Supplementary Appendix: For Online Publication Only

In this supplementary appendix we provide an extended analysis that supports our main paper, "Investments in social ties, risk sharing, and inequality," henceforth referred to as the main paper.

A. More general environments and surplus sharing rules

This section provides a slightly more general and comprehensive treatment of the issues studied in the corresponding section of the main paper—Section 6.1. The main difference is that here we allow for multiple groups, while in Section 6.1 attention was restricted to the one-group case. Because of the generalization, we present a complete and self-standing analysis, even though there is much overlap with Section 6.1. We number replicated assumptions and results so that they correspond to those in Section 6.1, and new results with the prefix SA.

The purpose of this section is to examine under what conditions our main conclusions extend to more general utility functions, income distributions, and surplus division rules. The environment with CARA utilities and jointly normally distributed incomes facilitates a convenient transferable (expected) utilities environment that is particularly tractable to analyze when social surplus is divided in accordance with the Myerson value. While analytical tractability requires a series of strong assumptions, below we show that some of the main qualitative insights of the model extend to much more general specifications.

For general specifications of the model, expected utilities are non-transferable and the simple, costless means of redistributing surplus via state-independent transfers we used before is no longer available, hence we need to take a more general approach to risk sharing. Let $v_i(c_i)$ be the utility function for agent i, mapping second-period consumption into utility. We assume that $v_i = v_j$ for all i and j in the same group, and that v_i is strictly increasing and strictly concave for all $i \in \mathbf{N}$.[1] Let $\mathcal{P}_k$ be the distribution the incomes of agents in group $k \in \mathbf{M}$ are drawn from.

Let $\mathcal{L}$ be the set of all possible networks for agents in $\mathbf{N}$. We assume there is a unique risk-sharing arrangement that will be implemented for any possible network $L \in \mathcal{L}$, and that agents correctly anticipate the risk-sharing agreement that will obtain. These risk-sharing arrangements, which depend on the social network, might be dictated by social conventions, or they can be outcomes of negotiation processes for transfer arrangements once the network is formed. Let $\tau(L)$ be the transfer arrangement, and let $u_i^\tau(L)$ be the expected second-period consumption utility of agent i implied by $\tau(L)$.[2]

We continue to assume that for every $L \in \mathcal{L}$, $\tau(L)$ specifies a pairwise-efficient risk-sharing arrangement $\tau_{ij}(L)$ for every pair of agents i, j that are linked in L. As shown earlier, this

[1] These properties imply that for any number of agents more than one, and for any point of the Pareto frontier of feasible consumption plans that can be reached via risk-sharing arrangements, there is a direction along the Pareto frontier in which a given agent's expected utility is strictly increasing.

[2] More precisely, the utility function v_i, the distribution of income realizations, and the transfer arrangement $\tau(L)$ jointly determine $u_i^\tau(L)$.

is equivalent to $\tau(L)$ being Pareto efficient at the component level. Agent i maximizes the difference between expected utility from the second-period risk sharing (given by u_i^τ) and her costs of establishing links.

Let $C_i(L)$ be the set of agents on the same component as i given L, and recall that G is a function that maps agents in $\mathbf{N}$ to groups in $\mathbf{M}$.

Next, we impose a series of assumptions on $\tau(\cdot)$. We do not claim that the above assumptions hold universally when informal risk-sharing takes place, but they are relatively weak requirements that are natural in many settings. Our main objective is to demonstrate that our qualitative results hold for a much broader class of models than the CARA-normal setting with surplus division governed by the Myerson value.

The first assumption requires that establishing a link always strictly increases the connecting agents' expected consumption utilities.

Assumption 11(a). *For every $i \in N$, $u_i^\tau(L \cup \{l_{ij}\}) > u_i^\tau(L)$ for every $L \in \mathcal{L}$ and all $j \in \mathbf{N}$ such that $l_{ij} \notin L$.*

The next assumption requires that establishing an essential link does not impose a negative externality on other agents. This implies that while both i and j privately benefit from essential link l_{ij} in terms of second-period expected utility, they do not benefit over and above the enhancement of risk-sharing opportunities that the link facilitates.

Assumption 11(b). *For every $k \in N$, $u_k^\tau(L \cup \{l_{ij}\}) \geq u_k^\tau(L)$ for every $L \in \mathcal{L}$ and all $i, j \in \mathbf{N}$ such that $C_i(L) \neq C_j(L)$.*

Next, we extend the idea that the private benefit that two agents receive from establishing a link should be increasing in the distance between them in the absence of the link. In the previous analysis these private benefits depended specifically on the Myerson distance between the two agents, while here we allow for a general class of distance measures. Before defining the class of distance measures we allow for, some additional notation is required. For two sets S and S', we define $\mathcal{M}(S, S')$ as the set of matching functions $\mu : S \to S' \cup \{\emptyset\}$ such that for $s \in S$, if $\mu(s) \neq \emptyset$ then $\mu(s) \neq \mu(t)$ for all $t \in S \setminus \{s\}$. Thus every $\mu \in \mathcal{M}(S, S')$ maps each element of S either to a different element of S' or to the empty set. Also, let $\overline{\mathbf{N}}^2 = \{(i,j) | i, j \in \mathbf{N}, i \neq j\}$.

A *distance measure* is a mapping $d : \overline{\mathbf{N}}^2 \times \mathcal{L} \to \mathbb{R}_{++}$ that has the following properties:

Assumption 11(c) (i)–(iii).

(i) *If i and j are in different components on L, then $d(i, j, L) = \overline{d}$, with $\overline{d}$ strictly greater than the maximum possible distance between any two path-connected agents.*

(ii) *d depends only on paths (thus ignoring walks with cycles).*

(iii) *Let S_{ij} be the set of paths between i and j, and let S_{kl} be the set of paths between k and l. We assume that $d(i,j;L) > d(k,l;L)$ if there exists a matching function $\mu \in \mathcal{M}(S,S')$ such that each path between i and j is matched to a shorter path between k and l, and that all such paths between k and l are independent (do not pass through any of the same nodes as each other).*

Assumption 11(c) (i)–(iii) places only weak restrictions on the distance measure. In particular, part (iii) in general provides only a very weak partial ordering of the distances between agents. However, there is a special case in which the ordering is complete. On a tree network, there is a unique path between any two agents, so this determines the ordering of distances between pairs of agents. In what follows, let $d(\cdot)$ be any distance measure satisfying the above requirements.

While we will use the concept of distance between agents in the general case of multiple groups, we first focus on extending our earlier results for the case of homogeneous agents. Then we make assumptions on how distance in the absence of a link influences the private benefits of two agents within the same group establishing that link.

The next assumption requires that if all agents are from the same group, then the private benefit two agents receive when establishing a link depends only (positively) on their distance in the absence of the link and on the sizes of the components they are on. Recall that in our benchmark model in the CARA-normal setting these private benefits depended only on the Myerson distance between the agents. The requirement below allows the private benefit to depend on different distance measures and also on the sizes of the agents' components (which for general utilities influences the difference between the Pareto frontiers of feasible consumption plans with and without the link).

Assumption 11(c). *If $G(i) = G(j)$ for all $i, j \in \mathbf{N}$ such that $l_{ij} \notin L$, then*

$$u_i^\tau(L \cup \{l_{ij}\}) - u_i^\tau(L) = g(d(i,j,L), |C_i(L)|, |C_i(L \cup \{l_{ij}\})|),$$

Moreover, $g(d(i,j,L), |C_i(L)|, |C_i(L \cup \{l_{ij}\})|)$ is increasing in $d(i,j,L)$.

Note that Assumption 11(c) differs slightly from the corresponding assumption in the main paper: Now that we are permitting multiple groups, a qualification is made for this assumption to apply only when all agents are in the same group. In the multiple-group case the composition of each component, in terms of the groups the constituent agents come from and their network positions, can matter.

The last assumption we need for recreating the results of the benchmark model for homogeneous agents is that the cost of link formation within a group is sufficiently small relative to the private benefits from establishing an essential link. In the CARA-normal framework with the surplus allocated according to the Myerson value and all agents being homogeneous, a pair forming an essential link received the full social surplus created by the link. This implies that the social and private benefits coincide in the benchmark model for essential links, and therefore there is no within-group underinvestment for any cost of link formation. For general utility functions and surplus allocation rules, such equivalence does not hold; therefore, a lack of within-group underinvestment cannot be expected to hold for all possible costs of link formation. However, for any specification of the general model that satisfies the assumptions above (in particular that the private benefit of establishing any link is always strictly positive), there is no within-group underinvestment *if* the cost of establishing a link between agents from the same group is small enough.

Assumption 11(d). *For all networks L,*

$$\kappa_w/2 < \min_{i,j:C_i(L)\neq C_j(L)} u_i^{\tau}(L \cup \{l_{ij}\}) - u_i^{\tau}(L).$$

Assumption 11(d) immediately implies that if all agents are from the same group, then in all stable networks there is a single component. The next proposition shows that the same holds for all efficient networks. For the rest of this section, the above assumptions are maintained.

A network is Pareto efficient if there is a feasible transfer agreement that could be reached on that network such that there is no other pair consisting of a network and a feasible transfer agreement in which all agents are weakly better off and some agents are strictly better off.

Proposition 12. *If all agents are from the same group, then a network is Pareto efficient if and only if it is a tree that connects all agents.*

Proof. First, we consider the "only if" direction. In any Pareto-efficient network, every component has to be a tree. This is because if any component is not a tree, then a link could be deleted and the same risk-sharing arrangement as before could be achieved, but the costs of establishing the link would be saved. Now suppose there are two components of a Pareto-efficient network L that are not connected. Let agents i and j be on different components. By Assumption 11(d), the total expected utilities (that is, taking into account the costs of network formation too) of both i and j are strictly higher for network $L \cup \{l_{ij}\}$ than for network L, while by Assumption 11(b) all other agents' total expected utilities are weakly higher for $L \cup \{l_{ij}\}$ than for L. This contradicts the Pareto efficiency of L.

We now consider the "if" direction. Consider a tree network, and suppose we implement a risk-sharing agreement in which $c_i(\omega) = c_j(\omega)$ for all i and j and all states ω. As all agents' consumptions are equalized in all states, there is then no way in which link-formation costs can be redistributed, and the risk-sharing arrangement can be changed, without making someone worse off. Suppose, by way of contradiction, that we can redistribute the link-formation costs by forming a different tree network and find new feasible consumptions that together constitute a Pareto improvement. Holding consumptions fixed, the change in the network will make some agents worse off if any agents are made better off. Thus to achieve a Pareto improvement, consumptions will have to be changed. Let $c'(\omega)$ be the new consumption vector. As all agents in the same group have the same utility function $v_i(c_i) = v(c_i)$ and as the utility function $v(\cdot)$ is strictly concave, Jensen's inequality implies that

$$\frac{1}{n}\sum_i v(c_i'(\omega)) < v\left(\frac{1}{n}\sum_i c_i'(\omega)\right) = \frac{1}{n}\sum_i v(c_i(\omega))$$

for all ω. Thus the average expected utility from consumption will decrease. As total link-formation costs have remained constant, this implies that at least one agent must be worse off. This is a contradiction. □

Corollary SA1. *When all agents are from the same group, there is no underinvestment.*

Given Proposition 12, Corollary SA1 follows immediately from Assumption 11(d) and we omit a proof.

Note that for any non-essential link, $|C_i(L)| = |C_i(L \cup \{l_{ij}\})|$. Thus the marginal benefits i and j receive from forming a superfluous link depends only on the distance between i and j on L and the number of agents in their component. The latter is n for any efficient network, by Proposition 12. Thus the marginal benefit i and j receive from forming a superfluous link depends only on the distance between i and j and is increasing in this distance. Therefore an efficient network will be stable if and only if the maximum distance between any two agents is sufficiently small. The next corollary formally states this result.

Corollary 14. *If all agents are from the same group, then an efficient network is stable if and only if its diameter is sufficiently small.*

Proof. Consider an efficient network L. As L is efficient, there exists a unique path between i and j for all i and all $j \neq i$. Consider two such agents i and $j \neq i$. Assumption 11(c) (i)–(iii) implies that $d(i, j, L)$ is strictly increasing in the path length between i and j and that $d(i, j, L) = d(j, i, L)$. Further, as $|C_i(L)| = |C_i(L \cup \{l_{ij}\})| = n$, by Assumption 11(c) we have that

$$u_i^\tau(L \cup \{l_{ij}\}) - u_i^\tau(L) = g(d(i, j, L), n, n) = g(d(j, i, L), n, n) = u_j^\tau(L \cup \{l_{ij}\}) - u_j^\tau(L).$$

Moreover, by Assumption 11(c), $g(d(i, j, L), n, n)$ is strictly increasing in $d(i, j, L)$. Thus for all i and $j \neq i$, there exists a threshold $\hat{d}$ such that i and j benefit from forming a superfluous link if and only if $d(i, j, L) > \hat{d}$.

In the absence of any underinvestment (by Corollary SA1), a network L is stable if and only if no two agents can benefit from forming a superfluous link. As the agents furthest away from each other have the strongest incentives to form a superfluous link, L is stable if and only if $\max_{i,j} d(i,j,L) \leq \hat{d}$. As $d(i,j,L)$ is strictly increasing in the (unique) path length between i and j, this is equivalent to the diameter of L being sufficiently small. □

A network is *least* stable within a class of networks if its stability implies the stability of any other network in that class. A network is *most* stable within a class of networks if its instability implies the instability of any other network in that class.

Proposition 14. *If all agents are from the same group, then*

(i) *the most stable efficient network is the star,*
(ii) *the least stable efficient network is the line.*

Proof. By Corollary 14, an efficient network is stable if and only if its diameter is sufficiently small. It follows that if a network with diameter d is stable, all efficient networks with weakly lower diameter will also be stable. As the line network is the efficient network that maximizes the diameter, its stability implies the stability of all other efficient networks and it is the least stable efficient network. Similarly, if a network with diameter d is unstable, Corollary 14 implies that all networks with a weakly higher diameter are unstable. As the star network is the efficient network that minimizes the diameter, its instability implies the instability of all other efficient networks and it is the most stable efficient network. □

Inequality measures within the Atkinson class will often rank utility vectors differently. In the simpler setting with CARA utilities, normally distributed incomes, and the Myerson value allocation rule, we were able to identify the star as the least equitable network for any inequality measure in the Atkinson class. This was achieved by showing that any efficient network could be transformed into a star by rewiring it in such a way that, at each step of the rewiring the utility of the center agent increased, the utility of one other agent decreased, and the utility of the remaining agents remained constant. Specifically, the act of removing a link l_{ij} and adding a link l_{jk} increased the utility of agent k, decreased the utility of agent i and held constant the utilities of all other agents.

In the more general setting, this rewiring need not hold constant the utilities of the other agents. This creates problems. Consider the four-agent line network, and suppose that the utilities, after link-formation costs, are $(10, 25, 25, 10)$. Now suppose we remove link l_{34} and add link l_{24} to create a star network. In the more general model, utilities after this rewiring might be $(11, 35, 11, 11)$. These two vectors will be ranked differently by different inequality measures within the Atkinson class. However, if we make an additional assumption that this kind of rewiring affects only those agents who gain or lose a link, then we can relate inequality to network structure in the more general setting.

Proposition 15. *Suppose there is one group, and for all pairs of efficient networks L and L' such that $L' = (L \setminus \{l_{ij}\}) \cup \{l_{jk}\}$, the transfer arrangements satisfy $\tau_l(L) = \tau_l(L')$ for all $l \neq i, k$. Then for all inequality measures in the Atkinson class, among the set of efficient*

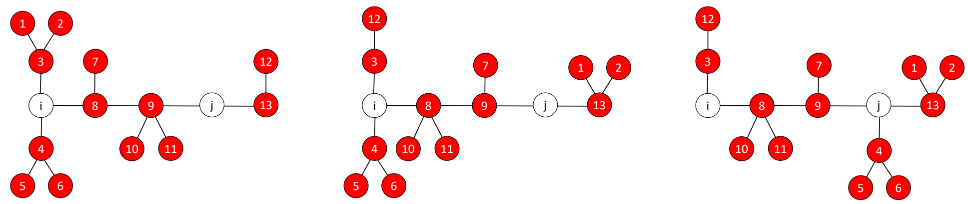

(A) Initial Network L

(B) L'

(C) Rewired network L''

FIGURE 1. An example of the rewiring used to find a contradiction in the proof of Lemma SA2 is shown. Panel (i) shows the initial network, Panel (ii) the interim network and Panel (iii) the final network after the rewiring is complete.

networks, star networks and only star networks maximize inequality, while line networks and only line networks minimize inequality.

Proof. We begin with a lemma:

Lemma SA2. *Suppose there is one group, and that for all pairs of efficient networks L and L' such that $L' = (L \setminus \{l_{ij}\}) \cup \{l_{jk}\}$, the transfer arrangements satisfy $\tau_l(L) = \tau_l(L')$ for all $l \neq i, k$. Then agents with a higher degree in L have a higher utility.*

Proof. Consider an efficient network L, and suppose agent i has higher degree than j. We will show that we can rewire a network in a way that weakly reduces i's utility and increases j's utility, but swaps the positions of i and j in the network such that on this new network i would have the same utility j had on the initial network. This will imply that i must have had a higher utility on the initial network.

Consider the following rewiring, an example of which is illustrated in Figure 1. As L is efficient, it is a tree by Proposition 12 and there is a unique path between i and j. If i is directly connected to j, we need not do any rewiring along this path. Otherwise, let there be $l \geq 1$ agents other than i and j on this path, and create the following two labelings of these agents: $i, i1, \ldots, il, j$ and $i, jl, \ldots, j1, j$. Thus $i1 = jl$, $i2 = j(l-1)$, and so on. Now if agent $i1$ has a link to an agent k on L and k is not on the path between i and j, we remove the link $l_{i1,k}$ and add the link $l_{j1,k}$. Repeat until all of $i1$'s links to agents not on the shortest path between i and j have been rewired. We now repeat for ik, with $k = 2, \ldots l$. Note that at each step of this rewiring we reach a connected tree network.

Consider now the neighbors of j not on the path between i and j. Match each of these neighbors to a different neighbor of i's who is also not on this path. As i has a higher degree than j, such a matching exists. For each such pair, we start with j's neighbor. Letting this neighbor of j be k, one by one we rewire each of k's links on L, except l_{ij}, to the neighbor of i that agent k was matched to. Let this agent be l. We then rewire each of l's links on L, except l_{il}, to agent k. Repeat for all of j's neighbors on L not on the path between i and j. Note again that at each step of this rewiring we reach a tree network. After all this rewiring, let the network that has been reached be denoted L'.

As in all the rewiring so far i and j have kept the same links, and as at each step an efficient network has been reached, by the premise of the Proposition, $\tau_i(L) = \tau_i(L')$ and $\tau_j(L) = \tau_j(L')$, so $u_i^\tau(L) = u_i^\tau(L')$ and $u_j^\tau(L) = u_j^\tau(L')$.

Finally, we consider the neighbors of i who were not on the shortest path to j and were not matched to any of j's neighbors. As i's degree is higher than j's, there exists at least one such agent. For all agents in this set, we remove their link to i and add a link to j. Let the network reached after this be denoted L''.

By Assumption 11(a), this increases j's utility and decreases i's utility, so $u_i^\tau(L) = u_i^\tau(L') > u_i^\tau(L'')$ and $u_j^\tau(L) = u_j^\tau(L') < u_j^\tau(L'')$. However, by construction, after this rewiring is complete i's position in L'' is identical to j's position in L (up to a relabeling of agents), while j's position in L'' is identical to i's position in L. Thus by Assumption 11(c), $u_i^\tau(L) = u_i^\tau(L'')$ and $u_j^\tau(L) = u_j^\tau(L'')$. We then have that

$$u_i^\tau(L) = u_i^\tau(L') > u_i^\tau(L'') = u_j^\tau(L).$$

□

We can now prove the Proposition. As shown in the proof of Proposition 6(ii), the star network can be reached from any efficient network L by rewiring links to the highest-degree agent in L. By Lemma SA2, the agent with the highest utility on L is the agent with the highest degree, and by Assumption 11(d) the net expected utility of this agent increases at each such step of the rewiring, while the net expected utility of all other agents weakly decreases. The argument from the proof of Proposition 6(ii) can then be applied again, and utilities become more unequal for any inequality measure in the Atkinson class.

The argument for the line network is equivalent. From any efficient network L, there is a rewiring to the line network that decreases the utility of the highest-degree agent at each step, which by Lemma SA2 is also the highest-utility agent, and increases the utility of all other agents. Thus utilities become more equal for any inequality measure in the Atkinson class. □

We will now consider the multiple-group case. With one group it was efficient for a network to form in which all agents are path connected to one another. We now make an assumption to ensure that this remains the case with multiple groups.

Assumption SA3 (Efficient Risk Sharing Across Groups). *For any network L with at least two components, there exist a risk-sharing agreement τ and a pair of agents i and $j \notin C_i$ such that all agents are weakly better off on $L \cup \{l_{ij}\}$ and some agents are strictly better off.*

Relative to the single-group case, agents from different groups provide each other with access to less correlated income streams. This increases the total surplus generated by risk sharing conditional on a given network being formed. Moreover, the presence of across-group links provides positive externalities to others insofar as it increases the marginal value of within-group links. This raises the question of how the additional surplus generated by across-group risk sharing should be split among the agents. We take a parsimonious approach to this issue by making two assumptions. The first assumption builds on the single-group analysis. It requires that agents receive at least the same marginal benefits they would receive if all agents were from the same group. The additional surplus generated must be split in such a way that each agent receives a weakly positive share.

Assumption SA4. *Consider a network L such that l_{ij} is essential on $L \cup \{l_{ij}\}$, and two allocations G, G' of the agents to groups. If all agents are from the same group under G, hence $G(k) = G(k')$ for all k, k', then*

$$u_i^\tau(L \cup \{l_{ij}\}, G') - u_i^\tau(L, G') \geq u_i^\tau(L \cup \{l_{ij}\}, G) - u_i^\tau(L, G).$$

This assumption requires that the essential within group links are weakly more valuable to agent i when the risk-sharing component i belongs to includes agents from multiple groups rather than all agents belonging to i's group.

Assumption SA5. *Consider two networks L and L' connecting the same sets of agents, and two allocations of the agents to groups G, G'. If L' can be reached from L by rewiring a link to i such that $L' = (L \setminus \{l_{jk}\}) \cup \{l_{ij}\}$, $i \neq j \neq k$, $l_{ij} \notin L$, $l_{jk} \in L$, G' contains agents from different groups, and under G all agents are from the same group, then*

$$u_i^\tau(L', G') - u_i^\tau(L, G') > u_i^\tau(L', G) - u_i^\tau(L, G).$$

Assumption SA5 is only a coarse partial ordering on utilities. While it implies that an agent's share of the additional surplus generated by across-group risk sharing increases as links are rewired to that agent, it makes no comparison between networks that cannot be reached by rewiring links to a single agent. In particular, following a rewiring to i it does not pin down how the payoffs of other agents change.

Proposition SA6. *Suppose all groups have the same utility function: $v_i = v_j$ for all i, j. With k different groups, there exist a $\bar{\kappa}_W > 0$ such that for all $\kappa_W < \bar{\kappa}_W$ a network is Pareto efficient if and only if it is a tree with $k-1$ across-group links.*

Proof. We begin by showing the "only if" direction. All Pareto-efficient networks are trees. First, by Assumption SA3, risk sharing among all agents is efficient, so L must connect all agents. Second, a Pareto improvement can be achieved on any connected non-tree network by implementing the same risk-sharing arrangement and deleting a superfluous link, thereby saving on costs.

We now show that efficient networks must also have exactly $k-1$ across-group links. We will show, by construction, that for any tree network with strictly more than $k-1$ across-group links there exists a Pareto improvement.

If there are more than $k-1$ across-group links in a tree network, we claim that there must exist an across-group link l_{ij} which, upon its removal, will result in a network $L' = L \setminus \{l_{ij}\}$ such that there exist two agents (k, l), with $G(k) = G(l)$ and $C_k(L') \neq C_l(L')$.

Suppose, by way of a contradiction, that there are $k' > k-1$ across-group links and that the claim in the previous paragraph is not true. As L is a tree network, removing all across-group links must then result in there being $k'+1$ components. If there are no agents from the same group in different components, this implies that there must be at least $k'+1 > k$ different groups, which is a contradiction. Thus there exist two components, each containing an agent from the same group. Denote these agents by k, l. As L is a tree, there exists a unique path between k and l on L, and as k and l are in different components following the removal of across-group links, there exists at least one across-group link on this path. Letting this link be l_{ij} proves the claim.

As k, l are in different components on L' but from the same group, the network $L'' = L' \cup \{l_{kl}\}$ will be a connected tree network with one less across-group link, and one more within-group link, than L.

On the network L'' we implement the same risk-sharing arrangement as before, with one exception. First, we identify the vector of consumptions that make them just as well off as on the original network and continue to satisfy the Borch rule:

$$\frac{\partial v_i(c_i(\omega))/\partial c_i(\omega)}{\partial v_i(c_i(\omega'))/\partial c_i(\omega')} = \frac{\partial v_j(c_j(\omega))/\partial c_j(\omega)}{\partial v_j(c_j(\omega'))/\partial c_j(\omega')} = \frac{\partial v_{i'}(c_{i'}(\omega))/\partial c_{i'}(\omega)}{\partial v_{i'}(c_{i'}(\omega'))/\partial c_{i'}(\omega')}$$

for all states ω, ω' and all $i' \neq k, l$.

As i and j save the cost of an across-group link, and utility is strictly increasing and concave in consumption, $c_i(\omega)$ and $c_j(\omega)$ must strictly decrease in all states ω. This additional consumption is passed on to agents k and l. As there is a strictly positive amount of remaining consumption in all states of the world, and utilities are strictly increasing in consumption, there exist feasible consumption vectors for agents k and l that strictly increase $\mathrm{E}(v(c_k))$ and $\mathrm{E}(v(c_l))$. Thus for all sufficiently small κ_w we have $\mathrm{E}(v(c_k)) > \kappa_w$ and $\mathrm{E}(v(c_l)) > \kappa_w$. We have therefore constructed a Pareto improvement.

We now show the "if" direction. Consider a tree network with $k-1$ across-group links. Suppose we implement a risk-sharing agreement in which $c_i(\omega) = c_j(\omega)$ for all i, j. As all agents' consumptions are equalized in all states, there is then no way in which link-formation costs can be redistributed, and the risk-sharing arrangement changed, without making someone worse off. Suppose, by way of a contradiction, that we can redistribute the link-formation costs, by forming a different tree network with $k-1$ across-group links, to generate a Pareto improvement. Holding consumption fixed, on the new network if some agents are better off, then some will be worse off. Thus to achieve a Pareto improvement, consumptions will have to be changed. Let $c'(\omega)$ be the new consumption vector. As the utility function $v(\cdot)$ is concave, Jensen's inequality implies that

$$\frac{1}{n}\sum_i v(c_i'(\omega)) < v\left(\frac{1}{n}\sum_i c_i'(\omega)\right) = \frac{1}{n}\sum_i v(c_i(\omega))$$

for all ω. Thus the average expected utility from consumption will decrease, and total link-formation costs have remained constant, so at least one agent must be worse off. This is a contradiction.

□

In our baseline model with CARA utilities, normally distributed incomes, and the Myerson value allocation rule, across-group underinvestment is possible but there is no within-group underinvestment. The same example establishes the possibility of across-group underinvestment in our more general setting. There is also never any within-group underinvestment in our more general setting, as we now show.

Proposition SA7. *There is never any within-group underinvestment.*

Proof. Consider any stable network L' and allocation to groups G'. Suppose, by way of a contradiction, that there is underinvestment within a group in L'. There must then be an essential link l_{ij} the planner could form to achieve a Pareto improvement. Stability of L' implies that either $u_i^\tau(L' \cup \{l_{ij}\}, G') - u_i^\tau(L', G') < c_w$ or $u_j^\tau(L' \cup \{l_{ij}\}, G') - u_j^\tau(L', G') < c_w$. Without loss of generality, suppose $u_i^\tau(L' \cup \{l_{ij}\}, G') - u_i^\tau(L', G') < c_w$. Consider now the alternative grouping G in which all agents are from the same group. In this case, by Assumption 11(d) and as l_{ij} is essential, $u_i^\tau(L' \cup \{l_{ij}\}, G) - u_i^\tau(L', G) \geq c_w$. Thus combining inequalities, $u_i^\tau(L' \cup \{l_{ij}\}, G) - u_i^\tau(L', G) > u_i^\tau(L' \cup \{l_{ij}\}, G') - u_i^\tau(L', G')$. This contradicts Assumption SA4. □

Consider the partial ordering in which an agent i is more central in a network L' than in network L if and only if L' can be reached from L by rewiring links only to i. The following result generalizes the result in the benchmark model that more centrally located agents within a group have higher incentive to create across-group links.

Proposition SA8. *Suppose that*

(i) *if there is one group, then for all efficient networks $L \cup \{l_{ij}\}$, $g(\overline{d}, |C_i(L)|, |C_i(L \cup \{l_{ij}\})|) = g(\overline{d}, |C_j(L)|, |C_j(L \cup \{l_{ij}\})|)$; and*

(ii) *there are two groups.*

Then for any efficient network L with across group-link l_{ij}, if it is profitable for an agent i to form l_{ij}, and the alternative efficient network L' can be reached from L by rewiring within-group links to i, then it is also profitable for i to form the link $l_{ij} \in L'$.

Proof. Let G' be the grouping of agents. Agent i is weakly better incentivized to invest in the across-group link l_{ij} on the network L' than on the network L if and only if

$$u_i^\tau(L,G') - u_i^\tau(L\setminus\{l_{ij}\},G') \leq u_i^\tau(L',G') - u_i^\tau(L'\setminus\{l_{ij}\},G'). \tag{1}$$

As L and L' are efficient, and l_{ij} is an across-group link on both L and L', all agents who are path connected to i on $L\setminus\{l_{ij}\}$ are from the same group as i, as are all agents path connected to i on $L'\setminus\{l_{ij}\}$. Thus on the networks $L'\setminus\{l_{ij}\}$ and $L\setminus\{l_{ij}\}$, by Assumption SA4 agent i must then get exactly the same payoffs as he would in the one-group case: $u_i^\tau(L\setminus\{l_{ij}\},G') = u_i^\tau(L\setminus\{l_{ij}\},G)$ and $u_i^\tau(L'\setminus\{l_{ij}\},G') = u_i^\tau(L'\setminus\{l_{ij}\},G)$, where G is the grouping in which all agents are from the same group. We can therefore rewrite equation (1) as

$$\begin{aligned} u_i^\tau(L,G') - u_i^\tau(L,G) + u_i^\tau(L,G) - u_i^\tau(L\setminus\{l_{ij}\},G) \quad \leq \quad & u_i^\tau(L',G') - u_i^\tau(L',G) \\ & + u_i^\tau(L',G) - u_i^\tau(L'\setminus\{l_{ij}\},G). \end{aligned} \tag{2}$$

Repeatedly applying Assumption SA5, $u_i^\tau(L,G') - u_i^\tau(L,G) < u_i^\tau(L',G') - u_i^\tau(L',G)$. Thus a sufficient condition for equation (2) to hold is that

$$u_i^\tau(L,G) - u_i^\tau(L\setminus\{l_{ij}\},G) \leq u_i^\tau(L',G) - u_i^\tau(L'\setminus\{l_{ij}\},G).$$

As we are in the one-group case and l_{ij} is essential on both L and L', $u_i^\tau(L,G) - u_i^\tau(L\setminus\{l_{ij}\},G) = u_i^\tau(L',G) - u_i^\tau(L'\setminus\{l_{ij}\},G) = g(\bar{d})$. This completes the proof.

□

B. Supported risk sharing

As with the previous section, this section provides a slightly more general and comprehensive treatment of analysis than in the main paper. This time the corresponding section of the main paper is Section 6.2. Again, we number replicated assumptions and results so that they correspond to those in Section 6.2 of the main paper, while new results are labeled with the prefix SA.

In this section we extend the model to capture the idea that having friends in common can reduce an agent's incentives to renege on an agreement. This might be because the friend in common is able to monitor actions and identify the guilty party in a dispute, or because reneging on the agreement will lead to a damaging reputation loss with the friend in common. While it is beyond the scope of this paper to fully explore these issues, and there is a vibrant literature that focuses on network-based enforcement of agreements (see, for example, Jackson et al. (2012), Wolitzky (2012), Ali and Miller (2013, 2016), Ambrus et al. (2014), Nava and Piccione (2014), and Ambrus et al. (2016)), in this section, motivated by this literature, we model the value of friends in common for enforcement by assuming that risk sharing between two agents is possible if and only if those two agents have a friend in common. This is known as closure (Coleman, 1988) and has long been thought important for cooperation because it enables collective sanctions to be imposed on a deviating agent—if an agent cheats on one of their neighbors, there are friends in common that can also punish the deviating agent.

A link in L is supported and can be used for risk-sharing if and only if it is part of a triangle (i.e., the complete network among three agents). Let $L'(L)$ be the spanning subgraph of L which contains only supported links. An illustration of this is provided in Figure 2. Risk-sharing agreements and rent distribution are as in Section 2 of the main paper. The only difference is that now risk sharing takes place on the network $L'(L)$ instead of L (but agents continue to pay to form links in L).

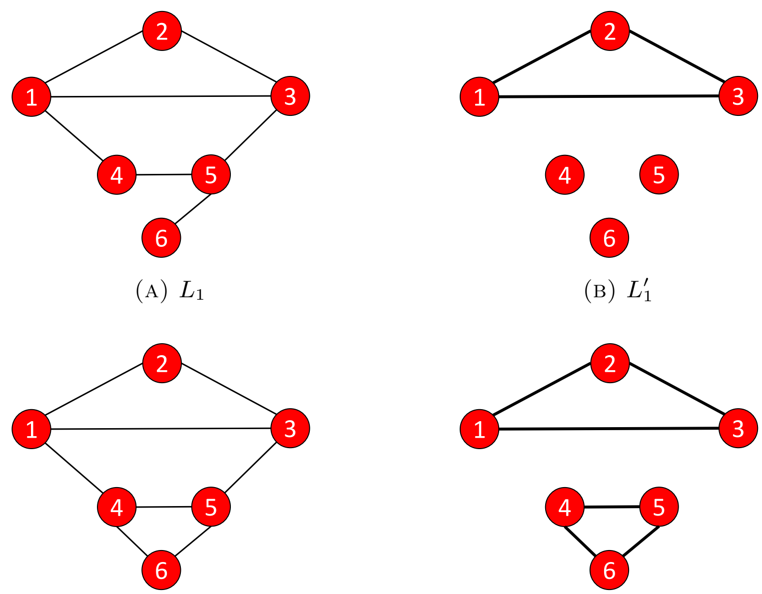

(A) L_1 (B) L'_1 (C) L_2 (D) L'_2

Figure 2. (A) Example of a network, L_1, among six villagers. (B) The links on L_1 that support risk sharing. (C) Example of a different network, L_2, among six villagers. (D) The links on L_2 that support risk sharing.

Before we can state our main result for this section, we need some new terminology. A network L is a *tree union of triangles* if it can be expressed as the union of $m \geq 2$ (non–node-disjoint) subnetworks, ordered as $L(N_1), \ldots, L(N_m)$, such that $\cup_{i=1}^{k} N_i \cap N_{k+1} = 1$ and each subnetwork $L(N_i)$ is a triangle. Thus each subnetwork in the sequence is a triangle that has exactly one node in common with the union of all the nodes in the subnetworks preceding it in the sequence. Two different tree unions of triangles are illustrated in Figures 3 and 4a. The tree union of triangles illustrated in Figure 3 is known the *Friendship graph* or Windmill network; in that network all triangles have the same node in common.

We will focus on risk sharing within a village. We denote the cost of forming a link by $\kappa = \kappa_w$. As before, we continue to focus on the parameter range for which risk sharing among

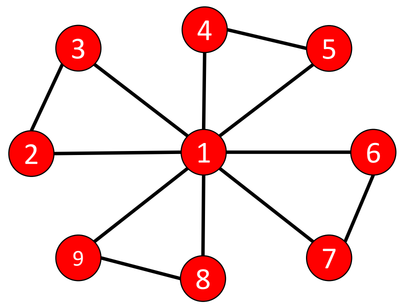

(A)

FIGURE 3. The Friendship graph on 9 nodes

all agents is efficient. As before, the surplus obtained from enabling risk sharing among two groups of agents is V. Proposition 16 shows that it is efficient for all agents to risk share if and only if $V \geq 3\kappa$, and that the efficient networks are then tree unions of triangles. Thus in comparison to Section 2 of the main paper, where agreements didn't need to be supported to be enforceable, tree unions of triangles play the role of tree networks.

Proposition 16. *Suppose the number of villagers $n \geq 3$ is odd.*

(i) *If risk sharing among all n agents is efficient, then the efficient risk-sharing networks are tree unions of triangles.*
(ii) *Risk sharing among all n agents is efficient for all n if and only if $V \geq 3\kappa$.*

The proof of Proposition 16 is fairly long, and is deferred until Section B.1. Here we offer some intuition. First, observe that any link that is not supported is costly to form but cannot be used for risk sharing. While in principle such a link might still be valuable as a means for supporting an agreement on another link, this requires a triangle to be formed with the other link, which would make it supported. Thus in an efficient network every link must be supported and must be part of some triangle. Given this, the most efficient way to organize links (among an odd number of agents) is to form a tree union of triangles. This creates distinct triangles in which no link is shared by two triangles. This might seem inefficient, but it is not, because it economizes on the number of triangles required. As a comparison, consider the tree union of triangles shown Figure 4a and the alternative network shown in Figure 4b; in the later, villagers 1 and 2 are connected to all other villagers and there are no other links. In the alternative network there are $n-2=7$ triangles, while there are just $(n-1)/2 = 4$ triangles in the tree union of triangles. Thus although the triangles in the alternative network all share the link l_{12}, meaning that for $n-3=6$ of the triangles only two additional links are required, there are more links in the alternative network than in the tree union of triangles ($3(n-2)-(n-3) = 15$ links in the alternative network, in comparison to $3(n-1)/2 = 12$ links in the tree union of triangles).

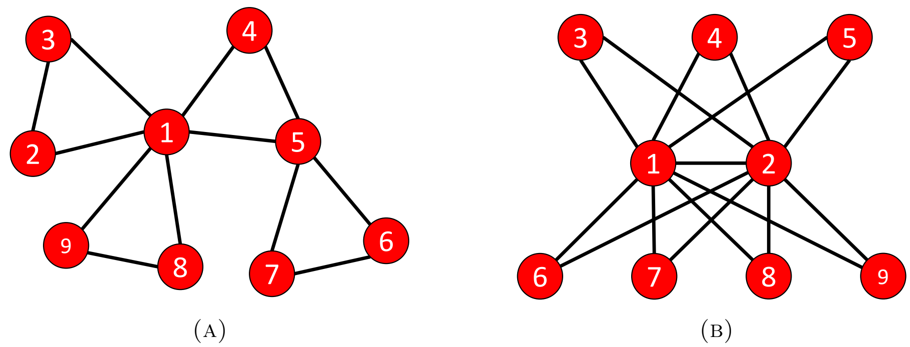

FIGURE 4. (A) A tree union of triangles connecting nine villagers. (B) An alternative network connecting nine villagers in which all villagers are able to risk share but all triangles share a common link, l_{12}, hence it isn't a tree union of triangles.

Jackson et al. (2012) find a class of networks they call social quilts to be those that can supporting risk-sharing agreements based on renegotiation proofness. Interestingly, tree unions of triangles are social quilts. The networks we identify through efficiency considerations based on the very simple condition of support for risk sharing to be possible would also be renegotiation proof in their setting. This provides further motivation for the simple approach to enforcement we take.

We now consider the stability of the efficient risk-sharing networks. Unlike the corresponding result in Section 4 of the main paper, all tree unions of triangles are equally pairwise stable and the empty network is now always pairwise stable. As risk sharing now requires three agents for an agent to extricate herself from an agreement while not leaving unsupported links, two links must be deleted at once. Thus we consider networks that are not only pairwise stable but also stable to multiple-link deletions. Such a network L must be pairwise stable and, for all agents i, $u_i(L) \geq u_i(L')$ for all L' that can be obtained by removing any of i's links in L.

As before, we let V be the constant value of reducing the number of risk-sharing groups by 1.

Proposition SA9. *In a tree union of triangles, agent i receives a net payoff $|N(i;L)|((V/3)-\kappa)$. A tree union of triangles L is pairwise stable if and only if $3V/5 \leq 3\kappa \leq 2V$. A tree union of triangles L is pairwise stable, and also stable to multiple-link deletions if and only if $3V/5 \leq 3\kappa \leq V$. The empty network is always pairwise stable.*

The full proof is in Section B.1. Analogously to before, on efficient networks all links are essential and make the same expected contribution to total surplus for a random arrival order of the agents (as can be used to calculate the Myerson value). Moreover, these benefits are shared equally among two agents when they have a link. Collectively, a triangle of links contributes an amount $2V$ to total surplus. In a tree union of triangles, each link is part of only one triangle, and thus each link contributes on average $2V/3$. As these benefits are split evenly among the agents forming the link, they each get $V/3$, while it costs each agent κ to form a link. Hence each agent receives a net payoff of $|N(i;L)|((V/3)-\kappa)$, which is again proportional to her degree.

If an agent deletes a link, exactly one of her other links becomes unsupported. Thus the agent's payoff decreases by $2V/3$, but there is a saving of only κ in costs. Thus a network is stable to individual link deletions if and only if $2V \geq 3\kappa$, while it is stable to multiple-link deletions if and only if $V \geq 3\kappa$ (which holds by the maintained assumption that it is efficient for all agents to risk share with one another).

Consider an agent's incentives to form an additional (superfluous) link. In any network, agents can benefit only from forming links that would be supported, so that they can be used for risk-sharing. The key to the proof is showing that on a tree union of triangles, for any superfluous link that would be supported upon its formation, there are the same incentives to deviate from it. Thus there is a profitable deviation to form any superfluous link in any tree union of triangles if and only if it is profitable to form the link, as shown in Figure 5. As it is profitable to form this additional link if and only if $V/2 \geq 3\kappa$, a tree union of triangles is robust to the pairwise addition of a link if and only if $V/2 \leq 3\kappa$.

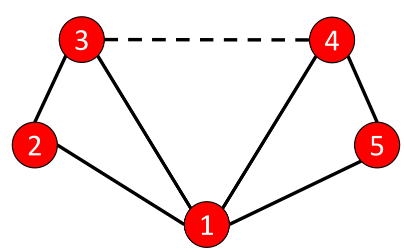

(A)

Figure 5. The Friendship graph on five nodes with a possible deviation shown by the dashed line

Finally, to see that the empty network is always stable, just note that on this network an additional link will not be supported, and so will not facilitate any risk sharing; thus there are no incentives to form any link. The stability of the empty network and the need for groups of at least three agents to support risk sharing suggest that it might be reasonable to permit coalitions of three agents to form links among themselves. We do so with the minimal possible extension to pairwise stability that facilitates such deviations.

A network is tripletwise stable with respect to expected utilities $\{u_i(L)\}_{i\in\mathbf{N}}$ if and only if it is pairwise stable and for all $i, j, k \in \mathbf{N}$, if two or more of l_{ij}, l_{ik}, l_{kj} are not in L and $\hat{L}$ is the union of network L with these three links, and then if $u_i(\hat{L}) \geq u_i(L)$ and $u_j(\hat{L}) \geq u_j(L)$ with at least one of these inequalities strict, then $u_k(\hat{L}) < u_k(L)$. In words, tripletwise stability requires a network to be pairwise stable and does not allow any set of three players to be able to benefit by forming the remaining links among themselves (thereby facilitating direct risk sharing among themselves).

Proposition 17.

(i) *If there exists an efficient tripletwise stable network then all friendship networks are tripletwise stable, and for a non-empty range of parameter specifications the only efficient networks that are tripletwise stable are friendship networks.*

(ii) *For all inequality measures in the Atkinson class, among the set of efficient networks, friendship networks maximize inequality and are the only efficient networks that maximize inequality.*

This result is analogous to results in Proposition 6 in Section 4 of the main paper. There, a star network was the most efficient stable network, but also the most unequal. Proposition 17 shows that this result generalizes to the case in which links must be supported to facilitate risk sharing, but with friendship networks taking the place of star networks.

The proof of Proposition 17 is in section B.1. The basic intuition for this result mirrors the intuition for the corresponding result in the main paper (Proposition 6). Groups of three agents have stronger incentives to deviate and form links among themselves to facilitate risk sharing when they are further apart. Among the set of efficient networks, the relevant distances are minimized by the friendship network. In terms of inequality, agents' net payoffs are again proportional to their degrees, and the total number of links is constant for all tree unions of triangles connecting n agents. Further, in any tree union of triangles all agents must have degree at least 2. The friendship network therefore minimizes the possible degree for all but one agent while maximizing the possible degree for the remaining agent. The star network did the equivalent thing in Section 4 of the main paper, and this was the key property of the star network that led it to generate the most inequality for any inequality measure in the Atkinson class. The argument establishing that the friendship network now generates the most inequality for any inequality measure in the Atkinson class is the same.

B.1. **Proofs.**

B.1.1. *Proof of Proposition 16.*

Proof. **Part (i)**: Consider an efficient network L. As the network is efficient, all agents are then in the same risk-sharing component, so $L'(L)$ is connected. Further, as the network is efficient, every link must be supported, so $L'(L) = L$. This means that the network can be decomposed into a set of triangles (where the triangles can share nodes and links with one another and every node is part of at least one such triangle). There may be more than one such decomposition for L. Moreover, as $L'(L) = L$ is connected, these triangles must be connected to one another so that there is a path from every triangle to every other triangle.

It is therefore possible to order the triangles in the decomposition so that as the triangles are added to the network in this sequence there is always a unique component.

Figure 6 gives an example of this triangle decomposition. In this example there is a redundant triangle such that the original network can be constructed from a set of triangles that excludes it. It doesn't matter which decomposition is selected, or whether the redundant triangle is included or not.

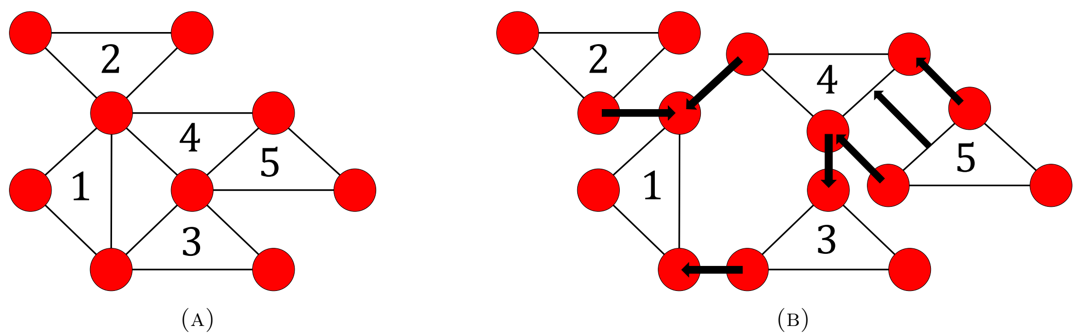

FIGURE 6. (A) A network in which every link is supported (i.e., $L'(L) = L$). (B) Representation of this network as a sequence of triangles; by combining triangles 1, 2, 3, 4, and 5, the original network is obtained. The arrows in (B) indicate which links and nodes are combined in this construction.

Consider an efficient network and an associated triangle decomposition. Suppose we create the network associated with the decomposition. Thus if there are k triangles in the decomposition, we are then left with a network consisting of k disjoint triangles (this will require creating duplicate nodes and links). This network has k components, $3k$ nodes, and $3k$ links. We then order these triangles, and recombine them to create the efficient network. We start with triangle 1, add triangle 2 so that 1 and 2 now form a network component, add 3 so that triangles 1, 2, and 3 form a component, and so on. Thus after each step in the sequence the number of components is reduced by one.[3] We consider how the number of links and nodes in the network must evolve along such a sequence.

When we connect an unconnected triangle to an existing set of connected triangles (which we term the component), the ways in which this might be done can be partitioned as follows: The new triangle can share 3 nodes with existing nodes, 2 nodes with existing nodes, or 1 node with existing nodes. In the case of sharing 3 nodes, no new nodes are being added to the network, but new links might be. As, by construction, all nodes in the component are already supported, it is without loss of generality to ignore such operations when searching for minimally connected networks that enable risk sharing among all agents (i.e., efficient

[3]For the example given in Figure 6, the sequence 1, 2, 3, 4, 5 results in a reduction in the number of components of one at each step, while the sequence 2, 5, 3, 1, 4 would not.

networks).[4] Figure 7 shows two examples of this. The addition of the triangle as shown in panel (A) has no effect on the number of links or nodes in the network (see panel (B)), while the addition of the triangle as shown in panel (C) increases the number of links but not the number of nodes in the network (see panel (D)).

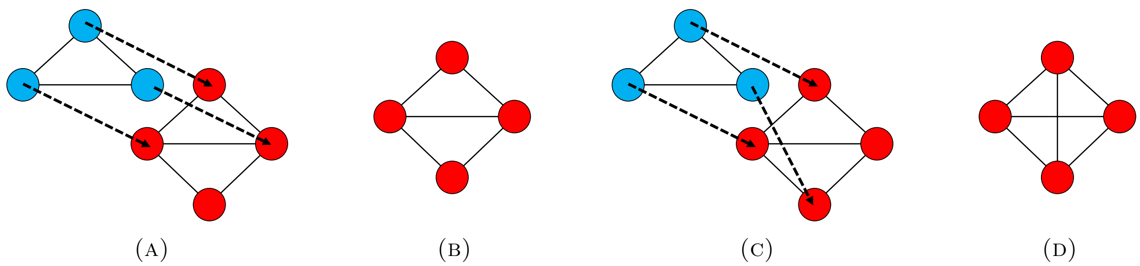

Figure 7. Panels (A) and (B) illustrate the addition of a triangle (blue) to a risk-sharing component (red) in which all nodes are shared by a single existing triangle. Panels (C) and (D) illustrate the addition of a triangle (blue) to a risk-sharing component (red) in which all nodes are shared by existing triangles.

When a triangle is added that shares two nodes, it can either share one link as well, or share no links. When a triangle is added that shares just one node, it cannot share any links. These three possibilities are enumerated below and illustrated in Figure 8.

(a) The triangle shares two nodes, one node with each of two different triangles. In this case, we increase the number of links in the component by 3 and increase the number of nodes in the component by 1.

(b) The triangle shares two nodes, both with the same triangle. In this case, we increase the number of links in the component by 2 and increase the number of nodes in the component by 1.

(c) The triangle shares one node. In this case, we increase the number of links in the component by 3 and increase the number of nodes in the component by 2.

Following the decomposition, along the sequence of recombining the triangles we do one of the above three operations at each of the $k-1$ steps. There are n nodes, where (by assumption) n is an odd integer. Suppose it is feasible to do any combination of the operations $(a)--(c)$, in any order, to arrive at n nodes. We always start with the component being a triangle, with three nodes and three links. This means that the number of nodes in the original network is $n = 3+a+b+2c$, where a is the number of (a) operations, b the number of (b) operations, and (c) the number of c operations. As the initial network is efficient, this sequence of operations must minimize the number of links in the resulting network, conditional on enabling all n agents to risk share. Assuming that any sequence of operations

[4]For example, the redundant triangle in Figure 6a could be added last, in which case it would share three nodes and three links with the component and its addition would add no new links or nodes to the component.

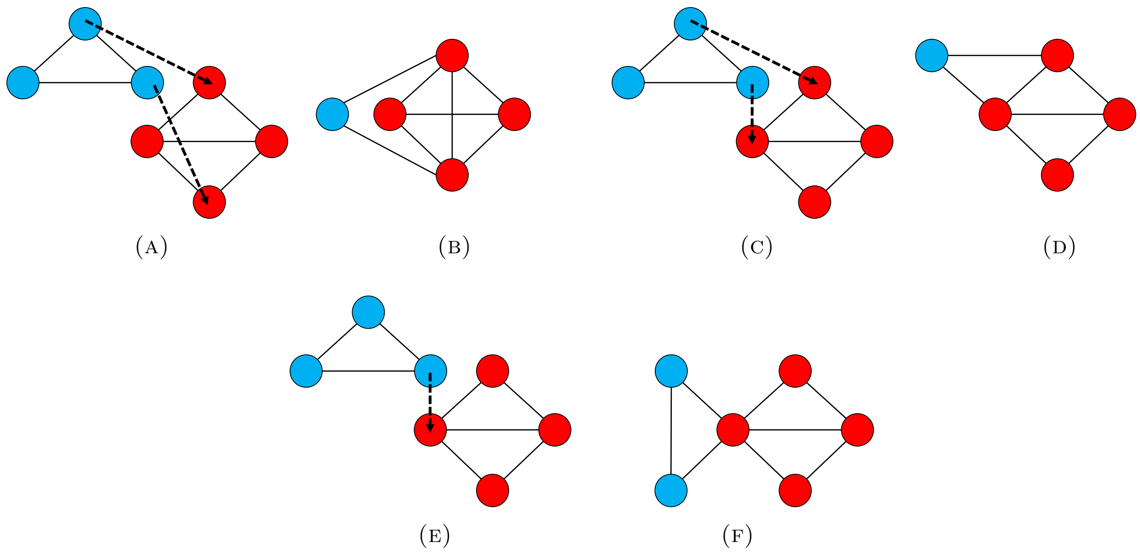

Figure 8. Panels (A)-(D) illustrate the possible ways in which triangles that share two nodes can be added, while panels (E) and (F) illustrate the possible ways in which triangles that share one node can be added.

is feasible, the sequence of operations must minimize $3+3a+2b+3c$ subject to $3+a+b+2c=n$. As n is odd, this is uniquely achieved by setting $a=b=0$ and $c=(n-3)/2$. (Incidentally, when n is an even number greater than 3, it can be seen that this is instead achieved by setting $a=0$, $b=1$, and $c=(n-4)/2$; thus when n is even the structure of the efficient networks is similar to the structure of the efficient networks when n is odd.) Note that the efficient network is constructed through sequentially adding triangles such that at each step in the sequence the added triangle shares exactly one node with the triangles already added. But this is just the definition of a tree union of triangles. This implies that this sequence of operations is feasible and that the efficient networks are tree unions. When there are n nodes, with n odd, any tree union of triangles is efficient, and no other network is efficient.

Part (ii): We have established that an efficient network is a tree union of triangles, so the number of links under full risk sharing is $3(n-1)/2$. Since every link incurs the cost κ for both agents, full risk sharing is therefore efficient if and only if $V(n-1) \geq 3(n-1)\kappa$.

□

B.1.2. *Proof of Proposition SA9.*

Proof. Let L be a tree union of triangles and consider one such triangle τ. Without loss of generality label the agents in this triangle $1,2$ and 3. Consider adding the agents to the network in an arbitrary permutation. For any such permutation, the last agent to be added from the set $\{1,2,3\}$ completes the triangle τ. As L is a tree union of triangles, prior to completion of this triangle, agents 1, 2 and 3 cannot risk share with each other and must be in different risk-sharing components of the network $L'(L)$. Thus, the completion of the triangle τ reduces the number of risk-sharing components by 2, generating additional value $2V$. So in the Myerson value calculation, the presence of the triangle τ generates an additional expected payoff for each of the agents $\{1, 2, 3\}$ equal to $2V/3$ (as each is last to arrive in $1/3$ of the permutations, and so each completes τ, thereby generating risk-sharing benefits of $2V$, in $1/3$ of the permutations). Thus, in a tree union of triangles an agent's payoff before link formation costs is $|N(i; L)|V/3$.

As after a link is deleted one sharing triangle is lost (as no triangles share links in a tree union of triangles) deleting a link causes that agent to lose benefits $2V/3$. Thus an agent does not want to delete any one of their links in a tree union of triangles if and only if $2V/3 > \kappa$.

Consider now the incentives of two unconnected agents i and j to form an additional link l_{ij}. As i and j are unconnected they are in different risk-sharing triangles. If the link l_{ij} does not create a new triangle with some agent k, then it does not facilitate any additional risk-sharing on any subnetwork that can be reached by adding the agents in sequentially. Hence, agents i's and agent j's Myerson value is unaffected, but they pay a cost κ each to form the link. As such deviations are unprofitable, we can restrict attention to link l_{ij} that would be part of a triangle once added. Let τ be the triangle on $L \cup \{l_{ij}\}$ between agents i, j and some other agent k. Thus $l_{ik} \in L$ and $l_{jk} \in L$. Upon its completion (i.e., when the last of i, j or k is added for a given arrival order) the triangle τ facilitates new risk sharing between agents i, j and k thereby reducing the number of risk-sharing components by 2, if and only if both i and k and j and k were not able to risk-share with each other before. As i and k are connected on L, and L is a tree union of triangles, they must be part of a risk-sharing triangle on L with another agent k'. Hence they are already risk-share with each other if and only if k' has already been added (i.e., k' is not the last agent to be added in the permutation among the four agents i, j, k, k'). This happens in $3/4$ of the permutations. Similarly, agents j and k must also already be part of a risk-sharing triangle with another agent $k'' \neq k'$ (were $k'' = k'$ this would imply that two risk-sharing triangles in L share a link $l_{k'k}$, but then L would not be a tree union of triangles). So risk sharing among agents j and k is also already possible if k'' is not last in the permutation among the four agents i, j, k, k'' (see Figure 9).

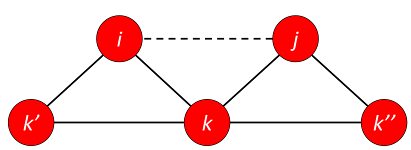

FIGURE 9. Adding a new link that is supported. The new link is the dashed link.

.

The probability that the new triangle τ generates benefits $2V$ upon being added is $2(3!)/5!$. There are $5!$ permutations of i, j, k, k', k''. There are $3!$ permutations of i, j, k. For each of these permutations, there are two permutations in which k'' and k' are the last two elements for a permutation of i, j, k, k', k''. Hence the probability that k'' and k' are both after all of i, j and k in a random permutation is $2(3!)/5! = 1/10$. The probability that τ generates benefits V is the probability that either k' is after all of i, j and k or k'' is after all of i, j and k, but k' and k'' are not both after all of i, j and k. The probability that k' is after all of i, j and k is $1/4$. The probability that k'' is after all of i, j and k is $1/4$. Thus the probability that τ generates benefits V is $1/2 - 1/10 = 2/5$. Thus the expected increase in surplus generated by the link l_{ij} is $2V/5 + 2V/10 = 3V/5$. These benefits accrue to agent i with probability $1/3$, to agent j with probability $1/3$ and to agent k with probability $1/3$. Thus agent i and j have a profitable pairwise deviation to form the link if and only if $V/5 > \kappa$. Thus a tree union of triangles is pairwise stable if and only if $3V/5 \leq 3\kappa \leq 2V$ as claimed.

When it is possible to delete multiple links at once, a lower bound on the benefit lost per link deleted in a tree union of triangles is $V/3$. Recall that in a tree union of triangles an agent's payoff before link formation costs is $|N(i; L)|V/3$. Thus, if after a deletion all remaining links still facilitate risk-sharing, only $V/3$ will be lost per link deleted. If after the deletion some of remaining links are not able to facilitate risk-sharing, the loss per link will be greater. The bound of $V/3$ is tight. For example, if an agent simultaneously deletes all their links this bound will be achieved. As the amount saved in link formation costs from deleting a link is κ, it then follows that a network is pairwise stable and also stable to multiple link deletions if and only if $3V/5 \leq 3\kappa \leq V$.

Finally, note that in the empty network the incremental benefits of forming a link l_{ij} are 0 as it does not permit any risk-sharing. Hence, the empty network is pairwise stable. □

B.1.3. *Proof of Proposition 17.*

Proof. **Part (i)**: By Proposition 16, efficient networks are tree unions of triangles and by Proposition SA9 all these networks are equally pairwise stable. Thus any difference in stability between the efficient networks in terms of stability must be due to tripletwise deviations that form at least two links among the three agents. Thus, there are two cases to consider—when a triplet deviates by adding two links and when a triplet deviates by adding three links.

We consider these cases shortly. Before that, it is helpful to define a new distance measure for tree unions of triangles. By the definition of a tree union of triangles, any L tree union of triangles can be decomposed into a sequence of triangles such that each triangle in the sequence shares a single node with triangles earlier in the sequence. Thus, for any two nodes i and j on a tree union of triangles L, there is a minimal subset of these triangles that must be added for i and j to be path connected. We define the *triangle distance* between i and $j \neq i$ on a tree union of triangles L to be the cardinality of this set of triangles and denote the distance by $\Delta(i, j; L)$. For example, in Figure 10 we have $\Delta(i, j; L) = 4$, $\Delta(i, k; L) = 5$ and $\Delta(k, j; L) = 1$.

.

Case A (two links): For the additional links to be valuable they must create a triangle. Thus, when the triplet adds two links, the other link must already be present. Without loss,

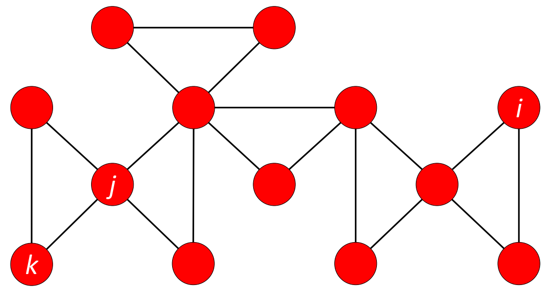

FIGURE 10. A tree union of triangles L

label this triplet i, j, k and suppose that $l_{ij} \in L$ is the link in this triangle that is already present. We let τ denote this triangle between i, j and k. As L is a tree union of triangles, l_{ij} must be supported and there must be an agent k' such that $l_{ik'} \in L$ and $l_{jk'} \in L$. Figure 11 shows the subnetwork of L among agents i, j, k and k', including the links that would be formed by the deviation.

If agents i, j and k deviate to form τ, the probability that agents i and j could risk-share without the links l_{ik} and l_{jk} at the time τ is completed, for a random arrival order, is the probability that agent k' has already been added—i.e., 3/4 (the probability that k' is not last to arrive out of i, j, k, k'). Figure 11 shows the subnetwork of L among agents i, j, k and k', and the links that would be formed by the deviation which are dashed.

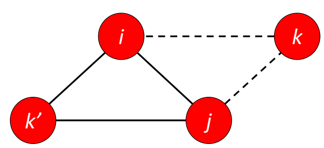

FIGURE 11. A subnetwork of a tree union of triangles L induced by agents i, j, k and k' is shown by the solid links. A possible tripletwise deviation among agents i, j and k through the creation of the links l_{ik} and l_{jk} is shown by the dashed links.

.

The probability that agents i and k can already risk share depends on whether there would be a supported path between them when the triangle τ is completed. Recall that $\Delta(i, k; L)$ is the triangle distance between i and k. Without loss, suppose that $\Delta(i, k; L) \geq \Delta(j, k; L)$. A supported path between i and k will exist upon the completion of τ if and only if all agents in the triangles counted in the triangle distance between i and k are already present. This requires $1+2\Delta(i, k; L)$ agents, including i, j and k to be present when τ is completed. Letting $x = 1 + 2\Delta(i, k; L)$, the probabilty of this is the probability that i, j or k arrive last in the arrival order among these x agents, i.e., $3(x-1)!/x! = 3/x$.

There are two possibilities to consider (given that $\Delta(i, k; L) \geq \Delta(j, k; L)$) when calculating the probability that agents j and k can already risk-share upon the completion of τ. First, we could have $\Delta(i, k; L) = \Delta(j, k; L)$, in which case agents j and k will be path connected upon the completion of τ if and only if i and k are path connected upon the completion of τ. Moreover, in this case, i and k (and thus also j and k) are path connected upon the completion of τ only if the triangle (i, j, k') is present upon the completion of τ. An example of this case is shown in panel (A) of Figure 12. Thus the probability that the triangle τ generates benefits $2V$ upon its completion is $1/4$ (i.e., the probabilty k' is last to arrive of i, j, k and k'), and the probability it generates benefits of exactly V upon its completion is $1 - 1/4 - 3/x$. So the expected benefits τ generates are $V(5/4 - 3/x)$.

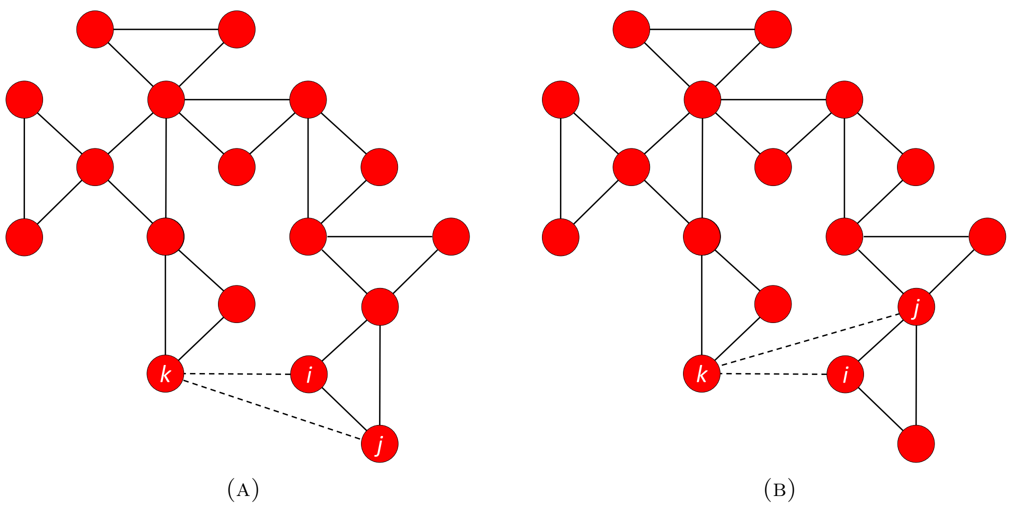

FIGURE 12. (A) A triplet deviation for agents i, j and k in which the triangle distance between both agents i and k and agents j and k is 6. (B) A triplet deviation for agents i, j and k in which the triangle distance between agents i and k is six and between agents j and k is five.

The second possibility is that $\Delta(j, k; L) = \Delta(i, k; L) - 1$. An example of this case is shown in panel (B) of Figure 12. Consider a labeling of agents consistent with l_{ik} and l_{jk} being the new links and l_{ij} being already present. If $\Delta(i, k; L)$ is the same for both possibilities, then the incentives to deviate in this case are always weaker. This is because we can match permutations such that permutation by permutation the risk-sharing value attributable to the new links, upon completion of τ, is weakly lower now than under the first possibility. For example, in Figure 12(a), consider any permutation in which k' is the last agent to be added and j is the second to last agent to be added. In this case, τ generates value $2V$ as upon the addition of j none of i, j or k would be able to risk-share with each other without the new links. Now consider the same sequence of agents for the example shown in Figure 12(b) (where, in this figure, agents j and k' have swapped position in comparison to before). Now, when j is added, agents k and j would be able to risk-share without the new links because they will be still be path connected. Hence the new links only generate additional risk-sharing benefits of V.

On any tree union of triangles L there are at least two leaf triangles (such that two of the agents in the triangle have degree 2). Thus, if the maximal triangle distance between any two nodes on L is z, there is a pair of connected nodes i, j whom are both triangle distance z from some other node k. Hence, for the triplet of agents with the strongest incentives to deviate by forming two links on any tree union of triangles L, the triangle distance between the agents without links will be equal to the maximum triangle distance in the network. Thus, the maximum incentives over all triplets in a tree union of triangles L, for them deviate by forming two links, is increasing in the maximum triangle distance on the network which we call the triangle diameter. For example, for the tree unions of triangles shown in Figure 12 the triangle diameter is 6 and for the deviation shown in panel (A) the triangle distance between agent i and k and between agents j and k are equal to 6. The friendship network has a triangle diameter of 2, which is strictly lower than for any other tree union of triangles that is not a friendship network. The incentives for some triplet to deviate on a tree union of triangles is therefore strictly lower on a friendship network than any other tree union of triangles.

Case B (three links): Again the additional links must create a triangle and facilitate risk sharing among the agents. Without loss, label these agents i, j and k and the triangle they create from their deviation τ. For the three links to be added, these agents must all initially be in different risk-sharing triangles. Moreover, as L is a tree union of triangles, there is a unique set of risk-sharing triangles among any two of them that connects them. Let X be the set of agents in the risk-sharing triangles connecting i and k, let Y be the set of agents in the risk-sharing triangles connecting j and k, and let Z be the set of agents in the risk-sharing triangles connecting i and j.

The triangle τ, upon its completion for a random arrival order, permits new risk sharing among the triplet generating value $2V$ if and only if none of the following conditions hold: (i) agent i or k is the last to arrive among the agents in the set X; (ii) agent j or k is the last to arrive among the agents in the set Y; (iii) agent i or j is the last to arrive among the agents in the set Z. This is a complex (although tractable) combinatorial calculation to write down. However, for our purposes, what matters are the following two facts: (a) this probability increases as additional agents are added to any of the sets X, Y or Z (whether these agents are present in the other sets or not); (b) this probability increases as the sets X, Y and Z become less overlapping holding their individual cardinalities fixed. For example, holding the sets Y and Z fixed, and the cardinality of X fixed, if $|X \cup Y|$ or $|X \cup Z|$ increases the probability increases.

The triangle τ, upon its completion, permits new risk sharing among the triplet generating value V or $2V$ if and only if at most one of the following conditions hold: (i) agent i or k is the last to arrive among the agents in the set X; (ii) agent j or k is the last to arrive among the agents in the set Y; (iii) agent i or j is the last to arrive among the agents in the set Z. Again, for our purposes, what matters is the following two facts: (a) this probability increases as additional agents are added to any of the sets X, Y or Z (whether these agents are present in the other sets or not); (b) this probability increases as the sets X, Y and Z become less overlapping holding their individual cardinalities fixed.

As i, j and k are in different risk-sharing triangles on L (and no triangles in a tree union of triangles share a link), we have the following inequalities on cardinalities:

(1) $|X|, |Y|, |Z| \geq 5$,
(2) $|X \cup Y|, X \cup Z|, |Y \cup Z| \geq 7$,
(3) $|X \cup Y \cup Z| \geq 7$.

For any given tree union of triangles, when considering the stability of it with respect to these deviations, we are interested in the triplet of agents that has the strongest incentives to deviate. These incentives are again minimized in the friendship graph. The friendship graph achieves the aforementioned bounds for any triplet of agents that can deviate in this way. Moreover, it is straight-forward to see that for any other tree union of triangles, the bounds are not achieved—there must exist two agents with a tree distance greater than 2, and without loss these agents can be labeled i and k such that $|X| \geq 7$.

Part (ii): By Proposition SA9 the payoff of each agent is proportional to its degree. Among tree unions of triangles the friendship graph maximizes the degree of the highest degree agent and set the degree of all remaining agents to 2. As all agents in all tree unions of triangles must have degree of at least 2 the argument used in the proof of Proposition 6(ii) in the main paper goes through unchanged.

□

C. Permitting some free links

This section replicates and then extends Section 6.4 in the main paper.

In practice, relationships are formed for many reasons, and there will be some relationships that exist for reasons unrelated to risk sharing but nevertheless permit risk sharing. These links might, for example, represent family relationships or close friendships formed in childhood. In effect, these are relationships are formed at no cost the purpose of risk sharing, providing another explanation for why real-world risk-sharing networks are denser than tree networks. We extend our baseline model to permit this possibility.

Let $\widehat{L}$ denote the exogenously given set of links that can be formed at no cost. As, by the Myerson value calculation, a link strictly increases the expected utility an agent receives in a risk-sharing arrangement, we assume that all such links are always formed. The network $\widehat{L}$ will consist of a set of components, each of which contains agents from the same group. For each component C, we identify an agent $i^*(C) \in \text{argmin}_i \max_j md_{ij}(C)$. This is an agent who has the lowest maximum Myerson distance to any other agent in component C. We will refer to agent $i^*(C)$ as the *Myerson distance central agent* in component C and let C_i denote the component to which i belongs. Considering all components, we then have a set of Myerson distance central agents $I^* = (i^*(C))_C$. Finally, we identify a Myerson distance central agent associated with the largest distance, $i^{**} \in \text{argmax}_{i^* \in I^*} \max_{j \in C_{i^*}} md_{i^*j}$.

We dub a network generated by forming all free links, and the links $l_{i^*i^{**}}$ for all $i^* \neq i^{**}$, a *central-connections network*. Suppose there are k different groups and $k' \geq k$ initial components. The set of efficient network then comprises the set of networks in which there is a single component and $k' - k$ within group links are formed (i.e., the minimal number of costly links that must be formed for there to be a single component).[5] Central-connections networks are always efficient. They are also the most stable networks within the class of efficient networks.

Proposition 18. *Suppose there is one group. If any efficient network is stable, then all central-connections networks are stable.*

Proposition 18 shows that when some within-group are formed at no cost, the most stable efficient network forms all additional links required for risk sharing with a single agent. As payoffs are proportional to degree, this again pushes villages toward inequitable outcomes.

We now prove Proposition 18.

Proof. Consider two components C and C'. For two agents i, j in component C, recall that $md(i, j, C)$ equals 1/2 less the probability that a path exists between i and j on C upon the arrival of i. Suppose now we take two components C and C'. Let agents i, k be in component C and agents j, k' be in component C', and form the bridging link $l_{kk'}$. The probability a path exists between i and j upon i's arrival is now is equal to the probability that a path exists between i and k on C multiplied by the probability that a path exists between k' and j on C'. This is because these events are independent, and when both path exist agents k and k' must have arrived before i and so the link $l_{kk'}$ must be present. It follows that

$$\underset{i,j}{\text{argmax}}\, md_{ij}(C \cup C' \cup \{l_{kk'}\}) = \{i, j : i \in \underset{l}{\text{argmax}}\, md_{lk}(C), j \in \underset{l}{\text{argmax}}\, md_{lk'}(C')\}.$$

Thus the network generated by forming all free links, and the links $l_{i^*i^{**}}$ for all $i^* \neq i^{**}$ minimizes the maximum Myerson distance on an efficient network and, by Lemma 4, is stable if any other efficient network is stable. □

When there are multiple groups, central-connections networks within group with the agent i^{**} providing the across group link(s) continue to work well. With multiple groups, agents' incentives to form superfluous within-group links depend on two things. First, as before, whether the link will be essential for a random arrival order, and second, unlike before, how many agents from other groups the link provides access to upon i's arrival when it is essential. Incentives to form a superfluous within-group links are increasing in the number of agents from other groups the link provides access to, and decreasing in the number of agents

[5] As before, the same set of risk-sharing arrangements can be implemented on any given component, and as expected utility is transferable, given that formation costs have been minimized, any point on the Pareto frontier can be obtained.

within-group the link provides access to. These considerations make superfluous links to the agent providing the across group link(s) particularly valuable. However, by construction the network generated by forming a central-connections network within-group, with the agent i^{**} providing the across group link(s), minimizes the maximum probability that a superfluous link to the agent providing the across group link(s) will be essential for a random arrival order. It thus minimizes the maximum incentives for an agent to form a superfluous link within-group to the agent providing the across group link(s).

Considering the incentives within a group to efficiently form an across-group essential link, a central-connections networks within-group is also likely to do well. By Lemma 9 more Myerson central agents have better incentives to form across group links. While central-connections networks maximize a slightly different notion of the centrality of the most central agent, in this case agent i^{**}, these measures of centrality are likely to be highly correlated. We therefore expect central-connections networks within-group to provide relatively good incentives for across group links to be formed.

D. General tensions between stability, efficiency and inequality

Like earlier sections, here we provide a more detailed treatment of a corresponding Section in the main paper. The corresponding section this time is Section 6.5.

The purpose of this section is to document a general fundamental tension between equality and efficient stable networks. We begin by relating different graph-theoretic concepts to stability, efficiency, and inequality.

D.0.1. *Equality.* We would like to say something general about inequality for all inequality measures in the Atkinson class on formed networks for *any* symmetric payoff function $u : L \to \mathbb{R}$. Unfortunately, without further restrictions on how network positions translate into payoffs, it is impossible to compare two networks in general. However, it is possible to pose and answer in general the question of when payoffs will be guaranteed to be perfectly equitable.

We proceed under the assumption that only agents' network positions matter for their payoffs—specifically, we require agents in identical network positions to receive the same payoffs. Intuitively, then, if all agents are in identical positions, they must receive equal payoffs. The set of networks for which this holds, thereby guaranteeing perfectly equitable outcomes, will be a useful benchmark that helps identify a general tension between equality and efficiency/stability.

In order to formalize the idea that agents are in identical network positions, we need to introduce some graph theory notation and terminology. We limit attention to connected networks. Every network is implicitly labeled, and we identify the set of labels with the set of nodes $\mathbf{N}$. Two networks L_1 and L_2 are called *isomorphic*, written $L_1 \sim_I L_2$, if they coincide up to labeling, that is, up to a permutation of $\mathbf{N}$. They are also *automorphic* if given the permutation of nodes associated with the isomorphism f, $l_{f(i)f(j)} \in L_2$ if and only if $l_{ij} \in L_1$.

When networks L_1 and L_2 are *automorphic* we write $L_1 \sim_A L_2$. A simple undirected binary graph $L \in \mathcal{L}$ is *vertex transitive* if for every given pair of nodes i and j in $\mathbf{N}$, there exists an automorphism $f : \mathbf{N} \to \mathbf{N}$ such that $f(i) = j$. Thus when a network is vertex transitive, we can take a node i and map it to the position of any other node j, by changing the label of j to i, and there exists a way of relabeling the other nodes such that all nodes have exactly the same neighbors as before and the structure of the graph is preserved. Thus the positions of any two nodes i and j in a vertex-transitive network are equivalent in a certain sense, and it is intuitive that the agents should receive the same payoff.

To formalize the idea that vertex transitivity is the key network symmetry condition for equal payoffs we show that for a large class of payoff functions mapping network positions into payoffs, payoffs are identical if and only if the network is vertex transitive. In principle, an agent's payoff can depend not only on their position in a network L, but also their position in subnetworks of L. Moreover, we might want to assign different subnetwork values to agents in the same subnetwork that vary with different orderings of the agents, and in particular, some notion of the marginal effect an agent has on the subnetwork. This gives us a rich basis for considering network payoffs.

Define $\mathcal{T}$ as the *ordered* set of permutations over the nodes $\mathbf{N}$. Note that with regards to any node $i \in \mathbf{N}$, every permutation $\tau \in \mathcal{T}$ maps one-to-one onto two specific induced subgraphs: one, the subgraph supported by nodes up to and *excluding* i, and two, the subgraph supported by nodes up to and *including* i. Let $\nu : \mathcal{T} \times \mathcal{V} \times \mathcal{L} \to \mathbb{R}$ be the function which assigns to every pair $\{\tau, i\}$ a "marginal value" with regards to such implied pairs of subgraphs in L. Let $S^i_\tau \subseteq L$ denote the induced subgraph supported by those nodes up to and including i in τ, while $S^{-i}_\tau = S^i_\tau \setminus \{i\}$ is the node-deleted subgraph of S^i_τ with regards to i. We require $\nu_i(\tau_k) = \nu_j(\tau_\ell)$ if the respective subgraphs including, respectively, i and j are isomorphic ($S^i_{\tau_k} \sim_I S^j_{\tau_\ell}$) *and* the respective subgraphs excluding, respectively, i and j are isomorphic ($S^{-i}_{\tau_k} \sim_I S^{-j}_{\tau_l}$). This confines node identity only to matter in so far as it corresponds to a position in a subnetwork. Let $V_i : \mathbb{R}^{n!} \to \mathbb{R}$ be the function which maps node i's multi-set of $n!$ marginal values $\{\nu_i(\tau)\}_T$ in L onto a "graph value", where $n!$ is simply the cardinality of $\mathcal{T}$. So the overall payoff we assign to an agent depends on all their possible marginal values. We restrict how these marginal values are mapped into payoffs by requiring only anonymity, i.e. any pre-image under V_i is closed under permutation. Hereafter, we let ν_i and V_i indicate the conditioning on some node i when convenient.

Let $\mathcal{V}$ denote the space of admissible functions V_i, and $\mathcal{F}$ denote the space of admissible functions ν_i. We will say that a result applies generically if it applies to all but a zero measure set of admissible functions ν_i and all but a zero measure set of admissible functions V_i. Of course, many non-generic mappings (V_i, ν_i) may be of interest. That non-withstanding it is of interest to study what network symmetry is needed in general for agents to receive equal payoffs.

We proceed to pin down the simple graphs for which expected payoffs V_i must be uniform for any payoff mapping (including the Myerson value).

Proposition SA10. *$V_i = V_j$ for all $i, j \in \mathbf{N}$ if L is vertex transitive and, generically, $V_i = V_j$ for all $i, j \in \mathbf{N}$ only if L is vertex transitive.*

Proof. **Sufficiency.** Consider any two nodes i and j as well as some permutation τ. Let $f(\tau) \in \mathcal{T}$ be the image of τ under the automorphism mapping i to j, which exists by the vertex transitivity of L. As automorphisms by definition preserve adjacency relations, $S_\tau^j \sim_I S_{f(\tau)}^{f(i)}$ and $S_\tau^{-j} \sim_I S_{f(\tau)}^{-f(i)}$ for all $i, j \in \mathcal{V}$. Hence, $\nu_j(\tau) = \nu_i(f(\tau))$ by the earlier requirement of identity in ν under isomorphism of the implied graph arguments, for any $\tau \in \mathcal{T}$. Fix the set of all of i's marginal values in L, written $\{\nu_i\}_\mathcal{T}$, in arbitrary order. By the foregoing argument, there exists a bijection between $\{\nu_i\}_\mathcal{T}$ and $\{\nu_j\}_\mathcal{T}$ through f. By anonymity of V, hence $V(\{\nu_j\}_\mathcal{T}) = V(\{\nu_i\}_{f\circ\mathcal{T}}) = V(\{\nu_i\}_\mathcal{T})$.

Necessity. We start with a well known result from graph theory. A simple undirected binary graph of finite order is vertex transitive if, and only if, its one-node deleted subgraphs are isomorphic (Thomassen, 1985). Thus if we have a graph that is not vertex transitive there exist nodes i and j such that $L \setminus \{i\} \not\sim_I L \setminus \{j\}$. Hence generically, for any permutation τ in which i is last, and any permutation τ' in which j is last, $\nu_i(\tau) \neq \nu_j(\tau)$. So generically, $V(\{\nu_i\}_\mathcal{T}) \neq V(\{\nu_j\}_\mathcal{T})$. □

Proposition SA10 shows that for a large space of payoff functions for which all that matters is agents' network positions, vertex transitivity guarantees equal payoffs and is also required for equal payoffs, generically. A non-generic payoff function in this space, that can be applied in the special case of transferable utilities (which is not assumed for the above result) is the Myerson value. As we have seen, vertex transitivity is sufficient but not necessary for equal payoffs under the Myerson value. With the Myerson value the weaker symmetry requirement of regular networks (so that each node has the same number of neighbors) is sufficient (see Section 4 of the main paper).

Proposition SA10 takes its informational basis for determining payoffs to be similar to that used by the Myerson value. However, this informational basis is very broad and any way of determining payoffs based on coarser information is covered by the result. For example, the result covers any payoff function that depends only on each agent's set of friends (neighbors), set of friends of friends, set of friends of friends of friends, and so on, for every possible subnetwork of L. As a more specific example, if payoffs were proportional to each agent's eigenvector centrality they would depend only on the structure of the network L, and so by Proposition SA10 agents would receive identical payoffs on a vertex transitive network.[6] Similarly, if payoffs were proportional to agents' marginal contributions to the spectral radius of the network L, then they would depend only on the structure of the network L and the

[6]As one of many ways in which this can be implemented, set $\nu_i(\tau)$ equal to i's eigenvector centrality on L for all τ in which i is last to arrive, and to 0 otherwise, and let V_i equal $\max_{\tau\in\mathcal{T}} \nu_i(\tau)$.

subnetworks $L \setminus \{i\}$ for all $i \in N$, and so, by Proposition SA10 agents' payoffs would be identical on a vertex transitive network.

D.0.2. *Efficiency.* A network $L = (n, L)$ is *Pareto efficient* if there is no network L' such that the payoffs of the agents on the network $L' = (n, L')$ Pareto dominate those on L (i.e., all agents receive weakly higher net payoffs on L' than L, and at least one agent receives a strictly higher payoff).

To get a handle on the set of *Pareto-efficient* networks, we assume that shorter path lengths facilitate weakly better risk sharing.

Assumption 19. *All Pareto-efficient networks $L = (n, L)$ have one component, and there is no alternative network $L' = (n, L')$ such that $|L'| \leq |L|$ and the path length distribution of L' first-order stochastically dominates the path length distribution of L.*

This enables us to eliminate some configurations as being Pareto efficient. We also make an assumption that risk-sharing relationships are sufficiently costly to maintain that dense risk-sharing networks are Pareto inefficient.

Assumption 20. *There are no Pareto efficient networks $L = (n, L)$ in which $|L| \geq \sqrt{n-1}$.*

To aid interpretation, a realistic lower bound on the size of a typical village is 100, while a realistic upper bound is 500. Thus when $n = 100$ this rules out Pareto-efficient risk-sharing networks with an average degree of more than about 10 links, while for $n = 500$ it rules out Pareto-efficient risk-sharing arrangements with an average degree of more than about 22 links. As a comparison, using the data collected by Banerjee, Chandrasekhar, Duflo and Jackson (2013) across 75 rural villages in southern India, the mean number of households in a village is 209 and the average number of risk-sharing relationships a household has is less than 4. Although there is no guarantee that the networks we observe in practice are efficient, or even close to efficient, the fact that risk-sharing networks are much sparser in practice than required by our upper bound is suggestive that the costs of forming links are sufficiently high to make Assumption 20 reasonable. Moreover, if there is no underinvestment in links in stable networks, as in our benchmark model, then the observed density of links is a valid upper bound for the density of links in Pareto-efficient networks.

D.0.3. *Stability.* Finally, we turn to stability. Since we want to make a point at a high level of generality, without a concrete model specification, we place no restrictions on stability. The efficient stable networks are of course constrained by the assumptions we've made on efficiency, and this tension with equality is sufficient for our impossibility result.

D.0.4. *A general tension.* The next result formalizes the general tension between efficient stable networks and equality by showing that, given the assumptions we've made, a network cannot be both Pareto efficient and regular—which, as argued above, is in general necessary but not sufficient for perfectly equitable outcomes.

Proposition 19. *Given Assumptions 19 and 20, there does not exist a Pareto-efficient and regular network.*

Proof. Towards a contradiction suppose there exists a regular *Pareto efficient* network $L = (n, L)$ of order r in which Assumptions 19 and 20 are satisfied.

By Assumption 19 all *Pareto efficient* networks have one component, and there is no alternative network L' such that $|L'| \leq |L|$ and the path length distribution of L' first order stochastically dominates the path length distribution of L. As L has one component, it contains at least $n-1$ links. Consider now a network $L' = (n, L')$ that has the same number of links as L and contains the star network on n nodes as a subnetwork. Such L contains at least $n-1$ links, such a network exists. The path length distribution for the network L' is that there are $n(n-1)/2 - |L|$ pairs of agents that have path length 2 and $|L|$ agents that path length 1. This is because nodes i and j have a minimum path length of 1 in L' if and only if $l'_{ij} \in L'$, while the remaining $n(n-1)/2 - |L|$ pairs of nodes must have a minimum path length of 2 because the star network is a subnetwork of L'. As L is Pareto efficient, it must have exactly the same path length distribution. It will also have $|L|$ minimum path lengths of 1 and to prevent its path length distribution from being first order stochastically dominated by L' the remaining minimum path lengths must be 2. Hence L has a diameter of 2.

In a regular network of order r, each node has r neighbors. A given node i then has r neighbors, and each of these have $r-1$ neighbors other than i. Were there any other nodes in the graph, the diameter would be more than 2. Hence, an upper bound on the number of nodes in the network L is $\bar{n} := 1 + r + r(r-1) = r^2 + 1$. Thus $n \leq r^2 + 1$. Rearranging this inequality, the order of L must be $r \geq \sqrt{n-1}$. However, this violates Assumption 20, so we have a contradiction. □

E. Timing of Negotiations

Here we provide a more detailed treatment of the analysis in Section 6.6 of the main paper. We begin by formally introducing a version of our basic model in which establishing links and negotiating risk-sharing agreements using these links happen at the same time. We refer to this model as the *one-stage negotiations model.* Let the socio-economic environment be as it is defined in Subsection 2.1 of the main paper. But now we consider the following network formation game ex ante (before the realization of endowments), in which agents simultaneously decide which other agents they would like to establish a link with, and what risk-sharing agreement to propose to these agents. Formally, each agent i selects a subset of agents $A(i) \subseteq N/\{i\}$ to approach, and for each $j \in N/\{i\}$ proposes a collection of risk-sharing agreements $\tau_{ij}(L, \omega)$ for every network L such that $l_{ij} \in L$, and for every $\omega \in \Omega$. Thus, a proposal specifies a proposed risk-sharing agreement with j for any possible network that can form. A link l_{ij} forms if and only if (i) $j \in A(i)$ and $i \in A(j)$; (ii) $\tau_{ij}(L, \omega) = -\tau_{ji}(L, \omega)$ for every network L such that $l_{ij} \in L$, and for every $\omega \in \Omega$. In words, the link is established if both agents approach each other and propose the same network-contingent and state-contingent transfer agreement with each other. As before, if a link is formed, both agents pay the same cost towards forming it. Successfully formed links comprise the realized network L, and after endowment realizations neighboring agents on L carry out the transfers specified by their agreements for network L and the particular endowment realization.

Pairwise stability can be extended in a straightforward manner to this more complicated network formation game. Let $\tau = (\tau_i)_{i \in N}$ denote the collections of proposals of different agents, and let $L(\tau)$ be the network established by τ.

We say that the collection of proposals τ and the resulting network $L(\tau)$ are *pairwise agreement stable* if and only if:

(i) there is no $i, j \in N$ with $l_{ij} \in L(\tau)$ such that i would strictly benefit from deviating from $A(i)$ to $A(i)/\{j\}$ while keeping the proposed τ_{ik} unchanged for every $k \in A(i)/\{j\}$;

(ii) there are no two agents i and j and transfer agreement τ'_{ij} who would both strictly benefit from a joint deviation of i approaching $A(i) \cup \{j\}$ and proposing an agreement with j of τ'_{ij} while keeping the proposal to all agents in $A(i)/\{j\}$ unchanged, and j approaching $A(j) \cup \{i\}$ and proposing an agreement with i of $-\tau'_{ji}$ while keeping the proposal to all agents in $A(j)/\{i\}$ unchanged.[7]

The next result establishes that networks and surplus divisions that are pairwise stable in our two-stage base model continue to be pairwise agreement stable in the one-stage model defined above. For every network L, let $\tau^*(L)$ be a collection of Pareto-efficient bilateral transfer agreements between neighboring agents in L that results in division of the total surplus according to the Myerson values of agents.

Proposition SA11. *Let L^* be a pairwise stable network in our baseline model in which the total surplus from risk sharing is divided according to the Myerson values of agents. Let the resulting vector of ex ante expected payoffs be $\underline{u}$. Then in the one-stage negotiations model there is a pairwise agreement stable collection of proposals τ with resulting network $L(\tau) = L^*$ such that the resulting ex ante expected payoffs of agents are $\underline{u}$.*

Proof. For every agent i, consider the strategy $A(i) = (j | l_{ij} \in L^*)$ and $\tau_{ij} = \tau^*_{ij}(L)$ for every $j \in A(i)$ and all L. Then the resulting network is L^* and the resulting ex ante payoffs are $\underline{u}$ by construction. We need to show that τ with the resulting network L^* is pairwise agreement stable. Condition (i) in the definition of pairwise agreement stability holds because given the above strategies, i not strictly benefitting from deviating from $A(i)$ to $A(i)/\{j\}$ while keeping the proposed τ_{ik} unchanged for every $k \in A(i)/\{j\}$ is equivalent to i not strictly benefitting in the first stage of the original model from cutting the link with j while keeping all other links. The latter holds because we assumed L^* to be pairwise stable in the original model.

[7]Note that this formulation allows both for $l_{ij} \in L(\tau)$ and $l_{ij} \notin L(\tau)$. In the former case the deviation is a renegotiation of the agreement between the two while in the latter case it is forming a new link with a new agreement.

We now turn to showing that condition (ii) in the definition of pairwise agreement stability also holds, for every $i, j \in N$. Consider first a pair i and j such that $l_{ij} \in L^*$. Since τ induces a Pareto-efficient risk-sharing arrangement on L^*, given all other bilateral risk-sharing agreements on L^* specified by τ, there is no transfer agreement τ'_{ij} between i and j that strictly benefits both of them relative to τ_{ij}. Consider now a pair i and j such that $l_{ij} \notin L^*$. In the original model, given L^*, a joint deviation by i and j to form a link is not strictly profitable for both of them, because L^* is pairwise stable. In the original model i and j divide the total extra surplus from established network $L^* \cup l_{ij}$ versus L^* equally as well as the cost of forming l_{ij}. By construction, the maximum total surplus they can achieve in the one-stage model, given the above strategies, when forming the link l_{ij} is the same as in the original model. Therefore i and j have no joint deviation in the one-stage model involving establishing l_{ij} with any transfer agreement between them, while keeping all transfer agreements with other agents unchanged, that strictly benefits both of them. $\square$

F. Overinvestment and Underinvestment Examples

In this Section we provide an example of over-investment within group in the unique stable network and a related example of underinvestment across group in the unique stable network.

We begin by assuming there is one group with s members connected by a network L. Equation 11 in the main paper implies that Myerson distance of two agents i, j such that $l_{ij} \notin L$ is greater than $1/2$, while the Myerson distanace between i and j if they form the link l_{ij} would be $1/2$. Thus i and j's gross payoff strictly increases if the link l_{ij} is added. So, for κ_w sufficiently close to 0, in all stable networks for any pair of agents i, j the link l_{ij} must be formed; The unique stable network is the complete network and there is overinvestment.

Suppose now there are two groups, g, g' both with s members and keep the same parameter values from the previous example. By equation 13 in the main paper, the incentives to form within group links are weakly increased by the presence of any across group links. Thus in all stable networks the network structure within-group must be complete networks; All possible within-group links must be formed. Suppose these are the only links formed so that no across-group links are formed. Denote this network L. From equation 13 the change in total variance achieved by connecting an agent i from group g to an agent j from group g' is strictly increasing in s (the size of both groups). Given the Myerson value calculation, this means that the marginal contribution of the link l_{ij} to total surplus (the certainty equivalent value of the variance reduction) is strictly greater on $L \cup \{l_{ij}\}$ than it is on any strict subgraph, including all those formed when the later of i and j arrives in the Myerson calculation. This implies that $(MV(i; L \cup l_{ij}) - MV(i; L)) + (MV(j; L \cup l_{ij}) - MV(j; L)) < TS(L \cup l_{ij}) - TS(L)$ for all $l_{ij} : i \in \mathbf{S}_g, j \in \mathbf{S}_{g'}$. So, setting κ_a such that

$$MV(i; L \cup l_{ij}) - MV(i; L) + MV(j; L \cup l_{ij}) - MV(j; L) < 2\kappa_a < TS(L \cup l_{ij}) - TS(L),$$

the network L is the unique stable network and there is underinvestment (in across-group links) in all stable networks.

References

Ali, N. and D. Miller (2013): "Enforcing cooperation in networked societies," mimeo UC San Diego.

Ali, N. and D. Miller (2016): "Ostracism and forgiveness,"The American Economic Review, 106.8 (2016): 2329-2348.

Ambrus, A., W. Y. Gao and P. Milan (2016): "Informal Risk Sharing with Local Information," Mimeo.

Ambrus, A., M. Mobius, and A. Szeidl (2014): "Consumption risk-sharing in social networks," American Economic Review, 104, 149-182.

Atkinson, A. B. (1970): "On the Measurement of Inequality," Journal of Economic Theory, 2, 244-263.

Coleman, J. (1988): "Social capital in the creation of human capital."American journal of sociology 94 (1988): S95-S120.

Jackson, M., T. Rodriguez-Barraquer and X. Tan (2012): "Social capital and social quilts: Network patterns of favor exchange," The American Economic Review 102.5 (2012): 1857-1897.

Nava, F., and M. Piccione, (2014): "Efficiency in repeated games with local interaction and uncertain local monitoring,"Theoretical Economics 9.1 (2014): 279-312.

Thomassen, C. (1987): "A characterization of locally finite vertex-transitive graphs,"Journal of Combinatorial Theory, Series B 43.1 (1987): 116-119.

Wolitzky, A. (2012): "Cooperation with network monitoring,"Review of Economic Studies, 80.1 (2012): 395-427.