. (4) We formalize two decentralized mechanisms leading to constrained efficient allocations. (5) We formally develop the extensions of our main results to the case where goods and friendship consumption are imperfect substitutes. (6) We explain the numerical methods used to simulate the model and to construct the geographic network representation of the real world Huaraz network.

A-1. Missing Proofs for Sections I to III

Proof that coalition-proof arrangements are robust to deviating subcoalitions

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

PROPOSITION 7: *Requiring coalitions to be robust to further coalitional deviations does not affect the set of coalition-proof allocations.*

PROOF:

Let $\mathcal{F}$ be a deviating coalition, and let $\mathcal{F}' \subseteq \mathcal{F}$ be a deviating subcoalition. Then $\mathcal{F}'$ is also a deviating coalition in the original set of agents $\mathcal{W}$. To see why, note that the capacities $\widetilde{c}'$ after the subcoalition $\mathcal{F}'$ deviates are exactly those associated with links within $\mathcal{F}'$, and hence these are also the capacities that remain when $\mathcal{F}'$ deviates in $\mathcal{W}$. Moreover, the allocation $\widetilde{\mathbf{x}}'$ of the subcoalition $\mathcal{F}'$ only uses the resources in $\mathcal{F}'$ and hence is also feasible when $\mathcal{F}'$ deviates from $\mathcal{W}$. These observations imply that the same allocation is available to all agents in $\mathcal{F}'$ if they consider a coalitional deviation from $\mathcal{W}$. Since these agents are better off with this allocation than they were in the coalition $\mathcal{F}$, where in turn they are better off than in the original allocation, it follows that $\mathcal{F}'$ is a profitable coalitional deviation in the original network as well. Hence minimal deviating

* E-mail: attila.ambrus@duke.edu, markusmobius@gmail.com, szeidla@ceu.hu .

coalitions are robust to further coalitional deviations. Since any allocation that has a deviating coalition also has a minimal one, requiring no deviating coalitions is equivalent to requiring no deviating coalitions that are robust to further group deviations.

Proof of Proposition 1

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

$$|\mathcal{F}|\, M - |\mathcal{F}|\, [|\mathcal{F}|\, M + (|\mathcal{W}| - |\mathcal{F}|)m]/\, |\mathcal{W}|$$

This amount has to be less or equal to the group's obligation which equals the perimeter $c[\mathcal{F}]$. Some algebra reduces this inequality to $a\,[\mathcal{F}] \geq \left(1 - \frac{|\mathcal{F}|}{|\mathcal{W}|}\right) S$. Hence the perimeter-area inequality implies that no group will want to deviate even in the worst case scenario. For the same reason, coalition proofness implies the perimeter-area inequality because the coalitional IC constraint $e_{\mathcal{F}} - x_{\mathcal{F}} \leq c\,[\mathcal{F}]$ has to hold for all states of the world.

Infinite networks in subsection II B

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

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

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

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

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

constructed a feasible allocation $\mathbf{x}'$ in H. If this allocation is coalition-proof for every finite subgraph H, then we say that the original allocation $\mathbf{x}$ is coalition-proof in the infinite network G.

Extending Theorem 1. An informal risk-sharing arrangement can be defined in the same way as before. We claim that in this infinite network environment, the statement of Theorem 1 holds word by word. As in the finite case, sufficiency is immediate. To prove necessity, let $\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \ldots$ be an increasing sequence of sets such that $\cup_k \mathcal{H}_k = \mathcal{W}$, and fix a coalition-proof allocation $\mathbf{x}$. For each k, construct the auxiliary network H_k as above. We can define $g_i = e_i - x_i$ for all $i \in \mathcal{H}_k$ as in the proof of Theorem 1, and let $g_z = -\sum_{i \in \mathcal{H}_k} g_i$; with this definition, we have constructed a finite implementation problem just like in the proof of Theorem 1. Since we have a coalition-proof allocation in H_k, Theorem 1 yields an informal risk-sharing arrangement t^k in H_k. For each (i, j) link we thus obtain a sequence of transfers $t^k_{ij} \in \left[-c\,(i, j)\,, c\,(i, j)\right]$ for the infinite sequence of k values for which $i, j \in \mathcal{H}_k$. Because there are only countably many links, we can select a subsequence that converges to some t^*_{ij} pointwise for each i and j. It is immediate that this transfer arrangement implements consumption allocation $\mathbf{x}$ and meets the capacity constraints.

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

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

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

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

$$DISP\,(\mathbf{x}) = \lim_{r \to \infty} \frac{1}{\left|\mathcal{B}^r_i\right|} \sum_{j \in \mathcal{B}^r_i} \left(x_j - \overline{e}\right)^2 .$$

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

Proof of Proposition 2

The following Lemma is used in the proof.

LEMMA 1: *Let Z be a random variable such that $|Z| \leq c$ almost surely. Then $\sigma_Z \leq c$.*

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

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

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

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

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

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

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

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

(N4) [Exponential growth.] There exist $1 < \underline{\gamma} < \overline{\gamma}$ constants such that whenever $\mathcal{F}_j^i \subseteq \mathcal{F}_{j'}^{i+1}$, we have $\underline{\gamma}\left|\mathcal{F}_j^i\right| \leq \left|\mathcal{F}_j^{i+1}\right| \leq \overline{\gamma}\left|\mathcal{F}_j^i\right|$.

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

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

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

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

(K) There exists $K > 0$ such that for all $\mathcal{G}$ finite, $\sigma_{\mathcal{G}} \leq K \cdot c_0 [\mathcal{G}]$.

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

PROPOSITION 8: *Under conditions (P1)-(P5), (N1)-(N5) and (K), there exist positive constants K' and K'' and a coalition-proof allocation $x(c)$ such that for every agent i, $Ex^2_i(c) \leq K' \exp\left[-K'' \cdot c^{2/3}\right]$.*

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

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

(K') For all $\mathcal{G} \subseteq \mathcal{F}^i_j$ with $|\mathcal{G}| \leq \left|\mathcal{F}^i_j\right|/2$, we have $\sigma_{\mathcal{G}} \leq c_0 [\mathcal{G}]_{\mathcal{F}^i_j}$.

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

Iterative logic. The unconstrained arrangement is constructed by first smoothing consumption within each $\mathcal{F}^1_j$ set; then smoothing consumption within each $\mathcal{F}^2_j$ set; and so on. When $i = 1$, all sets are singletons, so there is no need to smooth within a set. Now consider the step when we move from i to $i+1$. As we have seen, by (N4) the number of i level sets in $\mathcal{F}_j^{i+1}$ is bounded by a positive constant K. To simplify notation, denote $\mathcal{F}_j^{i+1}$ by $\mathcal{F}$, and denote the i-level sets $\mathcal{F}_{j'}^{i}$ that are subsets of $\mathcal{F}$ by $\mathcal{F}_1,...,\mathcal{F}_k$ where $k \le K$. We know from (N2) and (N3) that $\mathcal{F}_1,...,\mathcal{F}_k$ partition $\mathcal{F}$. We smooth consumption in $\mathcal{F}_j^{i+1}$ by first smoothing the total amount of resources currently present in $\mathcal{F}_1$ through the entire set $\mathcal{F}$; then smoothing the total amount currently in $\mathcal{F}_2$ through the set $\mathcal{F}$, and so on until $\mathcal{F}_k$. Note that the total consumption in $\mathcal{F}_1$ at this round is the same as the total endowment $e_{\mathcal{F}_1}$, because in each round i, we smooth all endowments within an i-level set. Having completely smoothed resources in $\mathcal{F}_1$ in the previous round, all agents in $\mathcal{F}_1$ are currently allocated $e_{\mathcal{F}_1} / |\mathcal{F}_1|$ units of consumption.

Auxiliary network flow. To smooth consumption over $\mathcal{F}$, we define an auxiliary network flow. This is a key step in the proof. For this flow, focus on the subgraph generated by $\mathcal{F}$ together with capacities c_0, and assume for the moment that each agent in $\mathcal{F}_1$ has $\sigma_{\mathcal{F}_1} / |\mathcal{F}_1|$ units of the consumption good (so the total in $\mathcal{F}_1$ is exactly $\sigma_{\mathcal{F}_1}$), while each agent in $\mathcal{F} \setminus \mathcal{F}_1$ has zero. We will show that a flow respecting capacities c_0 can achieve equal sharing in $\mathcal{F}$ from this endowment profile; and then use this flow to construct an unconstrained flow implementing the desired sharing over $\mathcal{F}$ for arbitrary shock realizations.

To verify that equal sharing can be implemented in the above endowment profile, note that equal sharing can be implemented through some IC transfer if for each set $\mathcal{G} \subseteq \mathcal{F}$ the excess demand for goods does not exceed the perimeter relative to $\mathcal{F}$ (this is where the key perimeter/area condition (K) plays it's role). What is this excess demand? Since we want equal sharing, we should allocate $\sigma_{\mathcal{F}_1} / |\mathcal{F}|$ to every agent in $\mathcal{G}$. But those agents in $\mathcal{G}$ who are also in $\mathcal{F}_1$ each have $\sigma_{\mathcal{F}_1} / |\mathcal{F}_1|$. So the excess demand for goods in the set $\mathcal{G}$ is

$$ed\,(\mathcal{G}) = |\mathcal{G}| \cdot \frac{\sigma_{\mathcal{F}_1}}{|\mathcal{F}|} - |\mathcal{G} \cap \mathcal{F}_1| \cdot \frac{\sigma_{\mathcal{F}_1}}{|\mathcal{F}_1|}. \tag{8}$$

Applying Theorem 1 to the finite network $\mathcal{F}$, there is a feasible flow if for every $\mathcal{G}$, we have $|ed\,(\mathcal{G})| \le c_0\,[\mathcal{G}]_{\mathcal{F}}$. To check that this holds, first assume that $|\mathcal{G}| / |\mathcal{F}| \ge |\mathcal{G} \cap \mathcal{F}_1| / |\mathcal{F}_1|$; then the first term in (8) is larger, and hence $|ed\,(\mathcal{G})| \le \sigma_{\mathcal{F}_1} \cdot |\mathcal{G}| / |\mathcal{F}|$. From (P4) we have $\sigma_{\mathcal{F}_1} \le \sigma_{\mathcal{F}}$, implying $|ed\,(\mathcal{G})| \le \sigma_{\mathcal{F}} \cdot |\mathcal{G}| / |\mathcal{F}|$. Now (P5) implies $\sigma_{\mathcal{F}} / |\mathcal{F}| \le \sigma_{\mathcal{G}} / |\mathcal{G}|$, and hence $|ed\,(\mathcal{G})| \le \sigma_{\mathcal{G}}$. Now recall the key condition (K') that $\sigma_{\mathcal{G}} \le c_0\,[\mathcal{G}]_{\mathcal{F}}$; it follows that $|ed\,(\mathcal{G})| \le c_0\,[\mathcal{G}]_{\mathcal{F}}$ as desired. We now check that (8) also holds when $|\mathcal{G}| / |\mathcal{F}| < |\mathcal{G} \cap \mathcal{F}_1| / |\mathcal{F}_1|$. In this case, the second term in (8) dominates, and hence $|ed\,(\mathcal{G})| \le \sigma_{\mathcal{F}_1} \cdot |\mathcal{G} \cap \mathcal{F}_1| / |\mathcal{F}_1|$. Since $\sigma_{\mathcal{F}_1} / |\mathcal{F}_1| \le \sigma_{\mathcal{G} \cap \mathcal{F}_1} / |\mathcal{G} \cap \mathcal{F}_1|$ by (P3), we can bound the right hand side by $\sigma_{\mathcal{G} \cap \mathcal{F}_1}$, which satisfies $\sigma_{\mathcal{G} \cap \mathcal{F}_1} \le \sigma_{\mathcal{G}} \le c_0\,[\mathcal{G}]_{\mathcal{F}}$ again verifying that $|ed\,(\mathcal{G})| \le c_0\,[\mathcal{G}]_{\mathcal{F}}$. This shows that the auxiliary flow can be implemented.

Smoothing with auxiliary flow. Denote the transfers associated with the auxiliary flow by t_1. To smooth the consumption of $\mathcal{F}_1$ over $\mathcal{F}$ for arbitrary shocks, we just use the transfers $t_1 \cdot \left(e_{\mathcal{F}_1} / \sigma_{\mathcal{F}_1}\right)$; that is, we scale up the above flow with the actual size of the shock in $\mathcal{F}_1$. This works, because t_1 was constructed to smooth a shock of exactly one standard deviation $\sigma_{\mathcal{F}_1}$. Extending this logic, to smooth the endowment of each other $\mathcal{F}_j$

through the set $\mathcal{F}$, we construct auxiliary flows t_2, ..., t_k analogously, and implement the total transfer given by $t_1 \cdot e_{\mathcal{F}_1}/\sigma_{\mathcal{F}_1} + ... + t_k \cdot e_{\mathcal{F}_k}/\sigma_{\mathcal{F}_k}$. This construction results in an unconstrained flow which smooths consumption in the entire set $\mathcal{F}$.

Note that while we used the capacities to construct the flow (this is how we got t_1,..., t_k), the actual flow is a stochastic object that may violate some capacity constraints, both because it is scaled by $e_{\mathcal{F}_1}/\sigma_{\mathcal{F}_1}$ and because it is summed over all j.

Iteration. We do the above step for all $i+1$-level sets $\mathcal{F}_j^{i+1}$; 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 $\mathcal{F}_j^n$, $j = 1,...,\infty$. Denote the unconstrained arrangement obtained in this way by t^U.

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

Link usage. Consider the step where we smooth the consumption of $\mathcal{F}_1$ over the entire set $\mathcal{F}$ using the flow $t_1 \cdot e_{\mathcal{F}_1}/\sigma_{\mathcal{F}_1}$. Fix some (u, v) link; then the use of this link in the flow at this round is $t_1(u, v) \cdot e_{\mathcal{F}_1}/\sigma_{\mathcal{F}_1}$. This is a random variable with mean zero and standard deviation $t_1(u, v)$, since $e_{\mathcal{F}_1}/\sigma_{\mathcal{F}_1}$ has unit standard deviation. Moreover, we know that $t_1(u, v) \leq c_0(u, v)$ because this is how t_1 was constructed (this is why it was important to construct t_1 such that it satisfies the capacity constraints c_0.) It follows from Lemma 1 that the standard deviation of link use at this step is at most $c_0(u, v)$.

Now consider link use as we smooth the consumption of all sets $\mathcal{F}_1$, ..., $\mathcal{F}_k$ over the set $\mathcal{F}$. As we have seen, smoothing for each of these sets implies adding a flow over the (u, v) link that has standard deviation of at most $c_0(u, v)$. Given that $k \leq K$ for some constant, the total standard deviation of the flow over (u, v) in each round of the algorithm is at the most $K \cdot c_0(u, v)$. Adding up these flows over all n rounds shows that the total standard deviation of the unconstrained arrangement over the (u, v) link is at most $nK \cdot c_0(u, v)$.

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

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

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

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

$$3 \int \bar{e}_{\mathcal{F}}^2 \, dP + 3 \int_{t^U(u,v)>c\cdot c_0(u,v)} \left[t(u,v) - c \cdot c_0(u,v)\right]^2 \, dP. \tag{9}$$

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

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

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

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

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

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

Bound on remaining variance. Using the bound on the tail probability, we can estimate the final term in (9). Let $z = t^U(u,v)$ which is a weighted sum of the y_j shocks by construction. Denoting the c.d.f. of z by $H(z)$ we have

$$\begin{aligned}
\int_{t^U(u,v)>c\cdot c_0(u,v)} \left[t(u,v) - c \cdot c_0(u,v)\right]^2 \, dP &= \int_{z=c\cdot c_0(u,v)}^{\infty} (z - c \cdot c_0(u,v))^2 \, dH(z) \\
&= -\int_{z=c\cdot c_0(u,v)}^{\infty} (z - c \cdot c_0(u,v))^2 \, d[1 - H(z)] = \\
= -\left[(z - c \cdot c_0(u,v))^2 (1 - H(z))\right]_{c(u,v)}^{\infty} &+ \int_{z=c\cdot c_0(u,v)}^{\infty} 2(z - c \cdot c_0(u,v))[1 - H(z)] \, dz
\end{aligned}$$

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

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

for appropriate constants K and K'.

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

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

To obtain a bound on $SDISP$, we first bound $DISP = SDISP^2$, which is just the average of Ex_u^2 over the entire network. We have seen that for each u,

$$Ex_u^2 \leq K_1\cdot\exp\left[-K_2 n\right] + K_3\cdot\exp\left[-K_4\cdot(c/n)^2\right]$$

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

Proof of Corollary 1

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

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

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

Now fix $c > 0$, and use a share $1/(10K^*)$ of capacities to implement between-squares risk-sharing using Proposition 2, taking $e_{\mathcal{G}_j}$ as the "endowment shocks" of the squares. The conditions of the proposition are easily seen to be satisfied, and hence we obtain between-squares dispersion which is exponentially small in $c^{2/3}$.

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

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

Proof of Proposition 3

We prove the following more general result.

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

Proof. Let $\mathcal{U}^* \subseteq \mathbb{R}^{|\mathcal{W}|}$ be the set of expected utility profiles that can be achieved by IC transfer arrangements: $\mathcal{U}^* = \{(v_i)_{i \in \mathcal{W}} \mid \exists$ IC allocation $\mathbf{x}$ such that $v_i \leq EU_i(x_i, c_i)$ $\forall i\}$. Our goal is to show that $\mathcal{U}^*$ 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 $\mathbf{e}$ and let $\mathbf{x}$ be an IC allocation. Consider an agent i, and for $r \geq 0$ define $y(r, x_i)$ to be the consumption level that makes i indifferent between his current allocation and reducing friendship consumption by r units, that is, $U(x_i, c_i) = U(y(r, x_i), c_i - r)$. For different values of r, the locations $(y(r, x_i), c - r)$ trace out an indifference curve of i. Note that $y(0, x_i) = x_i$ and that the IC constraint for the transfer between i and j can be written as

$$t_{ij} \leq y(c(i, j), x_i) - x_i \tag{10}$$

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

$$y_r(r, x_i) = \frac{U_c}{U_x}(y, c_i - r) \tag{11}$$

which is the marginal rate of substitution MRS_i. This is intuitive: MRS_i measures the dollar value of a marginal change in friendship consumption. Using the concavity of the MRS, we will show that $y(r, x_i)$ 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 (10), and consequently, that the set of IC profiles is convex.

To show that $y(r, x_i)$ is concave in x_i, let $\mathbf{x}^1$, $\mathbf{x}^2$ 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 $\overline{y}(r) = \alpha y\left(r, x_i^1\right) + (1-\alpha) y\left(r, x_i^2\right)$, so that $\left(\overline{y}(r), c_i - r\right)$ traces out the convex combination of the indifference curves passing through $\left(x_i^1, c_i\right)$ and $\left(x_i^2, c_i\right)$, and let $f(r) = U\left(\overline{y}(r), c_i - r\right)$, the utility of agent i along this curve. Clearly, $f(0) = U(x_3, c_i)$. Moreover, using (11),

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

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 $U\left(\overline{y}(r), c_i - r\right) \leq U\left(x_i^3, c_i\right)$, which implies that $y\left(x_i^3, r\right) \geq \overline{y}(r) = \alpha y\left(r, x_i^1\right) + (1-\alpha) y\left(r, x_i^2\right)$, and hence that $y(x, r)$ is concave.

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

Proof of Proposition 4

Fix realization $\mathbf{e}$, and let t denote the vector of transfers over all links in a given IC arrangement. Denote the planner's objective with a given set of weights λ_i by $V(t) = \sum_i \lambda_i U_i\left(e_i - \sum_j t_{ij}, c_i\right)$. Then the planner's maximization problem can be written as $\max_t V(t)$ subject to $t_{ij} \leq c(i, j)$ and $t_{ij} = -t_{ji}$ for all i and j. It is easy to see that Karush-Kuhn-Tucker first order conditions associated with this problem are those given in the Proposition. Since we have a concave maximization problem where the inequality constraints are linear, the Karush-Kuhn-Tucker conditions are both necessary and sufficient for characterizing a global maximum. For uniqueness, rewrite the planner's objective as a function of the consumption profile $\mathbf{x}$, $\overline{V}(\mathbf{x}) = V(t)$. This function is strictly concave in $\mathbf{x}$ 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 5

For each i and j, say that i and j are in the same equivalence class if there is an $i \to j$ path such that for all agents l on this path, including j, we have $\lambda_i U_i' = \lambda_l U_l'$. The partition generated by these equivalence classes is the set of risk-sharing islands $\mathcal{W}_k$. If $i \in \mathcal{W}_k$ and $j \notin \mathcal{W}_k$, 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 4 implies that $\left|t_{ij}\right| = c(i,j)$, as desired.

Proof of Proposition 6

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

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

Fix a constrained efficient arrangement, and suppose it is not monotone over all sets. Let $\mathcal{F}$ be a set where this property fails, and fix a connected component of the subgraph spanned by $\mathcal{F}$ that contains an agent i such that $x_i' > x_i$. Let $\mathcal{S}$ be the set of agents for whom $x_i' \leq x_i$, and $\mathcal{T}$ be the set of agents for whom $x_i' > x_i$ in this component. $\mathcal{S}$ is non-empty, because the total endowment available in any connected component of $\mathcal{F}$ has decreased, and $\mathcal{T}$ is non-empty by assumption. In addition, there exist $s \in \mathcal{S}$ and $t \in \mathcal{T}$ such that $t_{st}' > t_{st}$, because consumption in $\mathcal{T}$ is higher under $\mathbf{e}'$ than under $\mathbf{e}$. But $t_{st}' > t_{st}$ implies $c(s,t) > t_{st}$ and $c(t,s) > t_{ts}'$, and hence, by Proposition 4, $\lambda_s U_s'(x_s) \geq \lambda_t U_t'(x_t)$ in $\mathbf{e}$, and also $\lambda_s U_s'(x_s') \leq \lambda_t U_t'(x_t')$ in $\mathbf{e}'$. 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 $j \in \widehat{\mathcal{W}}(i)$ follows directly from this monotonicity condition combined with (ii) which is proved below.

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

Let $\mathcal{L}_i'$ denote the set of links connecting agents in $\mathcal{W}(i)$ with agents in $\mathcal{W} \setminus \mathcal{W}(i)$. For every $(k,l) \in \mathcal{L}_i'$, let t_{kl}' be such that $\lambda_k U_k'(x_k(\mathbf{e}) - t_{kl}') = \lambda_l U_l'(x_l(\mathbf{e}) + t_{kl}')$. In words, t_{kl}' is the amount of transfer between k and l that would equate the weighted marginal utilities of k and l. By Proposition 4 and by the definition of $\mathcal{W}(i)$, $t_{kl}' \neq 0\ \forall\ (k,l) \in \mathcal{L}_i'$. Let b' be the minimum amount of transfer that would equate the weighted marginal utilities of an agent in $\mathcal{W}(i)$ and a neighboring agent outside $\mathcal{W}(i)$: $b' = \min_{(k,l)\in\mathcal{L}_i'}\left|t_{kl}'\right|$.

We claim that the result holds for $\Delta = \min(b, b')$, that is whenever $\left|e_i - e_i'\right| <$

$\min(b,b')$, we have $\lambda_j U'_j(x_j(\mathbf{e}')) = \lambda_i U'_i(x_i(\mathbf{e}'))\ \forall\ j \in \widehat{\mathcal{W}}(i)$, and $U_j(x_j(\mathbf{e}')) = U_i(x_i(\mathbf{e}))$ $\forall\ j \notin \mathcal{W}(i)$. To see this, consider the restricted set of agents $\mathcal{W}(i)$, and endowments $x_i(\mathbf{e}) + e'_i - e_i$ for agent i, and $x_j(\mathbf{e})$ for $j \in \mathcal{W}(i)/\{i\}$ (where $x_i(\mathbf{e})$ still refers to the constrained efficient allocation given set of agents $\mathcal{W}$ and endowment realization $\mathbf{e}$). Let $\mathbf{x}^{\mathbf{e},\mathbf{e}'}$ denote this endowment vector on $\mathcal{W}(i)$. Consider now the consumption arrangement over $\mathcal{W}(i)$ that maximizes $\sum_{j\in\mathcal{W}(i)} \lambda_j U_j(x_j)$ subject to $\mathbf{x}$ being achievable from $\mathbf{x}^{\mathbf{e},\mathbf{e}'}$ by transfer scheme t' (over $\mathcal{W}(i)$) for which $|t_{jj'} + t'_{jj'}| \le c(j,j')\ \forall\ j,j' \in \mathcal{W}(i)$. Let this arrangement be denoted by $\mathbf{x}^{\mathcal{W}(i)}(\mathbf{e})$. Because $\lambda_j U'_j(x_j)$ is decreasing in x_j for all j, $|x_i^{\mathcal{W}(i)}(\mathbf{e}) - x_i(\mathbf{e})| \le \left|e_i - e'_i\right|$. Then there is a transfer scheme t' over $\mathcal{W}(i)$ that achieves $\mathbf{x}^{\mathcal{W}(i)}(\mathbf{e})$ from endowments $\mathbf{x}^{\mathbf{e},\mathbf{e}'}$, for which $|t'_{jj'}| \le \left|e_i - e'_i\right| < \Delta$. Since $\Delta < b$, all the capacity constraints in $\widehat{\mathcal{L}}_i$ are still slack. By Proposition 4 this means that $\lambda_j U'_j(x_j^{\mathcal{W}(i)}(\mathbf{e})) = \lambda_i U'_i(x_i^{\mathcal{W}(i)}(\mathbf{e}))$. Moreover, since $\Delta < b'$, all the capacity constraints in $\mathcal{L}'_i$ are still binding, in the same direction. Extend now $\mathbf{x}^{\mathcal{W}(i)}(\mathbf{e})$ to $\mathcal{W}$ such that $x_j^{\mathcal{W}(i)}(\mathbf{e}) = x_j(\mathbf{e})$ for $j \in \mathcal{W} \setminus \mathcal{W}(i)$. Similarly, extend transfer scheme t' to $\mathcal{W}$ such that $t'_{jj'} = 0$ whenever at least one of j an j' are not in $\mathcal{W}(i)$. Note that $t + t'$ is a direct transfer arrangement on $\mathcal{W}$ which meets the capacity constraints, and that $\mathbf{x}^{\mathcal{W}(i)}(\mathbf{e})$ satisfies the conditions of Proposition 4. Hence $\mathbf{x}^{\mathcal{W}(i)}(\mathbf{e})$ is the constrained efficient allocation given endowment realization $\mathbf{e}'$, and as shown above, satisfies the claims in (ii).

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

A-2. Microfoundations for link-level punishment

Consider the following multi-stage game.

Stage 1. An endowment vector $\mathbf{e}$ is drawn from a commonly known prior distribution.

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

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

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

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

Stage 4. The realized utility of agent i is $U_i\left(x_i', c_i'\right)$.

PROPOSITION 9: *An allocation* $\mathbf{x}(\mathbf{e})$ *is the outcome of a pure-strategy subgame-perfect equilibrium of this game if and only if it can implemented through an incentive-compatible informal risk-sharing arrangement.*

PROOF:

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

A-3. Background on the theory of network flows

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

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

(i) Skew symmetry: $f(u, v) = -f(v, u)$.

(ii) Capacity constraints: $f(u, v) \leq c(u, v)$.

(iii) Flow conservation: $\sum_w f(u, w) = 0$ *unless* $u = s$ *or* $u = t$.

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

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

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

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

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

A-4. Discussion of Dynamic Mechanisms Generating Constrained Efficiency

We now briefly discuss two intuitive dynamic mechanisms that provide foundations for constrained efficiency.

A decentralized exchange implementing any constrained efficient arrangement. We first consider a decentralized itarative procedure in which agents use a simple rule of thumb in helping those who are in need. In particular, we show that for any constrained efficient allocation, there exists a simple iterative procedure that uses, in each round of the iteration, only local information about the current resources of the parties involves, 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 (2007). 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 $\mathbf{e}$, and denote the efficient allocation corresponding to welfare weights λ_i by $\mathbf{x}^*$. Fix an order of all links in the network: l_1,...,l_L, and let i_k and j_k denote the agents connected by l_k. To initialize the procedure, set $x_i = e_i$ and $t_{ij} = 0$ for all i and j. Then, in every round $m = 1, 2, ...$, go through the links $l_1, ..., l_L$ in this order, and for every l_k, given the current values x_{i_k}, x_{j_k}, and $t_{i_k j_k}$, define the new values x'_{i_k} and x'_{j_k} and $t'_{i_k j_k} = t_{i_k j_k} + x'_{j_k} - x_{j_k}$ such that they satisfy the following two properties: (1) $x'_{i_k} + x'_{j_k} = x_{i_k} + x_{j_k}$. (2) Either $\lambda_{i_k} U'_{i_k}(x'_{i_k}) = \lambda_{j_k} U'(x'_{j_k})$, or $\lambda_{i_k} U'_{i_k}(x'_{i_k}) > \lambda_{j_k} U'_{j_k}(x'_{j_k})$ and $t'_{i_k j_k} = -c(i, j)$, or $\lambda_{i_k} U'_{i_k}(x'_{i_k}) < \lambda_{j_k} U'_{j_k}(x'_{j_k})$ and $t'_{i_k j_k} = c(i, j)$. This amounts to the agent with lower marginal utility helping out his friend up to the point where either their marginal utility is equalized, or the capacity constraint starts to bind. Once this step is completed for link k, we set $\mathbf{x} = \mathbf{x}'$ and $t = t'$ before moving on to link $k + 1$. For $m = 1, 2, ...$ let x_i^m denote the value of x_i, and let t_{ij}^m denote the value of t_{ij}, at the end of round m. Note that $\mathbf{x}_m$ meets the capacity constraints by design for every m.

PROPOSITION 10: *If consumption and friendship are perfect substitutes, then $\mathbf{x}^m \to \mathbf{x}^*$ as $m \to \infty$.*

PROOF:

Let $V(\mathbf{x})$ denote the value of the planner's objective in allocation $\mathbf{x}$. The above procedure weakly increases $V(\mathbf{x})$ in every round and for every link l_k. Hence $V(\mathbf{x}_1) \leq V(\mathbf{x}_2) \leq ..$, and since $V(\mathbf{x}) \leq V(\mathbf{x}^*)$ for all $\mathbf{x}$ that are IC, we have $\lim_{m\to\infty} V(\mathbf{x}_m) = V \leq V(\mathbf{x}^*)$. Since the set of IC allocations is compact, and $\mathbf{x}_m$ is IC for every m, there exists a convergent subsequence of $\mathbf{x}_m$, with limit $\mathbf{x}$ and associated transfers t. Clearly, $V(\mathbf{x}) = V$. If $V = V^*$ then $\mathbf{x} = \mathbf{x}^*$ since the optimum is unique. If $V < V^*$, then $\mathbf{x}$ 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 $\mathbf{x}$ and t. Then there is a transfer meeting the capacity constraints at $\mathbf{x}$ that increases the planner's objective by a strictly positive amount δ. But this means that for every $\mathbf{x}_m$ far along the convergent subsequence, the planner's objective increases by at least $\delta/2$ at that round, which implies that $V(\mathbf{x}_m)$ is divergent, a contradiction. Hence $\lim \mathbf{x}_m = \mathbf{x}^*$ along all convergent subsequences, which implies that $\mathbf{x}_m$ itself converges to $\mathbf{x}^*$.

Ex ante coalition-proofness of constrained efficiency. A second mechanism which yields constrained-efficient allocations is collective dynamic bargaining with renegotiation. Gomes (2000) shows that when agents can propose renegotiable arrangements to subgroups and make side-payments in a dynamic bargaining procedure, ultimately a Pareto-efficient arrangement will be selected.[2] We now show how to incorporate this result in our model by assuming that there is a negotiations phase prior to the endowment realization.

We say that a coalition-proof agreement $\mathbf{x}$ admits no ex ante coalitional deviations if there is no coalition $\mathcal{S}$ and coalition-proof risk-sharing agreement $\mathbf{x}'_{\mathcal{S}}$ within $\mathcal{S}$ such that all agents in $\mathcal{S}$ weakly prefer losing all their links to agents in $W \setminus \mathcal{S}$ and having agreement $\mathbf{x}'_{\mathcal{S}}$ to keeping all their links and having agreement $\mathbf{x}$, and at least one agent in $\mathcal{S}$ strictly prefer the former. Intuitively, an ex ante coalitional deviation implies that the agents of the deviating coalition leave the community (cut their ties with the rest of the community) and agree upon a new risk-sharing agreement among each other (using only their own resources).

PROPOSITION 11: *A coalition-proof agreement that admits no profitable ex ante coalitional deviations is constrained efficient. If goods and friendship are perfect substitutes then the set of coalition-proof agreements that admit no profitable ex ante deviations is equal to the set of constrained efficient agreements.*

PROOF:

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

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

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

A-5. Analysis with imperfect substitutes

We now explain how our results extend when goods and friendship are imperfect substitutes. With a general utility function $U(x, c)$, the definition of incentive compatibility (IC) of a transfer arrangement is the following:

DEFINITION 5: *A risk-sharing arrangement t is incentive compatible (IC for short) if*

$$U_i(x_i, c_i) \geq U_i\left(x_i + t_{ij}, c_i - c(i, j)\right) \tag{12}$$

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

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

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

A. The limits to risk-sharing with imperfect substitutes

With imperfect substitutes, the results in section II extend but the upper and lower bounds on risk-sharing are weakened by constant factors that depend on the degree of substitution. To obtain these extensions, we assume that the marginal rate of substitution

(MRS) is uniformly bounded. We continue to find that the first-best can only be achieved in highly expansive graphs where the perimeter-area ratio is bounded from below: we require $a[\mathcal{F}] \geq \underline{\sigma}/M$. Our findings about partial risk-sharing are about rates of convergence and hence they extend without modification; in particular, $SDISP$ converges exponentially for geographic networks.

Imperfect substitution also yields additional implications. If the MRS is increasing in consumption, then agents with low consumption value their friends less, reducing the maximum amount they are willing to give up. As a result, in a society that experiences a negative aggregate shock, the scope for insuring idiosyncratic risk is reduced. To formalize this point, we now also show that with an increasing MRS, the set of IC arrangements contracts after a negative aggregate shock.

PROPOSITION 12: *Assume that MRS_i is increasing in x_i for all i. Then for any pair of endowment realizations $\underline{\mathbf{e}}$ and $\overline{\mathbf{e}}$ such that $\underline{e}_i \leq \overline{e}_i$ for all i, an incentive compatible set of transfers in $\underline{\mathbf{e}}$ is also incentive compatible given $\overline{\mathbf{e}}$.*

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

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

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

The aggregate negative shock is thus a double burden: besides its direct negative effect on consumption, it also induces worse sharing of idiosyncratic risks, a finding consistent with Kazianga and Udry (2006), who document limited informal insurance during the severe draught of 1981-85 in rural Burkina Faso.

B. Constrained efficient arrangements

We begin with a summary of our results. The key novelty with imperfect substitutes is that changing the goods consumption of an agent affects his implied link values and hence incentive compatibility. To characterize constrained efficiency, we assume that the marginal rate of substitution MRS_i defined above is concave in x_i. When this holds, we can generalize Proposition 3, establishing the equivalence between constrained efficiency and the planner's problem.

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

Using the concept of marginal social gain allows us to extend the characterization of constrained efficient agreements in Proposition 4. Given this result, we can also partition the network into endogenous risk-sharing islands, such that marginal social utility is equalized within islands, and all links connecting the island to the rest of the community are blocked.

Finally, for an agent i who is not on the boundary of his risk-sharing island and hence has no links with binding constraints, the marginal social gain does equal λ_i times his marginal utility of consumption; hence, for such agents, the results of section III hold without modification. For example, weighted marginal utilities are equalized for any two such agents in the same risk-sharing island. Thus if risk-sharing islands are "large", then the results from the perfect substitutes case hold without modification for most agents.

C. Formal results

The equivalence between the planner's problem and constrained efficiency with general preferences and a concave MRS was established in Appendix A to the paper. To present our characterization result building on this equivalence, first we define a measure of marginal social welfare gain of transfers to agents. Fix an IC arrangement $\mathbf{x}$, and recalling the definition of acyclical transfer arrangements from the proof of Proposition 6, let t be an acyclical implementation of $\mathbf{x}$ in endowment realization e. Consider the following iterative construction. We say that the IC constraint from i to j binds if $U_i(x_i, c_i) = U_i(x_i + t_{ij}, \widehat{c}_{i,j})$. Let $\mathcal{W}^1 \subseteq \mathcal{W}$ denote the set of agents i for whom (i) there is no j such that $c(i, j) > 0$; and (ii) the IC constraint from i to j binds. Since t is acyclical, $\mathcal{W}^1$ is nonempty. For any $i \in \mathcal{W}^1$, let $\Delta_i = \lambda_i U_{i,x}(x_i, c_i)$ be the marginal benefit of an additional dollar to i. This is both the private and social marginal welfare gain, because no IC constraint binds for transfers from i.

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

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

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

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

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 13: *[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 $(\lambda_i)_{i \in \mathcal{W}}$ such that for every $i, j \in \mathcal{W}$ one of the following conditions holds:*

1) $\Delta_j = \Delta_i$
2) $\Delta_j > \Delta_i$ and the IC constraint binds for t_{ij}
3) $\Delta_j < \Delta_i$ and the IC constraint binds for t_{ji}.

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

Now we establish the necessity part of the proposition. Fix a constrained efficient arrangement, and let λ_i be the associated planner weights. Consider realization $\mathbf{e}$. We first show that the marginal value to the planner of an additional dollar to an agent i is Δ_i. Let $i \in \mathcal{W}^1$, 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 $j \in \cup_{l \le k-1} \mathcal{W}^l$ that the marginal social value of a dollar to j is Δ_j. Let $i \in \mathcal{W}^k$. For any j such that the IC constraint from i to j is binding, Δ_j is at least as large as the marginal social value of an additional dollar to i, because otherwise optimality requires reducing t_{ij}. Hence the marginal social value of a dollar to i is obtained when i transfers as much of the dollar as possible under incentive compatibility to some agent j. Given our above argument, i can transfer at most $1 - U_{i,x}(x_i, c_i)/U_{i,x}(\widehat{x}_{i,j}, \widehat{c}_{i,j})$ to j, hence the marginal welfare gain if he chooses to transfer to j will be δ_{ij}. Since i will choose to transfer the dollar to the agent where it is most productive, the marginal social gain will be the maximum of δ_{ij} over j, which is Δ_i.

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

For sufficiency, let now $\mathbf{x}$ denote the unique welfare maximizing consumption, let t be an IC transfer scheme achieving this allocation, and let $\widehat{\Delta}_i = \Delta_i(\mathbf{x}, t)$, for every $i \in \mathcal{W}$. Assume now that there exists another consumption vector $\mathbf{x}' \neq \mathbf{x}$ achieved by IC transfer scheme t' such that $(\mathbf{x}', t')$ satisfy conditions (1)-(3), and let $\Delta_i' = \Delta_i(\mathbf{x}', t')$, for every $i \in N$. Then there exists an acyclical nonzero transfer scheme t^d that achieves $\mathbf{x}$ from $\mathbf{x}'$, and which is such that $t' + t^d$ is IC. By definition of $\mathbf{x}$, t^d from $\mathbf{x}'$ improves social welfare. Let now $\mathcal{W}^d = \{i \in \mathcal{W} \mid \exists\, j \text{ such that } t_{ij}^d \neq 0\}$, and partition $\mathcal{W}^d$ into sets $\mathcal{W}_0^d, ..., \mathcal{W}_K^d$ the following way. Let $\mathcal{W}_0^d = \{i \in \mathcal{W}^d \mid -\exists\, j \in \mathcal{W}^d \text{ st. } t_{ij}^d > 0\}$. Given $\mathcal{W}_0^d, ..., \mathcal{W}_k^d$ for some $k \geq 0$, let $\mathcal{W}_{k+1}^d = \{i \in \mathcal{W}^d \setminus (\underset{l=0,...,k}{\cup} \mathcal{W}_l^d) \mid -\exists\, j \in \mathcal{W}^d \setminus (\underset{l=0,...,k}{\cup} \mathcal{W}_l^d) \text{ st. } t_{ij}^d > 0\}$. Note that $x_i' > x_i\ \forall\, i \in \mathcal{W}_0^d$, which together with there being no agent j such that $t_{i\,j}^d > 0$ implies that $\Delta_i' < \widehat{\Delta}_i$. Now we iteratively establish that $\Delta_i' < \widehat{\Delta}_i\ \forall\, i \in \mathcal{W}^d$. Suppose that $\Delta_i' < \widehat{\Delta}_i\ \forall\, i \in \underset{l=0,...,k}{\cup} \mathcal{W}_l^d$ for some $k \geq 0$. Let $i \in \mathcal{W}_{k+1}^d$. Note that by definition there is $j \in \underset{l=0,...,k}{\cup} \mathcal{W}_l^d$ such that $t_{i\,j}^d > 0$, and there is no $j' \in \mathcal{W}^d \setminus (\underset{l=0,...,k}{\cup} \mathcal{W}_l^d)$ such that $t_{i\,j'}^d > 0$. Suppose $\Delta_i' \geq \widehat{\Delta}_i$. This can only be compatible with $t_{i\,j}^d > 0$, $\Delta_j' < \widehat{\Delta}_j$, and (1)-(3) holding for both $(\mathbf{x}', t')$ and $(\mathbf{x}, t' + t^d)$ if $x_i > x_i'$. But $x_i > x_i'$, and $\Delta_{i'}' < \widehat{\Delta}_{i'}\ \forall\, i' \in \mathcal{W}$ such that $t_{ii'}^d > 0$ implies $\Delta_i' < \widehat{\Delta}_i$, a contradiction. Hence $\Delta_i' < \widehat{\Delta}_i\ \forall\, i \in \mathcal{W}_{k+1}^d$, and then by induction $\Delta_i' < \widehat{\Delta}_i\ \forall\, i \in \mathcal{W}^d$. But note that for any $i \in \mathcal{W}_K^d$ it holds that $x_i < x_i'$ and there is no $j \in \mathcal{W}$ such that $t_{ji}^d > 0$, and hence $\Delta_i' > \widehat{\Delta}_i$. This contradicts $\Delta_i' < \widehat{\Delta}_i\ \forall\, i \in \mathcal{W}^d$, hence there cannot be $(\mathbf{x}', t')$ satisfying (1)-(3) such that t' is IC and $\mathbf{x}' \neq \mathbf{x}$.

Proposition 6 can also be extended to the imperfect substitutes case. Fix a constrained efficient arrangement, and let $\mathbf{e}$ and $\mathbf{e}'$ be two endowment realizations such that $\mathbf{e}_i > e_i'$ for some $i \in \mathcal{W}$, and $\mathbf{e}_j = e_j'\ \forall\, j \in \mathcal{W} \setminus \{i\}$. Let $\mathbf{x}^*(\mathbf{x})$ be the consumption in the constrained efficient allocation after $\mathbf{e}$. Analogously to the perfect substitutes case, let $\widehat{\mathcal{W}}(i)$ the largest set of connected agents containing i such that all IC constraints within the set are slack given some transfer arrangement achieving the constrained efficient allocation after $\mathbf{e}_i$. For any endowment realization $\mathbf{e}$, let $\Delta_j(\mathbf{e})$ be Δ_j, as defined above, given any transfer scheme with the maximal number of links on which the IC constraints are slack, among the ones that attain the constrained efficient allocation. It is straightforward to show that there is a transfer scheme with a maximal number of links on which the IC constraints are slack, among the ones achieving the constrained efficient allocation, and that for all such transfer arrangements Δ_j is the same.

COROLLARY 1: *[Spillovers with imperfect substitutes] Assume that MRS_i is concave, then*

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

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

(iii) [More sharing with close friends] For any $j \neq i$, there exists a path $i \to j$ such that for any agent l along the path, $\Delta_l(\mathbf{e}') \geq \Delta_j(\mathbf{e}')$.

The proof of this result is analogous to the perfect substitutes case and hence omitted. Note that (ii) is weaker than in Proposition 6, because even small shocks can spill over the boundaries of the risk-sharing islands of agent hit by the shocks. Also note that since $\Delta_i = \lambda_i U_{i,x}$ for any agent not on the boundary of an island, (i) implies that consumption is monotonic in the endowment realization for such agents.

A-6. Numerical methods

Risk-sharing simulations. We use the following numerical approach for the simulations underlying Figure 5. We assume throughout that endowment shocks are uniformly distributed with support $[-1, 1]$. We build on Theorem 1 and express a SDISP-minimizing incentive-compatible risk-sharing arrangement as a cost-minimizing flow as follows. (1) Create two artificial nodes s and t as in the proof of Theorem 1. (2) Divide the shock support into K equal intervals. For each agent i, denote the subinterval into which i's endowment falls by k_i (treating $[-1, -1 + 2/K]$ as the first interval and $[1 - 2/K, 1]$ as the Kth interval). Create k_i links between s and i such that each link has capacity $2/K$ in the direction from s to i and zero in reverse direction. Define the "cost" of a flow going from s to i across any of these links to be j for the jth link of out k_i links. Similarly, create $K - k_i$ links between t and i. such that each link has capacity $2/K$ in the direction from i to t and zero in reverse direction. Define the cost of a flow going from i to t across any of these links to be j for the jth link of out k_i links. (3) Use Edmonds and Karp's (1972) algorithm to calculate a cost-minimizing flow in this augmented network. This solution induces an incentive-compatible risk-sharing arrangement that maximizes a piecewise linear approximation to the quadratic utility function assumed in the definition of $SDISP$, where the marginal utility of consumption for any agent is constant within each of the K intervals. Simulations (not reported) show that this approximation generates highly accurate predictions for $K = 20$. For the results presented in the text we set $K = 100$.

Geographic network representation. The algorithm used in the geographic representation constructed in Figure 6 is the following. For each household i, we first construct vectors $\mathbf{v}_j$ to every other households j in the unit square using households' initial (rescaled) geographic coordinates. We also calculate the length d_i of each of these vectors. Note, that the maximum distance between two households is $\sqrt{2}$. We then calculate a shift vector as the weighted sum $-\sum(\sqrt{2} - d_i)\mathbf{v}_j / \left\|\mathbf{v}_j\right\|$ and move each household in the direction of this shift vector. Shifts are larger if a household is closely surrounded by other households and the shift will push the household away from its neighbors. This procedure is repeated 23 times to obtain the representation in Figure 6E.

Geographic network representation of a circle. We apply the diffusion algorithm to a clearly non-geographic network to illustrate the validity of our approach. We use a circle with the same number of nodes and equivalent degree as the Huaraz network on which we based Figure 6. To equalize the degree distribution we assume a circle network where every agent interacts with r neighbors on each side such that $2r$ equals the average Huaraz degree (we randomize between r and $r + 1$ to overcome integer constraints). The diffusion algorithm imposes some randomness depending on the order of shocks that are applied to nodes: the standard deviations for neighboring square connections in the Huaraz and circle geographic representation are 2.3 and 2.1 respectively - hence the number of neighboring square connections is significantly different.

The diffusion algorithm provides the geographic representation shown in Figure 10 that has far more gaps (especially in the center): the average number of neighboring square connections is now only 23.0 which is less than half the number of neighboring connections in Figure 6E.

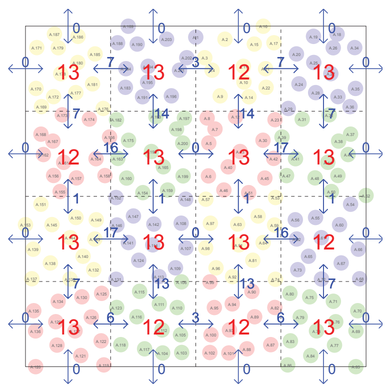

FIGURE 10. STRETCHING A HUARAZ-LIKE CIRCLE (WITH SAME NUMBER OF NODES AND EQUIVALENT AVERAGE DEGREE) TO CONSTRUCT A GEOGRAPHIC REPRESENTATION

REFERENCES

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

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

Bramoulle, Yann, and Rachel E. Kranton. 2007. "Risk-Sharing Networks." *Journal of Economic Behavior and Organization*, 64(3): 275 – 294.

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

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

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

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

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