

SOCIAL INVESTMENTS, INFORMAL RISK SHARING, AND INEQUALITY

ATTILA AMBRUS[†], ARUN G. CHANDRASEKHAR[‡], AND MATT ELLIOTT^{*}

ABSTRACT. This paper studies the formation of risk-sharing networks through costly social investments. First, individuals invest in relationships to form a network. Next, neighboring agents negotiate risk-sharing arrangements, in a generalized version of the model in Stole and Zwiebel (1996). This results in the social surplus being allocated according to the Myerson value. In particular, more centrally connected individuals receive higher shares. We find a novel trade-off between efficiency and equality. The most stable efficient network, which minimizes incentives to overinvest, also generates the most inequality. When individuals are split into groups and relationships across groups are more costly but incomes across groups are less correlated, there is never underinvestment into social connections within group, but underinvestment across groups is possible. More central agents have better incentives to form across-group links, reaffirming the efficiency inequality trade-off. Evidence from 75 Indian village networks is congruent with our model.

Date: March 16, 2015.

We thank the National Science Foundation, SES-1155302, for funding the research that collected the data we use. Ambrus, Chandrasekhar and Elliott acknowledge financial support from National Science Foundation (NSF) grant “The economic benefits of investing into social relationships,” SES-1429959. We thank Nageeb Ali, Ben Golub, Matt Jackson, Willemien Kets, Cynthia Kinnan, Rachel Kranton, Peter Landry, Horacio Larreguy, Melanie Morten, Kaivan Munshi and Luigi Pistaferri for helpful comments. We also thank seminar participants at the Calvo-Armengol prize conference, the Harvard-MIT theory seminar, Oxford, Cambridge, Warwick, Microsoft Research, 2012 ThRed conference in Budapest, Syracuse, Ohio State University, Duke, the Tinbergen Institute and the Santa Fe Institute.

[†]Department of Economics, Duke University; email: attila.ambrus@duke.edu.

[‡]Department of Economics, Stanford University, and NBER; email: arungc@stanford.edu.

^{*}Division of Humanities and Social Sciences, California Institute of Technology; email: mel-
liott@hss.caltech.edu.

1. INTRODUCTION

In the context of missing formal insurance markets and limited access to lending and borrowing, incomes may be smoothed through informal risk-sharing agreements that utilize social connections. A large theoretical and empirical literature studies this topic,¹ but considerably less attention has been paid to the social investments that enable risk sharing, how risk-sharing networks form in the presence of these costly investments, and the resulting implications for efficiency and equity.² At the same time, there is growing empirical evidence that risk-sharing networks respond to financial incentives,³ and that in general risk-sharing networks form endogenously, in a way that depends on the economic environment. A central question we address is whether such investments are efficient, and if not then whether too many or too few resources are allocated to maintaining relationships. We are also interested in whether the resulting equilibrium forces generate social inequality (asymmetries in network positions) in society, even when agents are ex-ante homogeneous, and if social inequality translates into financial inequality.

Both underinvestment and overinvestment in social capital are conceivable. Two people establishing a social connection to share risk gain access to a less stochastic income stream which might generate improved opportunities to share risk with their other connections. As these positive spillovers might not be fully taken into account when deciding whether to establish the link, underinvestment can prevail. On the other hand, if more socially connected individuals receive a higher share of the surplus generated by risk sharing, that can lead to overinvestment. Villagers may form links to redistribute the surplus towards themselves, rather than to increase the overall surplus generated. The empirical literature also suggests that both types of inefficiencies are possible, in different contexts. Austen-Smith and Fryer (2005) cites numerous references from sociology and anthropology, suggesting that members of poor communities allocate inefficiently large amounts of time to activities maintaining social ties, instead of productive activities. In contrast, Feigenberg et al. (2013) find evidence in a microfinance setting that it is relatively easy to experimentally intervene and create social ties among people that yield substantial benefits. One explanation for this finding is that there is underinvestment in social relationships.

It is important to study whether there is too little or too much investment into social relations, both to put related academic work (which often takes social connections to be

¹An incomplete list of papers includes Rosenzweig (1988), Fafchamps (1992), Coate and Ravallion (1993), Townsend (1994), Udry (1994), Ligon, Thomas and Worrall (2002), Fafchamps and Gubert (2007), Bloch, Genicot, and Ray (2008), Angelucci and di Giorgi (2009), Jackson, Rodriguez-Barraquer and Tan (2012), Ambrus, Mobius and Szeidl (2014).

²Previous works that do consider the network formation problem include Bramoullé and Kranton (2007a,b) in the theoretical literature and Attanasio et al. (2012) in the experimental literature. For a related paper outside the networks framework, see Glaeser et al. (2002).

³See recent work by Binzel et al. (2014) and Banerjee et al. (2014b,c), which in different contexts look at how social networks respond to the introduction of financial instruments such as savings vehicles or microfinance.

exogenously given) into context and to guide policy choices. Consider the example of microfinance. If there is overinvestment, microfinance has a greater scope for efficiency savings in terms of reducing people's allocation of time into social investments. With underinvestment, however, it has more scope for smoothing incomes. If there is neither under- nor overinvestment, it also tells us that informal risk sharing is working relatively well as a second-best solution. Understanding which regime applies can help anticipate policy implications and evaluate welfare impacts of interventions.

To explore efficiency and inequality, in this paper we consider a two-stage model of network formation and risk sharing, in a context in which agents with constant absolute risk aversion (CARA) utilities face uncertain income realizations. In the first stage, agents choose with whom to form connections. Link formation is costly, as in Myerson (1991) and Jackson and Wolinsky (1996). In the second stage, agents who have formed a connection commit to a bilateral risk-sharing agreement (transfers contingent on income realizations).⁴

To keep the model tractable, we abstract away from the issues of enforcement of risk-sharing agreements and assume that agents face no constraints in that dimension.⁵ We investigate agreements satisfying two simple properties. First, we require agreements to be pairwise efficient, in that no pair of directly connected agents leave gains from trade on the table. Second, following Stole and Zwiebel (1996), we require the agreements to be robust to "split-the-difference" renegotiations.⁶ An agreement between two agents must evenly split the surplus generated by their link over and above the expected payoffs they would receive in its absence. This can be thought of as the transfers being robust to renegotiation if renegotiation would result in the two agents benefitting equally from the link.⁷ From this exercise we obtain a recursive definition of how surpluses get divided on different networks. Applying the axiomatization of the Myerson value provided in Myerson (1980), we show that the unique division of surplus compatible with our two requirements is the Myerson

⁴Although we consider a model in which there is perfect risk sharing of income, we could easily extend the model so that some income is perfectly observed, some income is private, and there is perfect risk sharing of observable income and no risk sharing of unobservable income. This would be consistent with the theoretical predictions of Cole and Kocherlakota (2001) and the empirical findings of Kinnan (2011). In the CARA utilities setting, such unobserved income outside the scope of the risk-sharing arrangement does not affect our results.

⁵See, for example, Ambrus et al. (2014) for an investigation of such issues.

⁶Stole and Zwiebel (1996) model bargaining between many employees and an employer. This scenario can be represented by a star network with the employer at the center.

⁷The process is decentralized, as it envisages pairwise renegotiations. The result in Stole and Zwiebel (1996) that we extend is Theorem 1. Their Theorem 2 can also be extended to our setting, and this would provide fully noncooperative foundations. Indeed, related noncooperative foundations are provided by Fontenay and Gans (2013), while Navarro and Perea (2013) take a different approach to microfounding the Myerson value. Slikker (2007) also provides noncooperative foundations, although the game analyzed is not decentralized: offers are made at the coalitional level.

value,⁸ a network-specific version of the Shapley value.⁹ The transfers required to implement the agreements we identify are particularly simple. Each agent receives an equal share of aggregate realized income (as in Bramoullé and Kranton, 2007a) and then state independent transfers are made.

Although we derive the prediction of agents sharing the surplus according to the Myerson value using a decentralized approach, such a split of surplus might also be appropriate when negotiations are centralized. In such situations the Myerson value has normative appeal on the grounds of fairness: two agents benefit equally from an investment they make, and receive benefits proportional to their average marginal contributions to total surplus (from establishing costly links).¹⁰

A key implication of the Myerson value determining the division of surplus is that more centrally connected agents receive a higher share of the surplus. Moreover, in our risk-sharing context it implies that agents receive larger payoffs from providing “bridging links” to otherwise socially distant agents than from providing local connections.¹¹ Empirical evidence supports this feature of our model—see Goyal and Vega-Redondo (2007), and references therein from the organizational literature: Burt (1992), Podolny and Baron (1997), Ahuja (2000), and Mehra et al. (2001).

In the network formation stage, we study the set of pairwise-stable networks (Jackson and Wolinsky, 1996).¹²

Our general analysis considers a community comprised of different groups where all agents within each group are ex-ante identical, and establishing links within groups is cheaper than across groups. We also assume that the income realizations of agents within groups are more positively correlated than across groups. Groups can represent different ethnic groups or castes in a given village, or in different villages. We find that there can be overinvestment within groups but not underinvestment, whereas across groups underinvestment is likely to be the main concern.

To see the intuition about overinvestment within groups, we first consider the case of homogeneous agents, that is, when there is only one group. Using the inclusion–exclusion principle from combinatorics,¹³ we provide a complete characterization of stable networks.

⁸The Myerson value is defined for a transferable utility environment. However, agreements are made before income realizations and determine expected utilities, and in our CARA setting, expected utilities are transferable through state-independent transfers.

⁹For investigations of the division of surplus in social networks in other contexts, see Calvo-Armengol (2001, 2003), Corominas-Bosch (2004), Manea (2011), and Kets et al. (2011).

¹⁰These motivations make the Myerson value a commonly used concept in the network formation literature. See a related discussion on pp. 422–425 of Jackson (2010).

¹¹More precisely, in Section 4 we introduce the concept of Myerson distance to capture the social distance between agents in the network, and show that a pair of agents’ payoffs from forming a relationship are increasing in this measure.

¹²Results from Calvo-Armengol and Ilkilic (2009) imply that under some parameter restrictions—for example when agents are ex ante identical—the set of pairwise-stable outcomes is equivalent to the (in general more restrictive) set of pairwise Nash equilibrium outcomes.

¹³See Chapter 10 in van Lint and Wilson (2001).

We show that in this case there can never be underinvestment in social connections, as agents establishing an essential link (connecting two otherwise unconnected components of the network) always receive a benefit exactly equal to the social value of the link. However, overinvestment, in the form of redundant links, is possible, and becomes widespread as the cost of link formation decreases.

We find a trade-off between efficiency and equality. Among all possible efficient network structures, we find that the most stable (in the sense of being stable for the largest set of parameter values) is the star, which also results in the most unequal division of surplus. The intuition is that the star network minimizes the incentives of peripheral agents to establish redundant links. Conversely, the least stable efficient network entails the most equal division of surplus among all efficient networks. Although agents are ex-ante identical, efficiency considerations push the structure of social connections towards asymmetric outcomes that elevate certain individuals. Socially central individuals emerge endogenously from risk-sharing considerations alone.

Turning attention to the case of multiple groups, we find that across-group underinvestment becomes an issue when the cost of maintaining links across groups is sufficiently high.¹⁴ The reason is that the agents who establish the first connection across groups receive less than the social surplus generated by the link, providing positive externalities for peers in their groups. This gap between private and social benefits is smaller for agents located more centrally in their own group, providing a second force for some agents within a group to be more central. For example, with two groups, we show that the most stable efficient network structure involves stars within groups, connected by their centers. This reinforces the trade-off between efficiency and equality in the many-groups context.

Using data from 75 Indian villages, we provide some supporting evidence for our model. We split the villagers into two groups, by caste.¹⁵ From the theoretical analysis, risk-sharing links are most valuable when they bridge otherwise unconnected components. And when a link does not provide such a bridge, its value depends on how far apart, suitably defined, the agents would otherwise be on the social network. We call this distance between a pair of agents their Myerson distance. Our theory predicts that there is an upper bound on the Myerson distance between any two unconnected agents within the same group, beyond which the pair of agents would have a profitable deviation by forming a link. In addition, we predict that there will be inequality in social positions and that more central agents within their group will form across-group links. However, there are many alternative stories consistent with these predictions. We therefore also generate more subtle predictions to consider: how changes in the economic environment in terms of income variability or correlation correspond to changes

¹⁴While across-group overinvestment remains possible, the main concern when across-group link costs are relatively high is underinvestment.

¹⁵There is an extensive literature that examines caste as a main social unit where risk sharing takes place (Townsend (1994), Munshi and Rosenzweig (2009), and Mazzocco and Saini (2012)).

in network structure under our model, while differencing out the friendship network from the risk-sharing network.

Consider two villagers from the same caste. As income variability increases or within-caste incomes become less correlated, all else equal, the value of a risk-sharing link between these villagers increases; the Myerson distances that can be observed in a stable network therefore decrease, which in turn implies the following: (i) income variability is positively associated with lower Myerson distances between unconnected agents, and (ii) within-caste income correlation is associated with higher Myerson distances. The theory also predicts that villagers have to be sufficiently central within their own caste (the threshold depending again on income variability and within- versus across-caste income correlation) to be incentivized to provide a risk-sharing link across castes. This yields our final predictions: (iii) in villages with more income variability, more agents will have sufficient incentives to form an across-caste link, and so the association between within-caste centrality and who provides across-caste links is weaker, and (iv) in villages with more within-caste income correlation relative to across-caste income correlation, more agents will again have sufficient incentives to form across-caste links, and so the association between within-caste centrality and having an across-caste link will again be weaker. Because working with the exact Myerson distances is computationally infeasible, we develop an approximation which is exact for tree graphs, and also check that our results are robust to other notions of network sparsity. We demonstrate that our predictions are borne out in our data.

To strengthen our results, we exploit the fact that we have multigraph data. Not only do we have complete financial network data for every household in every village, but we have complete friendship network data as well. As our theory pertains only to the financial network, we are able to take a difference-in-differences approach. For example, for predictions (i) and (ii), we look within villages, across network type, and ask whether the association with economic environmental parameters (income variability and within-caste income correlation) differentially vary with the Myerson distance of the financial network compared to that of the social network. This allows us to take out arbitrary village-level fixed effects.

Ultimately, our empirical approach allows us to be more conservative than similar studies in the literature (e.g., Karlan et al. (2009), Ambrus et al. (2014), and Kinnan and Townsend (2012)). While the studies above have access to just a few of networks (e.g., 2, 1 and 16, respectively), we have 75 networks and also multigraph data. Most studies, therefore, are forced to do statistical inference within networks, which limits the number degree of correlated shocks they are able to handle. Relative to this approach, our focus on the village level and differencing out the social network is extremely conservative.

On the theory side, the studies on social networks and informal risk sharing that are most related to ours include Bramoullé and Kranton (2007a,b), Bloch et al. (2008), Jackson et al. (2012), Billand et al. (2012), Ali and Miller (2013a,b), and Ambrus et al. (2014). Many of these papers focus on the enforcement issues we abstract from and investigate how social

capital can be used to sustain cooperation for lower discount factors than would otherwise be possible. We take a complementary approach and instead focus on the distribution of surplus and the incentives this creates for social investments. One way of viewing our approach to enforcement is that we instead consider the part of the parameter range where the discount factor is sufficiently high to sustain cooperation without leveraging social capital. Among the aforementioned papers, Bramoullé and Kranton (2007a,b) and Billand et al. (2012) investigate costly network formation. Bramoullé and Kranton’s (2007a,b) model assumes that the surplus on a connected income component is equally distributed, independently of the network structure. This rules out the possibility of overinvestment, and leads to different types of stable networks than in our model. Instead of assuming optimal risk-sharing arrangements, Billand et al. (2012) assume an exogenously given social norm, which prescribes that high-income agents transfer a fixed amount of resources to all low-income neighbors. This again leads to very different predictions regarding the types of networks that form in equilibrium.

More generally, network formation problems are important. Establishing and maintaining social connections (relationships) is costly, in terms of time and other resources. However, on top of direct consumption utility, such links can yield many economic benefits. Papers studying formation in different contexts include Jackson and Wolinsky (1996), Bala and Goyal (2000), Kranton and Minehart (2001), Hojman and Szeidl (2008), and Elliott (2014). Notable in this literature is a lack of empirical work, which can be attributed to a number of innate difficulties that taking these models to data presents. One common problem is a multiplicity of stable networks. But perhaps most important is that the networks in question can only very rarely be partitioned into a sizeable number of separate networks that can reasonably be treated as independent. Since predictions are often at the level of the overall network structure, this makes testing extremely challenging. Our many observations of social networks that are relatively independent of each other, coupled with our approach to circumventing data limitations, allow us to provide a first step towards testing predictions based on the overall network structure. And although we study a specific network formation problem tailored to risk sharing in villages, the general structure of our problem is relevant to other applications.¹⁶

The remainder of the paper is organized as follows. Section 2 describes risk sharing on a fixed network. In Section 3 we introduce a game of network formation with costly link formation. We focus on network formation within a single group in Section 4 and then turn to the formation of across-group links in Section 5. We then generate comparative static

¹⁶For a different and more specific application, suppose researchers can collaborate on a project. Each researcher brings something heterogeneous and positive to the value of the collaboration, so that the value of the collaboration is increasing in the set of agents involved. Collaboration is possible only when it takes place among agents who are directly connected to another collaborator and surplus is split according to the Myerson value (as in our work, motivated by robustness to renegotiations). Such a setting fits into our framework.

results in Section 6, which we use to take our model to the data in Section 7. Section 8 concludes.

2. PRELIMINARIES AND RISK SHARING ON A FIXED NETWORK

To study social investments and the network formation problem, first we need to specify what risk-sharing arrangements take place once the network is formed. Below we introduce an economy in which agents face random income realizations, introduce some basic network terminology, and discuss risk-sharing arrangements for a given network.

2.1. The socio-economic environment. We denote the set of agents in our model by \mathbf{N} , and assume that they are partitioned into a set of groups \mathbf{M} . We let $G : \mathbf{N} \rightarrow \mathbf{M}$ be a function that assigns each agent to a group; i.e., if $G(i) = g$ then agent i is in group g . One interpretation of the group partitioning is that \mathbf{N} represents individuals in a village, and the groups correspond to different castes. Another possible interpretation is that \mathbf{N} represents individuals in a larger geographic region (such as a district or subdistrict), and groups correspond to different villages in the region.

Agents in \mathbf{N} face uncertain income realizations. For tractability, we assume that incomes are jointly normally distributed, with expected value μ and variance σ^2 for each agent.¹⁷ We assume that the correlation coefficient between the incomes of any two agents within the same group is ρ_w , while between the incomes of any two agents not in the same group it is $\rho_a < \rho_w$.¹⁸ That is, we assume that incomes are more positively correlated within groups than across groups, so that all else equal, social connections across groups have a higher potential for risk sharing.

Although we introduce the possibility of correlated incomes in a fairly stylized way, our paper is one of the first to permit differently correlated incomes between different pairs of agents. Such correlations are central to the effectiveness of risk-sharing arrangements, as shown below.

We refer to possible realizations of the vector of incomes as *states*, and denote a generic state by ω . We let $y_i(\omega)$ denote the income realization of agent i in state ω .

Agents can redistribute realized incomes; hence their consumption levels can differ from their realized incomes. We assume that all agents have constant absolute risk aversion (CARA) utility functions:

$$(1) \quad v(c_i) = -\frac{1}{\lambda} e^{-\lambda c_i},$$

where c_i is agent i 's consumption and $\lambda > 0$ is the coefficient of absolute risk aversion. The assumption of CARA utilities, together with jointly normally distributed incomes, greatly

¹⁷This specification implies that we cannot impose a lower bound on the set of feasible consumption levels. As we show below, our framework readily generalizes to arbitrary income distributions, but the assumption of normally distributed shocks simplifies the analysis considerably.

¹⁸It is well-known that for a vector of random variables, not all combinations of correlations are possible. We implicitly assume that our parameters are such that the resulting correlation matrix is positive semidefinite.

enhances the tractability of our model: as we show below it leads to a transferable utility environment in which the implemented risk-sharing arrangements are relatively simple. This utility formulation can also be considered a theoretical benchmark case with no income effects. The empirical relevance of predictions from this benchmark model is examined in Section 6.

2.2. Basic network terminology. Before proceeding, we introduce some standard terminology from network theory. A social network L is an undirected graph, with nodes \mathbf{N} corresponding to the different agents, and links representing social connections. Abusing notation we also let L denote the set of links in the network. We will refer to the agents linked to agent i , $\mathbf{N}(i; L) := \{j : l_{ij} \in L\} \subset \mathbf{N}$, as i 's *neighbors*. The *degree centrality* of an agent is simply the number of neighbors she has (i.e., the cardinality of $\mathbf{N}(i; L)$). An agent's neighbors can be partitioned according to the groups they belong to. Let $\mathbf{N}_g(i; L)$ be i 's neighbors on network L from group g . A *path* is a sequence of agents $\{i, k, k', \dots, k'', j\}$ such that every pair of adjacent agents in the sequence is linked. A *minimal path* between i and j is a path where i is the first agent in the sequence, j is the last and there is no other path between i and j that is a subsequence of that path. The *path length* of a path is the number of agents in the path.

We will sometimes refer to subsets of agents $\mathbf{S} \subseteq \mathbf{N}$ and denote the subgraphs they generate by $L(\mathbf{S}) := \{l_{ij} \in L : i, j \in \mathbf{S}\}$. A subset of agents $\mathbf{S} \subseteq \mathbf{N}$ is *path connected* on L if, for each $i \in \mathbf{S}$ and each $j \in \mathbf{S}$, there exists a path connecting i and j . For any network there is a unique partition of \mathbf{N} such that there are no links between agents in different partitions but all agents within a partition are path connected. We refer to the cells of this partition as *network components*. A *shortest path* between two path-connected agents i and j is a path connecting i and j with a lower path length than any other. The *diameter* of a network component $C \subset L$ is $d(C)$, the maximum value—taken over all pairs of agents in C —of the length of a shortest path. A network component is a *tree* when there is a unique minimal path between any two agents in the component. A *line network* is the unique (tree) network, up to a relabeling of agents, in which there is a path from one (end) agent to the other (end) agent that passes through all other agents exactly once. A *star network* is the unique tree network, up to a relabeling of agents, in which one (center) agent is connected to all other agents.

2.3. Risk-Sharing Agreements. We assume that income cannot be directly shared between agents $i, j \in \mathbf{N}$ unless they are connected, i.e., $l_{ij} \in L$. However, we let the state ω (and hence every agent's income realization) be publicly observed so agents can make transfer arrangements contingent on it. We consider this environment with perfectly observable incomes as a benchmark model, which is a relatively good description of village societies in which people closely monitor each other. It is also straightforward to extend the model so

that some income is publicly observed (and shared) while the remaining income is privately observed (and never shared). Results are very similar for this more general setting.¹⁹

Formally, a risk-sharing agreement on a network L specifies transfer $t_{ij}(\omega, L) = -t_{ji}(\omega, L)$ between neighboring agents i and j for every possible state ω . Abusing notation where there should be no confusion we sometimes drop the second argument and write $t_{ij}(\omega)$ instead of $t_{ij}(\omega, L)$. The interpretation is that in state ω agent i is supposed to transfer $t_{ij}(\omega)$ units of consumption to agent j if $t_{ij}(\omega) > 0$, and receives this amount from agent j if $t_{ij}(\omega) < 0$. Given a transfer arrangement between neighboring agents, agent i 's consumption in state ω is $c_i(\omega) = y_i(\omega) - \sum_{j \in \mathbf{N}(i)} t_{ij}(\omega)$. It is straightforward to show that state-contingent consumption plans $(c_i(\cdot))_{i \in \mathbf{N}}$ are feasible, that is they can be achieved by bilateral transfers between neighboring agents, if and only if $\sum_{i \in C} c_i(\omega) = \sum_{i \in C} y_i(\omega)$ for every state ω and network component C .

A basic assumption we make in our model is that given all other risk-sharing arrangements, an agreement reached by linked agents i and j must leave no gains from trade on the table. There must be no other agreement that can make both i and j strictly better off holding fixed the agreements of other players. We call such transfers *pairwise efficient*.²⁰

By the well-known Borch rule (see Borch (1962), Wilson (1968)) a necessary and sufficient condition for this property is that for all neighboring agents i and j ,

$$(2) \quad \left(\frac{\partial v_i(c_i(\omega))}{\partial c_i(\omega)} \right) / \left(\frac{\partial v_j(c_j(\omega))}{\partial c_j(\omega)} \right) = \left(\frac{\partial v_i(c_i(\omega'))}{\partial c_i(\omega')} \right) / \left(\frac{\partial v_j(c_j(\omega'))}{\partial c_j(\omega')} \right)$$

for every pair of states ω and ω' . But if this holds for all neighboring agents i and j then the same condition must hold for all pairs of agents on a component of L , independently of whether they are directly or indirectly connected. The results in Borch (1962) and Wilson (1968) imply that the risk-sharing agreement is Pareto-efficient at the component level. For this reason, below we establish some important properties of Pareto-efficient risk-sharing arrangements on components.

Proposition 1 shows that the CARA utilities framework has the convenient property that expected utilities are transferable, in the sense defined by Bergstrom and Varian (1985). This can be used to show that ex-ante Pareto efficiency is equivalent to minimizing the sum of the variances, and it is achieved by agreements that at every state split the sum of the incomes on each network component equally among the members and then adjust these shares by state-independent transfers. The latter determine the division of the surplus created by the

¹⁹Kinnan (2011) finds evidence that hidden income can explain imperfect risk sharing in Thai villages relative to the enforceability and moral hazard problems we are abstracting from. Cole and Kocherlakota (2001) show that when individuals can privately store income, state-contingent transfers are not possible and risk sharing is limited to borrowing and lending.

²⁰More formally, transfers $\{t_{ij}(\omega, L)\}_{\omega \in \Omega, ij: l_{ij} \in L}$ are pairwise efficient for a network L if there is no pair of agents $ij: l_{ij} \in L$ and no alternative transfers $\{t'_{ij}(\omega, L)\}_{\omega \in \Omega, ij: l_{ij} \in L}$ such that $t'_{kl}(\omega, L) = t_{kl}(\omega, L)$ for all $kl \neq ij$ and all $\omega \in \Omega$, that gives both i and j strictly higher expected utility.

risk sharing agreement. We emphasize that this result does not require any assumption on the distribution of incomes, only that agents have CARA utilities.

Proposition 1. *For CARA utility functions certainty-equivalent units of consumption are transferable across agents, and if $L(\mathbf{S})$ is a network component, the Pareto frontier of ex-ante risk-sharing agreements among agents in \mathbf{S} is represented by a simplex in the space of certainty-equivalent consumption. The ex-ante Pareto-efficient risk-sharing agreements for agents in \mathbf{S} are those that satisfy*

$$\min \sum_{i \in \mathbf{S}} \text{Var}(c_i) \quad \text{subject to} \quad \sum_{i \in \mathbf{S}} c_i(\omega) = \sum_{i \in \mathbf{S}} y_i(\omega) \quad \text{for every state } \omega,$$

and they are comprised of agreements of the form

$$c_i(\omega) = \frac{1}{|\mathbf{S}|} \sum_{k \in \mathbf{S}} y_k(\omega) + \tau_i \quad \text{for every } i \in \mathbf{S} \text{ and state } \omega,$$

where $\tau_i \in \mathbb{R}$ is a state independent transfer.

The proof of Proposition 1 is in Appendix B. Proposition 1 implies that the total surplus generated by efficient risk-sharing arrangements is an increasing function of the reduction in aggregate consumption variance (the sum of consumption variances). For a general distribution of shocks, this function can be complicated. However, if shocks are jointly normally distributed then $c_i = \frac{1}{|\mathbf{S}|} \sum_{k \in \mathbf{S}} y_k + \tau_i$ is also normally distributed, and $E(v(c_i)) = E(c_i) - \frac{\lambda}{2} \text{Var}(c_i)$.²¹ Hence in this case the total social surplus generated by efficient risk-sharing agreements is proportional to the aggregate consumption variance reduction. This greatly simplifies the computation of surpluses in the analysis below.

We use $TS(L)$ to denote the expected total surplus generated by an ex-ante Pareto-efficient risk-sharing agreement on network L , relative to agents consuming in autarky:

$$(3) \quad TS(L) := CE\left(\Delta \text{Var}(L, \emptyset)\right),$$

where, for $L' \subset L$, $\Delta \text{Var}(L, L')$ is the additional variance reduction obtained by efficient risk-sharing on network L instead of L' , and $CE(\cdot)$ denotes the certainty-equivalent value of a variance reduction.

2.4. Division of Surplus. The assumption that neighboring agents make pairwise efficient risk-sharing agreements pins down agreements up to state-independent transfers between neighboring agents, but does not constrain the latter transfers (hence the division of surplus) in any way. To determine these transfers, we follow the approach in Stole and Zwiebel (1996) and require that agreements are robust to split the difference renegotiations. This implies that the transfer is set in a way such that the incremental benefit that the link provides to the two agents is split equally between them. The motivation for this as provided by Stole and Zwiebel is to suppose that before agreements have been finalized any agent can opt out of an agreement with a neighbor and then renegotiate it so that he benefits from it as much as his

²¹See, for example, Arrow (1965).

neighbor.²² Of course, finding an agreement that i and j can reach without either wanting to renegotiate it requires them splitting the incremental benefits equally. And determining this split requires calculating the expected payoffs i and j would receive if they did not have an agreement. We therefore have to consider what agreements would prevail on the network without l_{ij} to find the risk sharing agreements i and j can reach on L , and so on. This results in a recursive system of conditions.

More formally, for a network L a contingent transfer scheme

$$(4) \quad \mathcal{T}(L) := \{t_{ij}(\omega, L')\}_{\omega \in \Omega, L' \subseteq L, ij: l_{ij} \in L},$$

specifies all transfers made in all subnetworks of L in all states of the world. The expected utility of agent i on a network $L' \subseteq L$ given a contingent transfer scheme $\mathcal{T}(L)$ is denoted $u_i(L', \mathcal{T}(L))$. Where there should be no confusion, we will abuse notation and drop the second argument.

For any network L , the expected utility vector $(u_1, \dots, u_{|\mathbf{N}|})$ is *robust to split the difference renegotiation* if there is a contingent transfer scheme $\mathcal{T}(L)$ such that $u_i = u_i(L, \mathcal{T}(L))$ for every $i \in \mathbf{N}$ and the following conditions hold:

- (i) $u_i(L') - u_i(L' \setminus \{l_{ij}\}) = u_j(L') - u_j(L' \setminus \{l_{ij}\})$ for every $l_{ij} \in L'$ and $L' \subseteq L$;
- (ii) transfers $\{t_{ij}(\omega, L')\}_{\omega \in \Omega, ij: l_{ij} \in L'}$ are pairwise efficient for all $L' \subseteq L$.

Suppose we want to find payoffs robust to split the difference renegotiation for the line network shown in Figure 1a. A first necessary condition is that agents 1 and 2 benefit equally from their link so that $u_1(L) - u_1(L \setminus \{l_{12}\}) = u_2(L) - u_2(L \setminus \{l_{12}\})$. But in order to ensure this condition is satisfied, we need to know $u_1(L \setminus \{l_{12}\})$ and $u_2(L \setminus \{l_{12}\})$. Without the link l_{12} agent 1 is isolated so $u_1(L \setminus \{l_{12}\}) = 0$. However, to find $u_2(L \setminus \{l_{12}\})$ we need to find payoffs for the three node network in Figure 1b. For this network robustness to split the difference renegotiation requires that $u_2(L \setminus \{l_{12}\}) - u_2(L \setminus \{l_{12}, l_{23}\}) = u_3(L \setminus \{l_{12}\}) - u_3(L \setminus \{l_{12}, l_{23}\})$. While $u_2(L \setminus \{l_{23}, l_{23}\}) = 0$, we need to consider the two node network shown in Figure 1c to find $u_3(L \setminus \{l_{12}, l_{23}\})$. For this network, payoffs must satisfy $u_3(L \setminus \{l_{12}, l_{23}\}) - u_3(L \setminus \{l_{12}, l_{23}, l_{34}\}) = u_4(L \setminus \{l_{12}, l_{23}\}) - u_4(L \setminus \{l_{12}, l_{23}, l_{34}\})$. As $u_3(L \setminus \{l_{12}, l_{23}, l_{34}\}) = u_4(L \setminus \{l_{12}, l_{23}, l_{34}\}) = 0$, the above condition simplifies to $u_3(L \setminus \{l_{12}, l_{23}\}) = u_4(L \setminus \{l_{12}, l_{23}\}) = V/2$, where the last equality follows from pairwise efficiency. Considering the three node network again, we now have the condition $u_2(L \setminus \{l_{12}\}) = u_3(L \setminus \{l_{12}\}) - V/2$. As the link l_{23} generates an incremental surplus of V to be split between agents 2 and 3, pairwise efficiency implies that $u_2(L \setminus \{l_{12}\}) = V/2$ and $u_3(L \setminus \{l_{12}\}) = V$. Finally, returning to the line network, we now have $u_1(L) = u_2(L) - V/2$. As the link l_{12} generates incremental surplus of V , $u_1(L) = V/2$ and $u_2(L) = V$.²³

²²For a detailed motivation of this assumption, and for noncooperative microfoundations, see Stole and Zwiebel (1996).

²³This argument only outlines why the payoffs $u_1(L) = V/2$ and $u_2(L) = V$ are necessary for robustness to split the difference renegotiations. By considering all other subnetworks, it can be shown that the payoffs




FIGURE 1. To find (gross) expected utilities that are robust to split the difference renegotiations on the (formed) line network shown we need to consider the expected utilities that would be obtained on all subnetworks.

Below we show that the requirement of robustness to split the difference renegotiation implies that the total surplus created by the risk-sharing agreement is divided among agents according to the Myerson value (Myerson 1977, 1980). The Myerson value is a cooperative solution concept defined in transferable utility environments that is a network-specific version of the Shapley value. The basic idea behind it is the same as for the Shapley value. For any order of arrivals of the players, the incremental contribution of an agent i to the total surplus can be derived as the difference between the total surpluses generated by the subgraph of $L(\mathbf{S})$ the subgraph $L(\mathbf{S} \setminus \{i\})$ if agents $\mathbf{S} \setminus \{i\}$ arrive before i . It is easy to see that, for any arrival order, the total surplus generated by L gets exactly allocated to the set of all agents. The Myerson value then allocates the average incremental contribution of a player to the total surplus, taken over all possible orders of arrivals (permutations) of the players, as the player's share of the total surplus. Thus, agent i 's Myerson value is²⁴

$$(5) \quad MV_i(L) := \sum_{\mathbf{S} \subseteq \mathbf{N}} \frac{(|\mathbf{S}| - 1)! (|\mathbf{N}| - |\mathbf{S}|)!}{|\mathbf{N}|!} (TS(L(\mathbf{S})) - TS(L(\mathbf{S} \setminus \{i\}))).$$

Proposition 2. *For any network L , any risk-sharing agreement that is robust to split the difference renegotiation yields expected payoffs to agents equal to their Myerson values: $u_i(L) = MV_i(L)$.*

Proof. Theorem 1 of Myerson (1980) states that there is a unique rule for allocating surplus for all subnetworks of L that satisfies the requirements of efficiency at the component level (note that this is an implicit requirement in Myerson's definition of an allocation rule) and, what Myerson (1980) defines as the equal-gains principle. Moreover, the expected payoff the above rule allocates to any player i is MV_i . Requirement (i) in our definition of robustness

$u_1(L) = u_4(L) = V/2$ and $u_2(L) = u_3(L) = V$ are the unique payoffs that are robust to split the difference renegotiations.

²⁴Our assumption that there is perfect risk sharing among path-connected agents ensures that a coalition of path connected agents generates the same surplus regardless of the exact network structure connecting them. This means that we are in the communication game world originally envisaged by Myerson. We do not require the generalization of the Myerson value to network games proposed in Jackson and Wolinsky (1996), which somewhat confusingly is also commonly referred to as the Myerson value.

to split the difference renegotiation is equivalent to the equal-gains principle as defined in Myerson (1980). Theorem 1 of Wilson (1968) implies that efficiency at the component level is equivalent to pairwise efficiency between neighboring agents, which is requirement (ii) in our definition of robustness to split the difference renegotiation. The result then follows immediately from Theorem 1 of Myerson (1980). \square

Proposition 2 is a direct implication of Myerson's axiomatization of the value. A special case of Proposition 2 is Theorem 1 of Stole and Zwiebel (1996), which in effect restricts attention to a star network.²⁵ Our contribution is to point out that their connection between robustness to split the difference renegotiations and the Shapley value can be extended to apply to all networks.

The above result shows that any decentralized negotiation procedure between neighboring agents that satisfies two natural properties (not leaving surplus on the table, and robustness to split the difference negotiations) leads to the total surplus created by risk-sharing divided according to the Myerson value, and to state-independent transfers between neighboring agents that implement this surplus division. Hence, from now on we assume that in the network formation process, all agents expect the surplus to be divided according to the Myerson value implied by the network that eventually forms.

Although we followed a decentralized approach to get to the implication that surplus is divided by the Myerson value, we note that on normative grounds such a division is also cogent in contexts in which there is a centralized community level negotiation over the division of surplus. This is because the Myerson value is a formal way of defining the fair share of an individual from the social surplus, as his average incremental contribution to the total social surplus (where the average is taken across all possible orders of arrival of different players, in the spirit of the Shapley value).

3. INVESTING IN SOCIAL RELATIONSHIPS

Having defined how formed networks map into risk-sharing arrangements, we can now consider agents' incentives to make social investments. We begin by providing the overall framework for the analysis. Then we look at a special case of our model, in which there is a single group. Building on these results we then consider the multiple group case.

In this section we formalize a game of network formation in which establishing links is costly, define efficient networks and identify different types of investment inefficiency.

We consider a two-period model in which in period 1 all agents simultaneously choose which other agents they would like to form links with, and in period 2 agents agree upon the ex-ante Pareto-efficient risk-sharing agreement specified in the previous section (i.e., the total surplus from risk sharing is distributed according to the Myerson value), for the network

²⁵Relative to Myerson's axiomatization, Stole and Zwiebel (1996) generate the key system of equations through considering robustness to renegotiations as we describe above, while Myerson wrote down the system of equations based only on fairness considerations. Stole and Zwiebel (1996) also provide non-cooperative bargaining foundations that underpin this system.

formed in the first period.²⁶ Implicit in our formulation of the timing of the game is the view that relationships are formed over a longer time horizon than that in which agreements are reached about risk sharing. By the time such agreements are being negotiated, the network structure is fixed, and investments into forming social relationships are sunk.

Formally, in period 1 we consider a network formation game along the lines of Myerson (1991): all agents simultaneously choose a subset of the other agents, indicating who they would like to form links (relationships) with. A link is formed between two agents if and only if they both want to form it (i.e., if both agents select each other). When agent i forms a link, he pays a cost $\kappa_w > 0$ if the link is with someone in the same group and $\kappa_a > \kappa_w$ if the link is with someone from a different group.

The collection of links formed in period 1 becomes social network L . Normalizing the utility from autarky to 0, we abuse notation²⁷ and let agent i 's *net* expected utility if network L forms be

$$(6) \quad u_i(L) = MV_i(L) - |\mathbf{N}_{G(i)}(i; L)|\kappa_w - \left(|\mathbf{N}(i; L)| - |\mathbf{N}_{G(i)}(i; L)| \right) \kappa_a.$$

The solution concept we apply to the simultaneous-move game described above is pairwise stability. A network L is *pairwise stable* with respect to expected utilities $\{u_i(L)\}_{i \in \mathbf{N}}$ if and only if for all $i, j \in \mathbf{N}$, (i) if $l_{ij} \in L$ then $u_i(L) - u_i(L \setminus \{l_{ij}\}) \geq 0$ and $u_j(L) - u_j(L \setminus \{l_{ij}\}) \geq 0$; and (ii) if $l_{ij} \notin L$ then $u_i(L \cup l_{ij}) - u_i(L) > 0$ implies $u_j(L \cup l_{ij}) - u_j(L) < 0$. In words, pairwise stability requires that no two players can both strictly benefit by establishing an extra link with each other, and no player can benefit by unilaterally deleting one of his links. From now on we will use the terms pairwise-stable and stable interchangeably.

Existence of a pairwise-stable network in our model follows from a result in Jackson (2003), stating that whenever payoffs in a simultaneous-move network formation game are determined based on the Myerson value, there exists a pairwise-stable network.

Our specification assumes that two agents forming a link have to pay the same cost for establishing the link. However, the set of stable networks would remain unchanged if we allowed the agents to share the total costs of establishing a link arbitrarily.²⁸ This is because for any link, the Myerson value rewards the two agents establishing the link symmetrically. Hence the agents can find a split of the link-formation cost such that establishing the link is profitable for both of them if and only if it is profitable for both of them to form the link with

²⁶For a complementary treatment of network formation when surplus is split according to the Myerson value, see Pin (2011).

²⁷In the previous section when investments had already been sunk we used $u_i(L)$ to denote i 's expected payoff before link formation costs.

²⁸More precisely, we could allow agents to propose a division of the costs of establishing each link as well as indicating who they would like to link to, and a link would then form only if both agents indicate each other and they propose the same split of the cost. A network would then be stable if it is a Nash equilibrium of this expanded network formation game and if there is no new link $l_{ij} \notin L$, and some split of the cost of forming this link, that would make both i and j strictly better off if formed.

an equal split of the cost. For this reason we stick with the simpler model with exogenously given costs.

A network L is *efficient* when there is no other network L' —and no risk sharing agreement on L' —that can make everyone at least as well off as they were on L and someone strictly better off. Let $|L_w|$ be the number of within-group links, and let $|L_a|$ be the number of across-group links. As expected utility is transferable in certainty-equivalent units, efficient networks must maximize the net total surplus $NTS(L)$:

$$(7) \quad NTS(L) := TS(L) - 2|L_w|\kappa_w - 2|L_a|\kappa_a,$$

Clearly, two necessary conditions for a network to be efficient are that the removal of a set of links does not increase $NTS(L)$ and the addition of a set of links does not increase $NTS(L)$. If there exists a set of links the removal of which increases $NTS(L)$, we will say there is *overinvestment* inefficiency. If there exists a set of links the addition of which increases $NTS(L)$, we will say there is *underinvestment* inefficiency.²⁹ A network is *robust to underinvestment* if there is no underinvestment inefficiency and no agent can strictly benefit from deleting a link would result in underinvestment inefficiency. A network is *robust to overinvestment* if there is no overinvestment inefficiency and no pair of agent i, j can both strictly benefit from creating the link l_{ij} .

We will say that a link l_{ij} is *essential* if after its removal i and j are no longer path connected while it is *superfluous* if after its removal i and j are still path connected.

Remark 3. *Preventing overinvestment requires that all links be essential. Superfluous links create no social surplus and are costly. In all efficient networks, therefore, every component must be a tree.*

In most of the analysis below, we focus on investigating the relationship between stable networks and efficient networks. Additionally, we investigate the amount of inequality prevailing in equilibria in our model. For this, we will use the Atkinson class of inequality measures (Atkinson, 1970). Specifically we consider a welfare function $W : \mathbb{R}^{|\mathbf{N}|} \rightarrow \mathbb{R}$ that maps a profile of expected utilities into the real line such that

$$(8) \quad W(u) = \sum_{i \in \mathbf{N}} f(u_i),$$

where $f(\cdot)$ is assumed to be an increasing, strictly concave and differentiable function. The concavity of $f(\cdot)$ captures the social planner's preference for more equal income distributions. Supposing all agents instead received the same expected utility u' , we can pose the question

²⁹Note that these definitions are not mutually exclusive (there can be both underinvestment and overinvestment inefficiency) or collectively exhaustive (inefficient networks can have neither underinvestment nor overinvestment inefficiency if an increase in the net total surplus is only possible by the simultaneous addition and removal of edges).

what aggregate expected utility is required to keep the level of the welfare function constant.³⁰ In other words we find the scalar $u' : |\mathbf{N}|f(u') = \sum_{i \in \mathbf{N}} f(u_i)$. Letting $\bar{u} = (1/|\mathbf{N}|) \sum_{i \in \mathbf{N}} u_i$ be the mean expected utility, Atkinson's inequality measure (or index) is given by

$$(9) \quad I(f) = 1 - \frac{u'}{\bar{u}} \in [0, 1].$$

We let \mathcal{I} be the set (class) of Atkinson inequality measures and note that any $I(f) \in \mathcal{I}$ equals zero if and only if all agents receive the same expected utility.³¹ Two different inequality measures from the Atkinson class can rank the inequality of two distributions differently. However, certain pairs of distributions are ranked the same way by all members of the class, such as when one distribution is a mean-preserving spread of the other one.

4. WITHIN-GROUP NETWORK FORMATION

In this section we assume that $|\mathbf{M}| = 1$, that is, that agents are ex-ante symmetric, and any differences in their outcomes stem from their stable positions on the social network. This will lay the foundations for the more general case considered in the next section.

We begin our investigation by proving a general characterization of the set of stable networks. Recall that a *minimal path* between i and j is any path between i and j such that no other path between i and j is a subsequence of that path. If there are K minimal paths between i and j on the network L , we let $\mathbf{P}(i, j, L) = \{P_1(i, j, L), \dots, P_K(i, j, L)\}$ be the set of these paths. For every $k \in \{1, \dots, K\}$, let $|P_k(i, j, L)|$ be the cardinality of the set of agents on the minimal path $P_k(i, j, L)$.³² We can now use these definitions to define a quantity that captures how far away two agents are on a network in terms of the probability that for a random arrival order they will be connected without a direct link when the second of the two agents arrives. We will refer to this distance as the agents' Myerson distance:

$$(10) \quad md(i, j, L) := \frac{1}{2} - \sum_{k=1}^{|\mathbf{P}(i, j, L)|} (-1)^{k+1} \left(\sum_{1 \leq i_1 < \dots < i_k \leq |\mathbf{P}(i, j, L)|} \left(\frac{1}{|P_{i_1} \cup \dots \cup P_{i_k}|} \right) \right).$$

This expression calculates the probability that for a random arrival order the link l_{ij} will be essential immediately after i arrives,³³ using the classic *inclusion-exclusion principle* from combinatorics. This probability is important because it affects i 's incentives to link to j .

As an illustration, suppose that there is a unique indirect path $P_1(i, j, L)$ between i and j that contains K agents, including i and j . We then have $md(i, j, L) = 1/2 - 1/K$. To

³⁰This exercise is analogous to the certainty equivalent exercise that can be undertaken for an agent facing stochastic consumption.

³¹As $f(\cdot)$ approaches the linear function the social planner cares less about inequality and $I(f) \rightarrow 0$. Nevertheless, strict concavity prevents $I(f)$ equaling 0 unless all agents receive the same expected utility.

³²For example, for a path $P_k(i, j, L) = \{i, i', i'', j\}$, $|P_k(i, j, L)| = 4$ and for a path $P_{k'}(i, j, L) = \{i, i', i''', i''', j\}$, $|P_{k'}(i, j, L)| = 5$. Finally, we will let $|P_k(i, j, L) \cup P_{k'}(i, j, L)| = 5$ denote number of different agents on path $P_k(i, j, L)$ or path $P_{k'}(i, j, L)$.

³³If for a given arrival order, agents $\mathbf{S} \subseteq \mathbf{N}$ arrive before i , then l_{ij} is essential immediately after i arrives if it is essential on the network $L(\mathbf{S} \cup \{i\})$.

see where this expression comes from, note that there are two reasons why l_{ij} might not be essential when i arrives. First, j might not yet have arrived. This occurs with probability $1/2$. Second, all other agents on the path $P_1(i, j, L)$, including j , might have already arrived. This occurs with probability $1/K$. The probability of both events occurring is 0 because collectively they require j to be both present and absent, so we can just sum them. Thus, the probability that l_{ij} is essential when i arrives, is $1 - 1/2 - 1/K = md(i, j, L)$.

Suppose now that there are two (minimal) paths between i and j , $P_1(i, j, L)$ and $P_2(i, j, L)$, on the network L . Suppose that $P_1(i, j, L) = \{i, i', i'', j\}$ and $P_2(i, j, L) = \{i, i', i''', i'''', j\}$. We need to find the probability that all the agents on at least one of these paths are present when i arrives. To avoid double counting, we need to add the probability that all the agents on $P_1(i, j, L)$ are present ($1/4$) to the probability all the agents on $P_2(i, j, L)$ are present ($1/5$) and then subtract the probability that all the agents on both paths are present ($1/6$).³⁴ So $md(i, j, L) = 1 - 1/2 - 1/4 - 1/5 + 1/6$. The Myerson distance calculation provides the general way of accounting for the probability that all the agents on at least one of the possible paths are present.

Lemma 4. *If all agents are from the same group network L is pairwise stable if and only if*

$$\begin{aligned} (i) \quad md(i, j, L \setminus \{l_{ij}\}) &\geq \kappa_w/V && \text{for all } l_{ij} \in L, \text{ and} \\ (ii) \quad md(i, j, L) &\leq \kappa_w/V && \text{for all } l_{ij} \notin L. \end{aligned}$$

The proof is relegated to Appendix B. Recall from equation 3 that the social benefits of a link is proportional to the variance reduction it generates. For a single group, if a link l_{ij} is essential in the network $L \cup \{l_{ij}\}$, then this variance reduction is $\Delta \text{Var}(L \cup \{l_{ij}\}, L) = (1 - \rho_w)\sigma^2$.

The crucial feature of this expression is that it does not depend on size of the network components the link l_{ij} connected on L . Although in general the size of these components does affect the consumption variance, two effects exactly offset each other.³⁵ On the one hand, in larger components there are more people to benefits from the essential link. On the other hand, people are already able to smooth there consumption effectively.

As the social value of a non-essential, or superfluous link, is always zero the total surplus generated by a network L takes a very simple form. Let $f(L)$ be the number of network components on L . Then $TS(L) = CE\left(\Delta \text{Var}(L, \emptyset)\right) = (|\mathbf{N}| - f(L))\frac{\lambda}{2}(1 - \rho_w)\sigma^2$. Since the surplus created by any essential link is

$$(11) \quad V := \frac{\lambda}{2}(1 - \rho_w)\sigma^2,$$

³⁴Note that all the agents on both $P_1(i, j, L)$ and $P_2(i, j, L)$ will be present if and only if i', i'', i''', i'''' and j are present before i arrives.

³⁵Let $L(\mathbf{S}_1)$ and $L(\mathbf{S}_2)$ be the network components of agent i and agent j on network $L \setminus \{l_{ij}\}$, and let $|\mathbf{S}_1| = s_1$ and $|\mathbf{S}_2| = s_2$. Then the sum of consumption variances on $L(\mathbf{S}_1)$ and $L(\mathbf{S}_2)$ (with Pareto-efficient risk sharing) are $\frac{s_1 + s_1(s_1 - 1)\rho_w}{s_1}\sigma^2$ and $\frac{s_2 + s_2(s_2 - 1)\rho_w}{s_2}\sigma^2$, respectively. Once S_1 and S_2 are connected through l_{ij} , the sum of consumption variances on $L(\mathbf{S}_1 \cup \mathbf{S}_2)$ becomes $\frac{s_1 + s_2 + (s_1 + s_2)(s_1 + s_2 - 1)\rho_w}{s_1 + s_2}\sigma^2$. This implies that the consumption variance reduction induced by the link l_{ij} is $\Delta \text{Var}(L \cup \{l_{ij}\}, L) = (1 - \rho_w)\sigma^2$.

the total gross surplus is equal to the latter constant times the number of network component reductions obtained relative to the empty network.

To consider individual incentives to form links we can use the definition of the Myerson value and consider the average marginal contribution an agent makes to total surplus over all possible arrival orders. Specifically, we want to consider the increase in i 's Myerson value due to a link l_{ij} . The link l_{ij} will reduce the number of components in the graph by one when i arrives relative to the counterfactual component reduction without l_{ij} , if and only if j has already arrived and there is no other path between i and j . In other words, the link increases i 's marginal contribution to total surplus if and only if it is essential when i is added. Moreover, for the permutations in which l_{ij} is essential it contributes V to i 's marginal contribution to total surplus. Averaging over arrival order, the value to i of the link $l_{ij} \in L$ is $md(i, j, L \setminus \{l_{ij}\})V$, while the value to establishing a new link $l_{ij} \notin L$ is $md(i, j, L)V$.

If a link l_{ij} is essential on L then for any arrival order, there will always be a component reduction of 1 when the later of i or j is added. Therefore, $md(i, j, L) = 1/2$, and l_{ij} will be formed as long as $V > 2\kappa_w$. As V is the social value of forming the link and $2\kappa_w$ is the total cost of forming it, when all agents are from the same group there is never underinvestment in a stable network or overinvestment in an essential link.

Proposition 5. *If all agents are from the same group then there is never underinvestment in a stable network. Furthermore, there is never overinvestment in an essential link.*

The proof is relegated to Appendix B. When all agents are from the same group Proposition 5 establishes that there is never overinvestment in an essential link, but overinvestment in superfluous links is possible. If the costs of link formation are low enough then agents will receive sufficient benefits from establishing superfluous links to be incentivized to do so. Even if a link l_{ij} is superfluous on L , for some arrival orders it will be essential on the induced subnetwork at the moment when i is added and make a positive marginal contribution to total surplus.³⁶ An example of such overinvestment is shown in Appendix A.

An immediate implication of Proposition 5 is that if all agents are from the same group and $2\kappa_w > V$ then the only stable network is the empty one and this network is efficient, while if $2\kappa_w < V$ then all stable networks have only one network component (all agents are path connected). For the remainder of the paper we focus on the parameter range for which the empty network is inefficient for a single group and assume $2\kappa_w < V$. We refer to this as our *regularity condition* and omit it from the statement of subsequent results.

Under this regularity condition the set of efficient networks are the set of tree networks in which all agents are path connected. In other words, all agents must be in the same component and all links must be essential. We will now focus on which, if any, of these efficient networks are stable. As by Proposition 5 there is never any underinvestment in any stable network the only reason an efficient network will not be stable is if two agents have

³⁶Consider, for examples, arrival orders in which i arrives first and j arrives second.




FIGURE 2. Three possible network structures for connecting 4 agents and the resulting net payoffs.

a profitable deviation by forming an additional (superfluous) link. We therefore focus on investigating what network structures minimize incentives for overinvestment. As we will see, this question is also related to the issue of inequality that different network structures imply.

Figure 2 illustrates three networks: A line (Figure 2a); a circle (Figure 2b) and a star (Figure 2c). While the line and star networks are efficient, the circle network is not as it includes a superfluous link. Among the two efficient networks, the star is more stable than the line. Applying Lemma 4, whenever the line is stable so is the star but there are parameter values for which the star is stable and the line is not. While the star is more stable than the line, it also results in more inequality. The expected utility distribution obtained on the line network can be generated from that obtained on the star network by the best off agent (agent 2) transferring $(V - 2\kappa_w)/2 > 0$ units of expected utility to one of the worst off agents (agent 3). This is enough to ensure that the expected utility distribution on the star is more unequal than the expected utility distribution on the line for any inequality measure in the Atkinson class. We generalize these insights in Proposition 6.

Proposition 6. *Suppose all agents are from the same group.*

- (i) *If there exists an efficient stable network then star networks are stable, and for a range of parameter specifications only star networks are stable. If a line network is stable then all efficient networks are stable.*
- (ii) *For all inequality measures in the Atkinson class, among the set of efficient network, star networks and only star networks maximize inequality, while line networks and only line networks minimize inequality.*

The proof is in Appendix B but we provide some intuition after we discuss the result. Proposition 6 states that, in a certain sense, among the set of efficient networks the star is the most stable but maximizes inequality, while the line minimizes inequality but is least stable. This indicates a novel tension between efficiency and inequality.

To gain intuition for Proposition 6, recall that Proposition 5 implies that an efficient network will be stable if and only if no pair of players have a profitable deviation in which

they form a superfluous link. By Lemma 4 the incentives for two agents to form such a link are strictly increasing in their Myerson distance. Thus, a network is stable if and only if the pair of agents furthest apart from each other, in terms of their Myerson distance, cannot benefit from forming a link. As efficient networks are tree networks, the Myerson distance between any two agents depends only the length of the unique minimal path between them.³⁷ The longest minimal path between any pair of agents is, by definition, the diameter of the network $d(L)$. So, an efficient network is stable if and only if its diameter is sufficiently small. More precisely, an efficient network L is stable if and only if its diameter is weakly less than $\bar{d}(\kappa_w, V)$, where $\bar{d}(\kappa_w, V)$ is increasing in κ_w , decreasing in V and integer valued.³⁸

Let $\mathcal{L}^e(\mathbf{N})$ be the set of efficient networks. Star networks have the smallest diameter among networks within this set, while line networks have the largest diameter among networks within this set. This establishes part (i) of Proposition 6.

To gain intuition for part (ii) a first step is noting that on any efficient network agents' net payoffs are proportional to their degree (i.e., the number of neighbours they have):³⁹ $u_i(L) = |\mathbf{N}(i; L)|(V/2 - \kappa_w)$. The key insight is then showing that for any network in the set $\mathcal{L}^e(\mathbf{N})$, the star network can be obtained by rewiring the network (deleting a link $l_{ij} \in L$ and adding a link $l_{ik} \notin L$) in such a way that at each step we increase the degree of the agent who already has the highest degree and reduce the degree of some other agent and obtain a new network in $\mathcal{L}^e(\mathbf{N})$. This process transfers expected utility to the agent with the highest expected payoff from some other agent. This increases inequality for any inequality measure within the Atkinson class. Likewise, we can obtain the line network from any network in the set $\mathcal{L}^e(\mathbf{N})$ by rewiring the network to decrease the degree of the agent with the highest degree at every step. This transfers expected utility from agent with the highest expected payoff to some other agent and decreases inequality for any inequality measure in the Atkinson class.

5. CONNECTIONS ACROSS GROUPS

We now generalize our model by permitting multiple groups. These different groups might correspond to people from different villages, different occupations, or different social status groups, such as castes. We will first show that (under our regularity condition) there is still never any underinvestment within a group. However, this does not apply to links that bridge groups. As, by assumption, incomes are more correlated within a group than across a group, there can be significant benefits from establishing such links and not all these benefits accrue to the agents forming the link. Intuitively, an agent establishing a bridging link to another group provides other members of his group with access to a less correlated income stream,

³⁷Suppose d is the number of agents on the unique path connecting i and j . The probability that this path exists when agent i arrives is $1/d$. In addition, if agent j has not yet arrived, which occurs with probability $1/2$, i would not benefit from the link l_{ij} , so i 's expected payoff from forming a superfluous link to j is $(1 - 1/2 - 1/d)V$. We also note that as d gets large, this converges to $V/2$ which is the value i receives from forming an essential link.

³⁸We show in the proof that $\bar{d}(\kappa_w, V) = \lfloor 2V/(V - 2\kappa_w) \rfloor$.

³⁹This is also known as an agent's degree centrality.

which benefits them. As agents providing such bridging links are unable to appropriate all the benefits these links generate, and these links are relatively costly to establish, there can be underinvestment.

To analyze the incentives to form links within a group, we first need to consider the variance reduction obtained by a within-group link. Such a link may now connect two otherwise separate components comprised of arbitrary distributions of agents from different groups. Suppose the agents in $\mathbf{S}_0 \cup \dots \cup \mathbf{S}_k$ and the agents in $\widehat{\mathbf{S}}_0 \cup \dots \cup \widehat{\mathbf{S}}_k$ form two distinct network components, where for every $i \in \{0, \dots, k\}$, the agents in \mathbf{S}_i and those in $\widehat{\mathbf{S}}_i$ are all from group i . Consider now a potential link l_{ij} connecting the two otherwise disconnected components. Letting s_0 be the number of agents in group 0, the variance reduction obtained is:⁴⁰

$$(12) \quad \Delta \text{Var}(L \cup l_{ij}, L) = \left[(1 - \rho_w) + \frac{\left(\sum_{i=0}^k \left(\hat{s}_i \sum_{j=0}^k s_j - s_i \sum_{j=0}^k \hat{s}_j \right)^2 \right)}{\left(\sum_{i=0}^k s_i \right) \left(\sum_{i=0}^k \hat{s}_i \right) \left(\sum_{i=0}^k s_i + \hat{s}_i \right)} (\rho_w - \rho_a) \right] \sigma^2.$$

The key feature of this variance reduction is that it is always weakly greater than $(1 - \rho_w)\sigma^2$, which is the variance reduction we found in the previous section when all agents were from the same group. Thus, the presence of across-group links only increases the incentives for within-group links to be formed. A within-group link can now give (indirect) access to less correlated incomes from other groups and so is weakly more valuable. This implies that there will still be no underinvestment under our regularity condition that $2\kappa_w < V$.⁴¹ The above reasoning is formalized by Proposition 7.

Proposition 7. *There is no underinvestment between any two agents from the same group in any stable network.*

The proof of Proposition 7 is in Appendix B. While underinvestment is not possible within group, it is possible across groups. An example of this is shown in Appendix A. Although when all agents are from the same group the value of an essential link does not depend on the sizes of the components it connects, the value of an essential link connecting two different groups of agents increases in the sizes of the components. To demonstrate this formally, consider an isolated group that has no across-group connections and consider the incentives for a first such connection to be formed. Thus the first component consist of agents from a single group, say group 0. We let the second component consist of agents from one or

⁴⁰By definition

$$\Delta \text{Var}(L \cup l_{ij}, L) = \text{Var}(L(\mathbf{S}_0, \dots, \mathbf{S}_k)) + \text{Var}(L(\widehat{\mathbf{S}}_0, \dots, \widehat{\mathbf{S}}_k)) - \text{Var}(L(\mathbf{S}_0 \cup \widehat{\mathbf{S}}_0, \dots, \mathbf{S}_k \cup \widehat{\mathbf{S}}_k)).$$

Recalling that

$$\text{Var}(L(\mathbf{S}_0, \mathbf{S}_1, \dots, \mathbf{S}_k)) = \left(\sum_{i=0}^k (s_i + s_i(s_i - 1)\rho_w) + 2\rho_a \sum_{i=0}^{k-1} \left(s_i \sum_{j=i+1}^k s_j \right) \right) \sigma^2 / \sum_{i=0}^k s_i,$$

some algebra yields the result.

⁴¹Recall that this regularity condition just requires that it is efficient for two agents in the same group, both without any other connections, to form a link.

more of the other groups (1 to k). The variance reduction obtained by connecting these two components is

$$(13) \quad \Delta \text{Var}(L \cup l_{ij}, L) = \left[(1 - \rho_w) + \frac{\hat{s}_0 \left(\left(\sum_{i=1}^k s_i \right)^2 + \sum_{i=1}^k s_i^2 \right)}{\left(\sum_{i=1}^k s_i \right) \left(\hat{s}_0 + \sum_{i=1}^k s_i \right)} (\rho_w - \rho_a) \right] \sigma^2,$$

which is increasing in \hat{s}_0 :

$$(14) \quad \frac{\partial \Delta \text{Var}(L \cup l_{ij}, L)}{\partial \hat{s}_0} = \frac{\left(\sum_{i=1}^k s_i \right)^2 + \sum_{i=1}^k s_i^2}{\left(\hat{s}_0 + \sum_{i=1}^k s_i \right)^2} (\rho_w - \rho_a) \sigma^2 > 0.$$

The inequality follows since $\rho_w > \rho_a$. Thus if agents i and j who connect two otherwise unconnected groups they receive a strictly smaller combined private benefit than the social value of the link. To see why, suppose that on the network L the link l_{ij} is essential, and without l_{ij} there would be two components, the first connecting agents from group $G(i)$ and the second connecting agents from group $G(j) \neq G(i)$. Consider the Myerson value calculation. For arrival orders in which i or j is last to arrive, the value of the additional variance reduction due to l_{ij} obtained upon the arrival of the later of i or j , is the same as its marginal social value, i.e., the value of variance reduction obtained by l_{ij} on L . For any other arrival order the value of variance reduction due to l_{ij} when the later of i or j arrives is strictly less. Averaging over these arrival orders, the link l_{ij} contributes less to i and j 's combined Myerson values than its social value, leading to the possibility of underinvestment.

Besides underinvestment, overinvestment is also possible across groups. Forming superfluous links will increase an agent's share of surplus without improving overall risk sharing and can therefore create incentives to overinvest. Nevertheless, when κ_a is relatively high, underinvestment rather than overinvestment in across-group links will be the main efficiency concern. In many settings, within-group links are relatively cheap to establish in comparison to across-group links. For example, when the different groups correspond to different castes, as in our data, it can be quite costly to be seen interacting with members of the other caste (e.g., Srinivas (1962), Banerjee et al. (2013b)). Motivated by this, and because across-group links are considerably sparser in our data (to be described in the next section) than within-group links, we focus our attention on this parameter region. More concretely, below we investigate what within-group network structures create the best incentives to form across-group links and what network structures minimize the incentives for overinvestment within group. Remarkably, we find that these two forces push within-group network structures in the same direction, and in both cases towards inequality in the society.

We begin by considering within-group overinvestment, which corresponds to the formation of superfluous links within-group links. We found in the previous section that when all agents are from the same group the star is the efficient network that minimized the incentives

for overinvestment. However, once we include links to other groups, the analysis is more complicated. The variance reduction a within-group link generates is still 0 if the link is superfluous, but when the link is essential it depends on the distribution of agents across the different groups the link grants access to. Moreover, the variance reduction may be decreasing or increasing in the numbers of people in those groups.⁴² This makes the Myerson value calculation substantially more complicated. When all agents were from the same group all that mattered was whether the link was essential when added. Now, for each arrival order in which the link is essential, we also need to keep track of the distribution of agents across the different groups that are being connected. Nevertheless, our earlier result generalizes to this setting, although the argument establishing the result is more subtle.

To state the result, it is helpful to define a new network structure. A *center-connected star* network is a network in which all within-group network structures are stars and all across-group links are held by the center agents in these stars. We denote the set of center-connected star networks by \mathcal{L}^{CCS} .

Proposition 8. *If any efficient network L is robust to overinvestment within group, then any center-connected star network $L' \in \mathcal{L}^{CCS}$ is also robust to overinvestment within group. Moreover, if $L \notin \mathcal{L}^{CCS}$, then for a range of parameter specifications any center-connected star network $L' \in \mathcal{L}^{CCS}$ is robust to overinvestment within group but L is not.*

The proof of Proposition 8 is in Appendix B. In Proposition 6 we found that when all agents are from the same group, incentives for overinvestment (within group) are minimized by forming a (within-group) star. However, the incentives to form superfluous within-group links are weakly greater when someone within the group holds an across-group link (see equation 12). We can therefore think of the incentives for over-investment we found in Proposition 6 as a lower bound on the minimal incentives we can hope to obtain once there are across-group links. A key step in the proof of Proposition 8 shows that this lower bound is obtained by all center-connected star networks.

Consider a center-connected star network L' . As the agent at the center of a within-group star, agent k , has a link to all agents within the same group, we can focus on the incentives of two non-center agents from the same group, i and j , to form a superfluous link. Consider any subset of agents $\mathbf{S} \subseteq \mathbf{N}$ such that $i, j \in \mathbf{S}$. On the induced subnetwork $L'(\mathbf{S})$ either l_{ij} is superfluous or else $k \notin \mathbf{S}$. This implies that no across-group links are present whenever the additional link l_{ij} makes a positive marginal contribution. Hence considering different arrival order, the average marginal contribution of such a link when it is added is the same on the star network with no across-group links as for a center-connected star network: The lower bound on within-group overinvestment incentives is obtained.

⁴²In the case of an essential across-group link that bridges two otherwise disconnected groups, the comparative statics are unambiguous. In this case, the variance reduction is increasing in the sizes of the groups connected, as shown by inequality (4).

We now consider the within-group network structures that maximize the incentives for an across-group link to be formed. We have already established that the marginal contribution of a first bridging link to the total surplus is increasing in the sizes of the groups it connects. By the Myerson calculation, the agents with the strongest incentives to form such links are then those who will be linked to the greatest number of other agents within their group when they arrive. The result below formalizes this intuition.

Let $\mathcal{A}(\mathbf{S}_k)$ be the set of possible arrival orders for the agents in \mathbf{S}_k . For any arrival order $A \in \mathcal{A}(\mathbf{S})$, let $\mathbf{T}_i(A)$ be the set of agents to whom i is path connected on $L(\mathbf{S}')$, where \mathbf{S}' is the set of agents (including i) that arrive weakly before i . Let $T_i^{(m)}$ be a random variable, taking values equal to the cardinality of $\mathbf{T}_i(A)$, where A is selected uniformly at random from those arrival orders in which i is the m -th agent to arrive.

We will say that agent $i \in \mathbf{S}_k$ is more *Myerson central* (from now on, simply more central, for brevity) within his group than agent $j \in \mathbf{S}_k$ if $T_i^{(m)}$ first-order stochastically dominates $T_j^{(m)}$ for all $m \in \{1, 2, \dots, |\mathbf{S}_k|\}$.⁴³ In other words, considering all the arrival orders in which i is the m -th agent to arrive, and all the arrival orders in which j is the m -th agent to arrive, the size of i 's component at i 's arrival is larger than that of j 's at j 's arrival in the sense of first-order stochastic dominance.⁴⁴ This measure of centrality provides a partial ordering of agents.

Lemma 9. *Suppose agents in \mathbf{S}_0 form a network component, and all other agents in \mathbf{N} form another network component. Let $i, i' \in \mathbf{S}_0$ and let $j \notin \mathbf{S}_0$. If i is more central within group than i' , then i receives a higher payoff from forming l_{ij} than i' receives from forming $l_{i'j}$:*

$$MV(i; L \cup l_{ij}) - MV(i; L) > MV(i'; L \cup l_{i'j}) - MV(i'; L)$$

The proof is relegated to Appendix B. The key step in the proof pairs the arrival orders of a more central agents with a less central agent, so that in each case the more central agent is connected to weakly more people in the same group upon his arrival, and to the same set of people from other groups. Such a pairing of arrival orders is possible from the definition of centrality, and in particular the first-order stochastic dominance it requires.

Lemma 9 shows that more central agents have better incentives to form intergroup links. We can then consider the problem of maximizing the incentives to form intergroup links by choosing the within-group network structures (networks containing only within-group links). We will say that the within-group network structures that achieve these maximum possible incentives are most robust to underinvestment inefficiency across groups.

⁴³We also use this notion of centrality to compare the within-group centrality of the same agent on two different network structures. To avoid repetition we do not state the slightly different definition that would apply this situation.

⁴⁴An alternative and equivalent definition is that i is more central than j if there exists a bijection $B : \mathcal{A}(\mathbf{S}_k) \rightarrow \mathcal{A}(\mathbf{S}_k)$ such that $|\mathbf{T}_i(A)| \geq |\mathbf{T}_j(B(A))|$ and $A(i) = A'(j)$, where $A(i)$ is i 's position in the arrival order A and $A' = B(P)$.




FIGURE 3. Center-connected within-group stars, in a context with two groups.

Proposition 10. *If any efficient network L is robust to underinvestment across group, then some center-connected star network $L' \in \mathcal{L}^{CCS}$ is also robust to underinvestment across group. Moreover, if $L \notin \mathcal{L}^{CCS}$, then for a range of parameter specifications the center-connected star network $L' \in \mathcal{L}^{CCS}$ is robust to underinvestment across group but L is not.*

The proof of Proposition 10 is in Appendix B. Intuition can be gained from Lemma 9. This Lemma shows that agents have better incentives to provide a bridging link across group when they are more central within their own group. Thus to maximize the incentives of an agent to provide an across-group link, we need to maximize the centrality of this agent within group. This is achieved by any network that directly connects this agent to all others in the same group. However, only one of these within-group network structures can be part of an efficient network, and this is the star network, with the agent providing the across-group link at the center.

Figure 3 shows a center-connected star network when there are groups. As long as it is efficient for these groups to be connected, center-connected star networks and only the center-connected star networks minimize the incentives for within-group overinvestment (by Proposition 9) and minimize the incentives for across-group underinvestment (by Proposition 10).

The above results further reinforce the tension between efficiency and equality. The within-group star not only minimizes the incentives for within-group overinvestment, it also minimizes the incentives for across-group underinvestment. If an agent i provides an across-group link, then of all the possible within-group network structures, the star with i at the center maximizes i 's incremental payoff from establishing the link.⁴⁵

6. COMPARATIVE STATICS

We now turn our attention to taking our model to the data. Recall that the broad predictions of our model thus far are that (1) there is endogenous centrality, (2) there is no underinvestment within groups, (3) agents cannot be too far away from each other (in terms of the Myerson distance), (4) agents that have across-caste links should be more Myerson

⁴⁵Nevertheless, for some (but not all) parameter values, a within-group star will be more equitable when the central agent forms an across-group link than without it, because the across-group link generates positive spillovers to the whole group.

central. However, these predictions provide neither a clean nor powerful test of our theory. Predictions (2) and (3) depend on an unobservable linking cost parameter and there are many alternative stories, including ones not directly connected to risk sharing, that generate similar predictions. For example, if individuals have heterogenous time budgets and made random links within and across groups, predictions (1) and (4) are mechanically generated.

We therefore turn to more subtle predictions that rely on the comparative statics under our model as we change parameters of the economic environment. We study four demanding predictions from the theory with richer empirical content.

The first two predictions look at how network structure, described by Myerson distances among agents, depends on income variability and correlation. The intuition is that in networks where the gains from risk sharing are higher (income is more variable/less highly correlated), the Myerson distance between any two agents cannot be too high in a stable network. Otherwise, a pair of individuals would be incentivized to form an additional link. When the gains from risk sharing are lower, the Myerson distance can be larger.

The latter two predictions describe the composition of across-group bridging links. Individuals with higher Myerson centrality have better incentives to form bridging links. However, when income variability or the within- versus across-group income correlation is high, many members of either group can find it worthwhile to form a cross-group bridging link. Thus, we expect the average centrality within their own group of forming the cross-group link is lower when the incentives to form such links are higher. We now formalize our predictions.

In the case of one group, Proposition 4 provides the key characterization of the set of pairwise-stable networks. This characterization yields an exact expression for increased pay-offs two agents would receive were they to form an additional link. For a risk-sharing network to be stable, these benefits should be less than the cost of forming the link: for every i and j that do not have a link,

$$(15) \quad md(i, j, L) \leq \frac{2\kappa_w}{(1 - \rho_w)\lambda\sigma^2}.$$

In our empirical setting, described below, we are interested in Indian village networks where there are multiple groups, given by caste. Nevertheless, inequality (15) provides an appropriate benchmark for within-group links. Recall that the left-hand side of the inequality captures the probability that the link is not essential in the case of a random arrival order and the right-hand side gives the value of the variance reduction obtained. The key complication is that when a within-group link is essential for a subgraph but it connects two otherwise separate components that contain people from multiple groups, then the value of the variance reduction will depend on the composition of the people within the two components.

Consider the variance reduction obtained by combining any two components, with agents in groups $\mathbf{S}_0, \dots, \mathbf{S}_k$ and $\hat{\mathbf{S}}_0, \dots, \hat{\mathbf{S}}_k$ respectively. From equation 12, for some function f , the certainty-equivalent value of this variance reduction is

$$[(1 - \rho_w) + f(s_0, \dots, s_k, \hat{s}_0, \dots, \hat{s}_k)(\rho_w - \rho_a)] \frac{\lambda \sigma^2}{2}.$$

This is hard to compute in general. We approximate it by assuming that within each component there are the same number of people from each group: $s_i = \alpha$ and $\hat{s}_i = \beta$ for $i = 0, \dots, k$.

Proposition 11. *The certainty-equivalent value of variance reduction obtained by linking a component with α people from each of groups $0, \dots, k$ to a component with β people from each of groups $0, \dots, k$ is*

$$\frac{(1 - \rho_w)\lambda\sigma^2}{2}.$$

The proof is in Appendix B. Proposition 11 shows that the variance reduction obtained by permitting two components containing agents from multiple groups to share risk is the same as when all agents are from the same group, as long the proportion of people from each group is the same in each component.

The above considerations lead to the following predictions:

- P1. In villages with higher σ^2 , the average Myerson distance should be smaller.
- P2. In villages with lower ρ_w , the average Myerson distance should be smaller.

While our first set of predictions looks at the relationship between the Myerson distance and the environmental parameters, our second set of predictions looks at the composition of the links. Our interest is in which agents provide the across-group links. Proposition 9 shows more central agents have better incentives to provide an across-group link. The importance of centrality will depend on the overall strength of the incentives to form across-caste links. When income variance is high, or within-caste income correlation is high relative to across-caste income correlation, the incentives to form an across-caste link will also be high, and so network position will be less important; villagers in more varied locations will have sufficient incentives to form across-caste links. More formally, from the variance reduction given in equation 12 it is straightforward to show that for an across-caste bridging link l_{ij} :

$$\frac{\partial \Delta \text{Var}(L, L \cup \{l_{ij}\})}{\partial \sigma^2} > 0, \quad \frac{\partial \Delta \text{Var}(L, L \cup \{l_{ij}\})}{\partial \rho_w} > 0, \quad \frac{\partial \Delta \text{Var}(L, L \cup \{l_{ij}\})}{\partial \rho_a} < 0.$$

This means that the incentives to form an across-caste link are increasing in σ^2 and $\rho_w - \rho_a$, leading to the following predictions:

- P3. In villages with higher σ^2 , the association between within-group centrality and the formation of across-group links is lower.
- P4. In villages with higher $\rho_w - \rho_a$, the association between within-group centrality and the formation of across-group links is lower.

P1-P4 provide a complex set of predictions that we test empirically in the next section. We can think about villages where households experience more variable income, e.g., due to rainfall, or more correlated income and ask whether changes in income variability or correlation correspond to changes in network structure.

7. EMPIRICAL ANALYSIS

We now turn to the data and see if there is support consistent with the predictions developed above.

We make use of a unique and detailed social network dataset from 75 villages in Karnataka, India, which is particularly well-suited for our analysis as (i) it involves numerous independent villages (essential for inference, though most network-based studies have just one or a handful of villages), (ii) it includes *complete* network data across both financial and social connections across almost all households in every village (network-based studies are notoriously subject to measurement error), and (iii) caste is salient in these communities.

7.1. Setting and Data. The data we use were collected in 2011 by Banerjee, Chandrasekhar, Duflo and Jackson (2013a, 2014a), by conducting surveys in 75 villages in Karnataka, India. The villages span 5 districts and range from 2- to 3-hour drive from Bangalore. They are far enough apart to be treated as independent systems (the median distance between them is 46 km, and a district has 1000–3000 villages). The survey included a village questionnaire, a census of all households, demographic covariates (including caste and occupation), as well as data on a number of amenities (e.g., roofing, latrine, or electricity access quality). A detailed individual-level survey was administered to most adults in every village. The survey included a networks module with twelve dimensions of relationships, including financial relationships, social relationships, and advice relationships.⁴⁶

Our analysis focuses on two types of networks: the *financial graph*, L^F , and the *social graph*, L^S . The financial graphs represent risk-sharing connections, and the social graph represents friendships/links used to socialize. We build “AND” networks, which say a link exists if it exists on every dimension being considered (various types of financial connections on the financial network, and various types of social connections on the social network).⁴⁷ The advantage of doing this is that it generates a network structure that is more robust to independent measurement error.⁴⁸ The event that a frivolous link is coded decreases exponentially if we require that it exist across multiple dimensions, which a priori helps in detecting effects when we look differentially across network-type. In some of our empirical analysis, we will explicitly consider how our predictions differentially play out in L^F relative to L^S , as our theory speaks to the former.

Table 1 provides some summary statistics for our data. The average number of households per village is 209 with a standard deviation of 80. The average degree of the financial graph is 3.3, and of the social graph is 3.6. Further, we see that the clustering in the financial graph is 0.19 (0.09), whereas for the social graph it is 0.11 (0.04). The fact that the clustering is

⁴⁶See Banerjee et al. (2014a) for more details. In total we have network data from 89.14% of the 16,476 households based on interviews with 65% of all individuals of age 18 to 55.

⁴⁷We say $l_{ij} \in L^F$ if i goes to j to borrow money in times of need; j goes to i to borrow money in times of need; i goes to j to borrow material goods such as kerosene, rice, or oil in times of need; and j goes to i to borrow material goods in times of need.

⁴⁸We say $l_{ij} \in L^S$ if i goes to j ’s house to socialize or vice versa.

nearly twice as high for the financial graph is of note. This is consistent with the results of, for instance, Jackson et al. (2012) that financial links may need to be supported/embedded in cliques to sustain cooperation. Both the financial and social networks exhibit relatively few cross-caste links. This is seen by looking at the ratio of the probability of having a cross-caste link relative to the probability of having a within-caste link. A within-caste link is three times as likely in the financial graph and less than twice as likely in the social graph. Finally, 67% of households are high caste (GM or OBC). We describe our main outcome variables later.

7.2. Defining Groups, and Boundaries of the Network. As our theory pertains to networks formed from multiple groups, here we make use of caste. Following Munshi and Rosenzweig (2006) and Banerjee et al. (2013), we partition our individuals into two broad caste groups: scheduled caste/scheduled tribes (SC/ST) and general merit/otherwise backward castes (GM/OBC). These are governmental designations used to condition the allocation of, for instance, school seating by caste and reflect a core fissure in the social fabric.

It is worth noting that our model makes simplifying assumptions that do not perfectly fit the data, and they can at most approximate it. One of these is that the linking costs between any two individuals in the same group are the same, and the linking cost between any two individuals from any two different groups are the same. These clearly do not hold exactly in our setting, no matter how we define groups.

An extensive literature documents that in India, caste plays a significant role in informal risk sharing. For instance, Morduch (1991, 1999, 2004) and Walker and Ryan (1990) show that informal insurance functions rather well within a caste (though it still is imperfect), while there is very limited insurance across caste. Morduch (2004) discusses this literature at large, which primarily uses the methodology developed in Townsend (1994) to describe the extent of risk sharing within and across caste groups in a village. More recent work by Munshi and Rosenzweig (2014) as well as Mazzo and Saini (2012) has begun to turn attention to risk-sharing at the level of subcastes, both within and across villages.

Our motivation for choosing caste (as opposed to for example sub-caste) as the defining unit of group comes from multiple sources. First, an extensive sociological literature (see for example, Srinivas (1962)) argues that it is particularly costly for people to form cross-caste links. Second, anecdotal evidence suggests that our chosen caste group division is the principle social fissure in these communities (e.g., evidenced by SC/ST “colonies” or sectors of the village).⁴⁹ Third, defining two large groups allows us to say more direct and precise statements about the network structure within and across each group.⁵⁰

⁴⁹We think focusing on caste in this aggregate way is a reasonable approximation to the setting described in our theory. For example, agents are about 3 times as likely to link within caste block than across caste block. They are also about 3 times as likely to link across subcaste, but within caste, as compared to across caste. On the other hand, they are twice as likely to have a within subcaste link as compared to a within caste, across subcaste link. It is also worth noting that if we partition all households into subcastes, in the typical village over 85% of nodes are members of at most 3 subcastes, along with a smattering of smaller subcastes.

⁵⁰Cutting the sample into very fine groups of can amplify measurement error. With a subset of only 5 or 10 nodes, the notion of who is central and who isn’t is remarkably sensitive to survey measurement error.

It is well-understood in the literature that an important component of risk-sharing comes from members of one’s subcaste, even across villages (see Munshi and Rosenzweig, 2014). From the vantage point of this paper, we take the stance that establishing outside links and outside risk-sharing arrangements takes a longer time than within village ones, because of the lower frequency of interaction with those outside agents. Because of this we take those arrangements given and incorporated into income realizations of agents in the village by the time our agents start interacting. This is of course a simplifying assumption, but we regard it as a reasonable approximation of reality. Restricting attention to risk-sharing within a village is also consistent with what is done in the literature (see, e.g., Townsend (1994), Kinnan (2011), Chiappori, Samphantharak, Schulhofer-Wohl and Townsend (2012), and Ambrus et al. (2014), among others).

7.3. Empirical Strategy. Our predictions are about how network structure varies with σ^2 , ρ_w , and ρ_a . The analysis is observational (not causal) and simply looks at the cross-sectional variation of network structure with these parameters through ordinary least squares (OLS).

We take two approaches to support our empirical claims. First, by focusing on different aspects of network structure (Myerson distance or the composition of cross-caste links), we provide evidence in support of our predictions from very different moments of the data.

Second, we exploit multigraph data. Our theory is built for risk-sharing networks and not for social connections. Exploiting this feature allows us to take a difference-in-differences approach and to study whether the correlations we document come from L^F , the financial graph, as opposed to L^S , the social graph. Since the theory is *differentially more informative* about the structure of financial links as opposed to social links, the difference in patterns across link types is informative. Furthermore, as unobserved endogeneity or homophily is likely to affect several dimensions of the multigraph at once, looking at the difference in risk-sharing link patterns versus social link patterns within a village allows us to address unobserved village-level endogeneity that enters additively through a fixed effect. To take a simple example for a confounder, consider P1. In villages where the weather is more variable, fewer days are suitable for working, so people may spend more time socializing, thereby spuriously generating shorter Myerson distances on average. Our difference-in-differences approach eliminates this sort of confound.

At the same time, we note that our empirical approach is more conservative than similar studies in the literature (e.g., Karlan et al. (2009), Ambrus et al. (2014), Kinnan and Townsend (2012)) in terms of statistical inference. First, these studies typically have very few networks (2, 1, and 16, respectively), and therefore consider node or link-level regressions with standard errors generated at that level. This effectively treats nodes or dyads as independent or loosely correlated, making inference preclude, essentially, village-level shocks. Correlation in any factors outside the model (e.g., incentives to form links for other reasons) as well as equilibrium selection at the village level are all precluded from econometric analyses that don’t study the theory where the entire graph is the unit of observation. We make no

such assumptions on the independence of nodes or dyads for valid statistical inference, and instead exploit the fact that we have 75 independent villages (recall that the median pairwise distance between them is over 46 km). By focusing on village-level variation, we are allowing for arbitrary correlation within graphs. In fact, in our most conservative specifications, we allow correlation at the subdistrict level.

Second, analyses in the previous literature usually do not have access to different types of edges—the multigraph—and therefore cannot employ our difference-in-differences approach. What we do, relative to this, is extremely conservative. By differencing across network types, we are asking whether the patterns in the graph which match our theory are differentially at play for the financial network relative to the social networks.

Third, another reason our approach is conservative is that typical models of multigraph link formation have a fixed-cost component. Thus, incentives driving the formation of risk-sharing links are likely to influence the structure of social or information links through this channel. Observe that by looking at the financial graph relative to the social graph, the variation coming from the fixed-cost component – which is consistent with our theory – is not even being used in our analysis in support of our theory.

In sum, our analysis is observational, but makes use of variation in the economic environment as well as the type of network we focus on, to grapple with more correlated unobservables (e.g., village level fixed unobservables) in a way that previous analysis could not.

7.4. Variable Construction.

7.4.1. Approximating the Myerson distance and centrality. Our next task is to compute the Myerson distance of every pair in every village and the Myerson centrality for all nodes. Unfortunately, this is computationally infeasible for the sample sizes of our data (see Algaba et al. (2007)), presenting a new challenge. Thus, we develop an approximation, described below.

Let $\mathbf{md}(L)$ be the matrix of Myerson distances and define $\mathbf{q}(L) := 1/2 - \mathbf{md}(L)$. So $\mathbf{q}(L)$ is a matrix with the ij th entry capturing the probability that, upon his arrival agent i will not be connected to agent j . It is difficult to directly characterize $\mathbf{md}(L)$ (or equivalently, $\mathbf{q}(L)$) as each village typically consists of around 230 households and the number of candidate paths between each i and j is exponential in the size of the network. Correctly accounting for paths that share nodes is computationally very intensive (see Proposition 4), and it has to be done for all pairs of agents without a direct connection.⁵¹ Instead, we develop a computationally feasible approximation of $\mathbf{md}(L)$, which is exact for trees.

To approximate \mathbf{q} , we use the following idea. The algorithm works by starting with a node, moving to its neighbors, then move to its neighbors' neighbors, and so on, never returning to a previously used node along a given path. This helps us to avoid counting non-minimal

⁵¹Further, due to presumed measurement error (see Banerjee et al. (2013)), there are likely to be missing paths. In fact, the data have occasional disconnected components, and so measures that are precisely based on exact paths or even maximal path lengths are likely to be problematic (Chandrasekhar and Lewis (2014)).

paths. All the while, we keep track of how many ways we have moved from the original node to any given node. We denote our approximation of \mathbf{q} by $\hat{\mathbf{q}}$.




FIGURE 4. The nodes i, j for which we are computing $md(i, j, L)$ have purple stripes. The tree contains a single minimal path (solid orange nodes), whereas the circle contains two paths (solid orange nodes and chequered blue nodes).

The inclusion–exclusion principle weights paths that are longer less and a path that shares many nodes with another less. With this in mind, we choose the following two approximation strategies. Let the shortest path between two nodes be of length l . We first count the paths of length l and length $l + 1$. We then count paths of length $l + 2$.⁵² If there are fewer than k such paths, we use them all. Otherwise, we consider only the k shortest and in practice we set $k = 4$.⁵³ Discarding longer paths in this way biases downwards our approximation of \mathbf{q} . As we cannot keep track of exactly which nodes feature in each path, we also have to make an assumption about the overlap of nodes in order to apply the inclusion–exclusion principle to these paths. Each path must share the same first and last node. We perform the inclusion–exclusion principle assuming that only these nodes are shared (see Section 4). Assuming no other nodes are shared introduces a second bias, but this time upwards in our approximation.

To explain these concepts, we provide some illustrations. Figure 4 presents two examples: a tree and a circle. The tree has a single minimal path between nodes 1 and 8, whereas the circle has two minimal paths between nodes 1 and 4. Figure 5 shows how links are removed for the case of a tree. Once a node has been reached, links back into that node are deleted before the nodes neighbors are “infected.” This ensures only minimal paths are included in the calculation.

In the case of the circle shown in Figure 4, our algorithm is also exact for paths between 1 and 4. There are two minimal paths (which in this case are both shortest paths too), and we find both in the initial run of our algorithm. Following the inclusion–exclusion principle, we add $1/4$ to $1/4$ and subtract $1/6$. In this case our assumption that the two paths share only two nodes is accurate. We are also exact for paths between 1 and 3, but in this case there is

⁵²Counting more paths greatly (exponentially) increases the running time of our algorithm.

⁵³We need a fixed (small) truncation. Otherwise both the memory requirements and the run-time of the algorithm grow exponentially. Results are not sensitive to the truncation point.




FIGURE 5. As the algorithm progresses, directed links into nodes that are reached are deleted. This ensures that only minimal paths are included. In this case, as in all tree networks, there is a unique minimal path from A to B.

a path of length $l + 2$. To find this path, we look for paths of length l from 1 to nodes other than 3. In this case there is one such path to node 5. We then look for paths from 5 to 3 that pass through one other node. There is one such path and so the calculation we perform is: $1/3 + 1/5 - 1/6$. While we are accurate for all pairs of nodes in the circle shown, in larger circles we will miss the longer paths.

The following algorithm finds the length of the shortest path between two nodes, how many paths of that length there are and how many paths there are that are one longer. From this information, we also find paths of length $l + 2$.⁵⁴

Algorithm 12 (Incoming Link Deletion). *Let e^i be the i th basis vector. This will represent the root (starting) node. Initialize $\hat{\mathbf{q}} = \text{zeros}(n, n)$, a matrix of zeros. Initialize $z^{t,i} = \text{zeros}(n, 1)$ and $x^{t,i} = \text{zeros}(n, 1)$ to be n -vectors of zeros, indexed by $i = 1, \dots, n$ and $t = 1, \dots, T$. Repeat steps 1–4 for each of (e^1, \dots, e^n) .*

(1) *Period 1: There is no identification or updating steps.*

(a) *Percolation: $x^{1,i} = \mathbf{A}e^i$.*

(Identifies who is connected to the root node)

(2) *Period 2, given $(x^{1,i}, \mathbf{A})$:*

(a) *Identification: $z^{2,i} = e^i$.*

(b) *Update graph:⁵⁵ $\mathbf{A}_2 = \text{zeros}(n, n)$, $\mathbf{A}_2(\neg z^{2,i}, :) = \mathbf{A}(\neg z^{2,i}, :)$.*

(Deletes links into the root node)

(c) *Percolation: $x^{2,i} = \mathbf{A}_2 x^{1,i}$.*

(Records number of paths from root node to other nodes passing through one other)

(3) *Period t , given $(x^{t-1,i}, \mathbf{A}_{t-1})$:*

⁵⁴For paths from i to j , this is done by looking at paths of length l to agents other than j , and then looking at paths from these agents to j .

⁵⁵Let $\mathbf{A}(:, v)$ denote $(A(1, j), \dots, A(n, j))$.

- (a) *Identification:* $z^{t,i} = \mathbf{1} \{ \sum_{s=3}^t x^{s-2,i} > 0 \}$.
(Identifies nodes already visited)
- (b) *Update graph:* $\mathbf{A}_t = \text{zeros}(n, n)$, $\mathbf{A}_t(\neg z^{t,i}, :) = \mathbf{A}_{t-1}(\neg z^{t,i}, :)$.
(Deletes links into all nodes that have already been visited)
- (c) *Percolation:* $x^{t,i} = \mathbf{A}_t x^{t-1,i}$.

By construction $x_j^{t,i}$, the j th entry of $x^{t,i}$, records paths from i to j that pass through t nodes. If t' is the lowest t with a positive entry in this matrix, then the shortest path from i to j passes through t' nodes. In this case, $x_j^{t',i}$ tells us how many such paths there are and $x_j^{t'+1,i}$ tells us how many paths there are that pass through one more node. However, by construction $x_j^{t'+k,i} = 0$ for all $k > 1$ and longer paths are not recorded. This is because the incoming links to node j will have been deleted by this step of the algorithm. Deletion of incoming links helps prevent non minimal paths from being recorded. Using this information for all seed nodes, the number of paths of length $t' + 2$ between i and j are also found as described above. The inclusion-exclusion principle is then applied to this combined set of paths, assuming each path shares only the first and last nodes, to calculate $\hat{\mathbf{q}}(L)$.

Proposition 13. *Let L be a tree. Then $\hat{\mathbf{q}}(L) = \mathbf{q}(L)$.*

See the Appendix B for the proof. To operationalize $\hat{\mathbf{q}}(L)$ in our regression analysis, we need a village-level measure of Myerson distances. We use $\tilde{q}(L) := \sum_{i < j} \hat{q}_{ij} / \binom{n}{2}$ which measures an appropriately weighted density of the network. Finally, to approximate Myerson centrality we use $\sum_j \hat{q}_{ij}$, as people are central when they are likely to be connected to others. Thus, their q_i terms are high (equivalently, they have a low Myerson distance to others).

A limitation of the Incoming Link Deletion algorithm is that longer paths are excluded. To address this, we construct an alternative algorithm. This Outgoing Link Deletion algorithm is identical to the one described, except that it deletes outgoing links instead of incoming links. The Outgoing Link Deletion algorithm finds longer paths, and does an especially good job of picking up longer paths that share few nodes with other paths. However, it also includes additional short non-minimal paths and is not exact for tree networks. As longer paths are found, we directly use the output of the algorithm without constructing any additional longer paths. Nevertheless, for the set of paths we find, it is computationally infeasible to compute the Myerson distances using the inclusion-exclusion principle. Censoring these paths would defeat the point of the Outgoing Link Deletion algorithm. Instead, we use an approximation of the inclusion-exclusion principle which makes the computation much simpler. This approximation treats every path as completely independent, assuming that no nodes are shared (even though we know at two must be). For example, if we find 3 paths from i to j that pass through l nodes, l' nodes and l'' nodes respectively, our approximation of q_{ij} will be $1/l + 1/l' + 1/l''$.

7.4.2. Construction of income variability and caste-income correlation. Our analysis requires estimates of σ^2 and $(\rho_w - \rho_a)$. However, due to data limitations—despite extremely detailed

multigraph network data across many villages, we lack financial records—we need to construct measures of these quantities.

For income variability, we merge National Oceanic and Atmospheric Administration (NOAA) data with our network data. Matsuura and Willmott (2012) construct a gridded monthly time series of terrestrial precipitation from 1900 to 2010. We match this to our villages using our GPS data. Once month fixed-effects are removed, the crucial variable is the standard deviation of rainfall by village.

Measuring income correlation is difficult. Ideally, we would have time-series data on incomes of all households, as well as plausible instruments allowing us to calculate the exogenous variation in income correlation both within and across castes for each village. While we do not have access to income data, we do have detailed data on occupation.⁵⁶ Thus, we make use of the relative within-caste to across-caste occupation correlation. The main idea is that shocks to individuals will be more highly correlated when they have the same occupation. We take two approaches to computing caste–income correlation.

One approach is to look at the correlation of being in the high-caste group with holding a given occupation, for all occupations in our survey. We then take the weighted average of these correlations, where the weight is the share of agents in the occupation. Thus, for a given village we consider

$$(16) \quad \widehat{\rho_w - \rho_a} := \sum_{k=1}^K \text{corr}(Caste, Occupation_k) \Pr(Occupation_k),$$

where $Caste$ is a $1 \times n$ vector of GM/OBC dummies and $Occupation_k$ is an $n \times 1$ vector of dummies for a household having a member in occupation k . This constructs a score which is 0 if there is no correlation between caste group and occupation, and 1 if caste group perfectly predicts occupational choices.⁵⁷

Another approach involves making simplifying assumptions about the structure of the income process. If the income of an individual in a given caste and occupation can be thought of as depending on a (caste–occupation)-specific mean, occupation-specific idiosyncratic iid shocks, and individual-specific idiosyncratic iid shocks, where the occupation shocks all have the same variance across occupations, then as shown in Appendix C, we can write

$$(17) \quad \widehat{\rho_w - \rho_a} := \sum_g \phi_g \Pr(o_{i,g} = o_{j,g}) - \Pr(o_{i,A} = o_{j,B}),$$

⁵⁶We recognize that occupations can have a choice component. Nevertheless, we proceed with these measures for three reasons. First, in rural villages the primary household occupation (agriculture or sericulture) is often passed on through generations. Second, it is possible to show that under a natural model with endogenous selection into occupation, the within- versus across-group income correlations are captured by our occupational choice measures, whereas the choice of occupation does not generate spurious correlations with the network in a manner consistent with P1–P4. Finally, this is the best possible approximation, given the severe income data limitations, as a necessary component of our analysis is the network data.

⁵⁷Unsurprisingly, the analogous variable constructed at the sub-caste level has a correlation of 0.94 with the one constructed using caste.

where $o_{i,g}$ denotes the occupation of i in caste $g \in \{A, B\}$ and ϕ_g is the population share of caste g . In sum, both measures are intuitive, but imperfect, proxies for the regressors we need in our analysis. Therefore we take a rough-and-ready approach, utilizing both in our analysis:

$$(1) \widehat{\rho_w - \rho_a}^I = \sum_g \phi_g \Pr(o_{i,g} = o_{j,g}) - \Pr(o_{i,A} = o_{j,B}).$$

$$(2) \widehat{\rho_w - \rho_a}^{II} = \sum_{k=1}^K \text{corr}(Caste, Occupation_k) \Pr(Occupation_k).$$

7.5. Results.

7.5.1. *Myerson distance as a function of income variability and correlation.* We begin with P1 and P2. Figure 6 presents results in the raw data. As predicted, villages variable individual incomes are associated with lower $\widetilde{md}(L^F)$, and villages with more within-caste income correlation are associated with higher $\widetilde{md}(L^F)$.

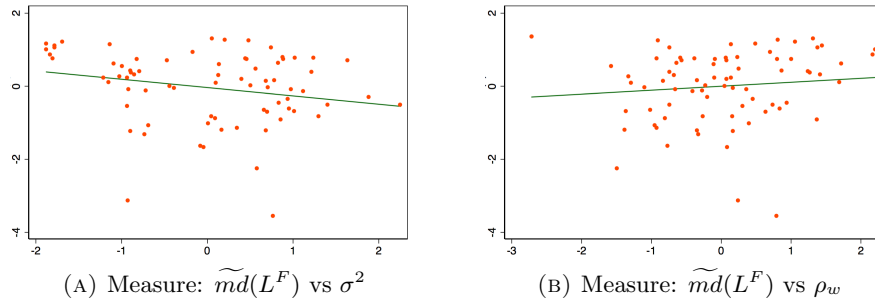


FIGURE 6. Myerson distance in standardized units against the variability of rainfall at the village level and the within-caste income correlation metric, both in standardized units.

Table 2 demonstrates the robustness of this graphical evidence in regression of $\widetilde{md}(L^F)$ on σ , ρ_w , and with district or subdistrict fixed effects as well as controls for caste composition. In all specifications here and throughout the paper, unless otherwise noted we use a Wild clustered bootstrap to account for subdistrict-level clustering in our inference (Cameron et al., 2008).⁵⁸

Columns 1-3 present regressions using $\widehat{\rho_w}^I$, whereas columns 4-6 present regressions using $\widehat{\rho_w - \rho_a}^{II}$. Panel A presents results using the Incoming Link Deletion algorithm whereas Panel B looks at Outgoing Link Deletion. We find a one standard deviation increase in Measure I is associated with a 0.292 or 0.296 standard deviation increase in $\widetilde{md}(L^F)$ (column 1, Panels A and B). However, there is no significant relationship between income variability and $\widetilde{md}(L^F)$ in these specifications. Column 4 presents the analogous results using $\widehat{\rho_w - \rho_a}^{II}$.

⁵⁸While these villages are essentially independent units, the median distance being 46 km apart, we take a conservative approach because geography is a determinant of rainfall variability and occupation. Our villagers are members of 12 subdistricts, and we therefore cluster our standard errors at that level. To deal with the finite sample bias, we use a Wild cluster bootstrap procedure for the t -test statistic, using Rademacher weights, to generate p values for hypothesis testing as well as standard errors.

A one-standard-deviation in Measure II is associated with a 0.205 or 0.199 standard deviation increase in $\widetilde{md}(L^F)$, but income variability has no detectable correlation with the outcome variable (column 4, Panels A and B). In columns 2 and 5 we add district fixed effects, and the results are mostly reflective of those in columns 1 and 3. A one standard deviation increase in income correlation corresponds to a roughly 0.1 or 0.18 standard deviation increase in Myerson distance, using Measure I (p -values 0.477 and 0.12 in Panels A and B) and Measure II (p -values 0.00 and 0.012 in Panels A and B) respectively. Inclusion of district fixed effects resolves some of the noise around the rainfall variability coefficient as well in the case of Panel A. Finally, when we include subdistrict fixed effects, our estimate under Measure I is too noisy to distinguish from zero, though we find similar evidence under Measure II (a 0.1 standard deviation effect, p -value 0.04 and 0.13 in Panels A and B). Further, when comparing villages within subdistricts we are able to identify an income variability effect. A one standard deviation increase in income-variability is associated with a 0.3 standard deviation decrease in the average Myerson distance (columns 3 and 6, respectively, of both panels).

Taken together, we show an increase in within-caste income correlation is associated with a change in Myerson distance in a manner consistent with our theory, though the evidence for income variability is considerably weaker.

7.5.2. Differences by network type. Next, we use a difference-in-differences approach and see whether the effects we are interested in are coming differentially from the financial graph as opposed to the social graph. Figure 7 presents the raw data, differenced, graphically. Villages in areas corresponding to more correlated within-caste income processes are associated with greater $\widetilde{md}(L^F) - \widetilde{md}(L^S)$. Similarly, an increase in income variability is associated with a differential decrease in the measure of network density in the financial graph as compared to the social graph.

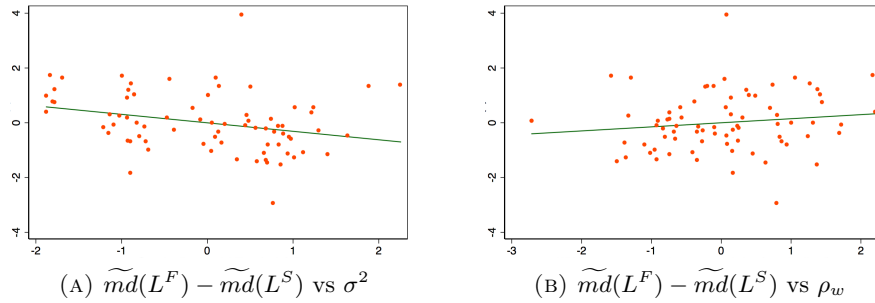


FIGURE 7. Myerson distance, differenced across network type, versus income variability and within-caste income correlation

Let v index village and $t \in \{F, S\}$ index network type. We use the following regression:

$$(18) \quad Y_{v,t} = \alpha + \beta \sigma_v \mathbf{1}_{\{t=F\}} + \gamma \rho_{w,v} \mathbf{1}_{\{t=F\}} + \mu_v + \mu_t + \delta X_v \mathbf{1}_{\{t=F\}} + \epsilon_{v,t},$$

where $Y_{v,t} = \widetilde{md}(L_v^t)$, μ_v is a possibly endogenous village fixed effect, μ_t is a network type fixed effect, and X_v are village-level controls. P1 and P2 correspond to $\beta < 0$ and $\gamma > 0$, since our theory pertains only to risk sharing networks. Thus, our effects should be differentially more predictive for the financial network than for the social network and should remain true when looking at relative effects.

Table 3 presents the results. Panel A presents the Incoming Link Deletion algorithm and Panel B the Outgoing Link Deletion algorithm. We include results for both measures of income correlation. We find that a one standard deviation increase in income variability differentially decreases $\widetilde{md}(L)$ by about 0.24 standard deviations more in financial networks than in social networks (columns 1-4), irrespective of the measure of income correlation used. We are unable to detect a statistically significant association between rainfall variability and the outcomes of interest.

Thus, even when we look within villages, by allowing for village fixed effects and yet allowing for subdistrict-level correlation in our error terms, we see that the financial graph behaves in a manner more consistent with the theory than the social graph.

7.5.3. Association between within-caste centrality and cross-caste links. We now look at how the composition of cross-caste links varies with these parameters. P3 says that in villages with higher variability we should see a greater association between within-caste centrality and having an across-caste link. P4 says that a similar effect occurs when we look at within-versus across-caste income correlation.

We present regressions where we look at the financial network only on σ_v and $\rho_w - \rho_a$ as well as difference-in-difference results. To make this concrete, the latter regression is

$$(19) \quad Y_{v,t} = \alpha + \beta \sigma_v \mathbf{1}_{\{t=F\}} + \gamma (\rho_{w,v} - \rho_{a,v}) \mathbf{1}_{\{t=F\}} + \mu_v + \mu_t + \delta X_v \mathbf{1}_{\{t=F\}} + \epsilon_{v,t},$$

where $Y_{v,t}$ is the average within-caste centrality of nodes with a cross-caste link in network t . The remaining terms are as before. P3 and P4 imply $\beta < 0$ and $\gamma < 0$.

Table 4 presents the results. Panel A shows regressions using just the financial networks on regressors of interest. Panel B displays results from the difference-in-difference. Columns 1–3 use the Incoming Link Deletion algorithm, while columns 4–6 use the Outgoing Link Deletion algorithm to compute the outcome of interest.

The main result is that we find a one standard deviation increase in the measure of the $\rho_w - \rho_a$, irrespective of which measure is used, corresponds to about a 0.15 standard deviation decline in the average Myerson centrality (within caste) of the individuals with across caste links. Again, however, when we look at rainfall variability, we are unable to confirm or reject our theory, as estimates are very noisy.

Thus, our evidence is partially in support of our theory. We find a negative association between income correlation and the centrality of those with a cross-caste link. However, we are unable to confirm or reject our hypothesis when looking at the variability of income.

7.6. Discussion. This section provided suggestive evidence, consistent with our theory, demonstrating the following:

- (1) Networks have higher Myerson distances when within-group income correlation is higher.
- (2) Although we only find at best mild evidence that networks exhibit lower Myerson distances when there is greater variability in rainfall, the rainfall data is noisy and there is no evidence of the opposite relationship.
- (3) Networks exhibit more negative associations between within-group centrality and across-group linking when within-group income correlation is higher, though we detect no relationship with variability of rainfall.
- (4) Results (1)–(3) are robust to a difference-in-differences approach. By differencing out the social graph, we can remove village-fixed endogenous factors and more convincingly argue that the underlying force concerns risk sharing.

Having several distinct predictions is useful, as we are conducting an observational analysis. We find that we cannot reject any of these predictions based on our data, and in fact, given the data limitations we view our results as surprisingly supportive of the theory.

Of course, there are several other stories consistent with the income-variability prediction. For instance, one may think that any theory where the returns to investing in risk-sharing relationships go up should predict networks of greater density. We note, however, that this prediction does not hold in Bramoullé and Kranton (2007a), which is perhaps the closest paper to ours. In their model, increased variability has no effect on the stable network structures.

Furthermore, our compositional story is new. More generally, P3–P4 are specific to either risk-sharing relationships between castes or differences in income correlations between castes, and so within the literature are unique to our model. These predictions are also more subtle. For instance, higher within-group correlation leading to more sparse risk-sharing networks is not a conclusion of other risk-sharing network formation stories.

We again emphasize that our findings are clearly observational (non-causal) and subject to measurement error. However, given the unique dataset we used—network data for 75 independent villages, including multigraph data that allows us to difference out by link-type to see if effects are driven by those consistent with our theory—this represents a first opportunity to tackle the types of questions we address. The data requirements are immense, and we are able to take a serious pass at looking at different types of cuts of our data—several network measures, several parameters of the economic environment, and several different results from our theory—to see if the data are consistent with our story. Additionally, we are extremely conservative in conducting statistical inference, by relying only on independence across villages (in fact, we are employing an even more conservative approach than that by using a Wild cluster bootstrap at the subdistrict level) and not exploiting any of the within-village observations (which are likely to be correlated).

8. CONCLUSION

We develop a relatively tractable model of network formation and surplus division in a context of risk sharing that allows for heterogeneity in correlations between the incomes of pairs of agents. Such correlations have a sizeable impact on the potential of informal risk sharing to smooth incomes. We investigate the incentives for relationships that enable risk sharing to be formed both within a group (caste or village) and across groups, giving access to less correlated income streams. We find that overinvestment into social relations is likely within a group, but there is potential underinvestment into more costly social connections that bridge different groups. We also find a novel trade-off between equality and efficiency.

Using a unique dataset of 75 Indian villages, we find empirical support for our model. First, we find that within group income correlation and, to a lesser extent, income variability are associated with the Myerson distance of the social network in a manner consistent with our theory. Second, more centrally connected individuals are more likely to establish across-group links when looking at networks with stronger within-group income correlation relative to across-group income correlation. Third, each these relationships are differentially stronger for the network of financial relationships than for the network of other types of social relationships. The evidence, while descriptive, is congruent with our model.

Although we focus our analysis on risk sharing, our conclusions regarding network formation could apply in other social contexts too, as long as the economic benefits created by the social network are distributed similarly to the way they are in our model—a question that requires further empirical investigation. Within the context of risk sharing, a natural next step would be to provide a dynamic extension of the analysis that allows for autocorrelation between income realizations.

REFERENCES

- Ahuja, G. (2000): "Collaboration networks, structural holes, and innovation: A longitudinal study," *Administrative science quarterly*, 45, 425–455.
- Algaba, E., JM Bilbao, J. Fernández, N. Jiménez and J. López (2007): "Algorithms for computing the Myerson value by dividends," *Discrete Mathematics Research Progress*, 1-13.
- Ali, N. and D. Miller (2013): "Enforcing cooperation in networked societies," mimeo UC San Diego.
- Ali, N. and D. Miller (2013): "Ostracism," mimeo UC San Diego.
- Ambrus, A., M. Mobius, and A. Szeidl (2014): "Consumption risk-sharing in social networks," *American Economic Review*, 104, 149-182.
- Angelucci, M. and G. De Giorgi (2009): "Indirect effects of an aid program: How do cash injections affect non-eligibles' consumption?," *American Economic Review*, 99, 486–508.
- Arrow, K. (1965): "Aspects of the theory of risk-bearing (Yrjo Jahnsson Lectures)," Yrjo Jahnssonin Saatio, Helsinki.
- Atkinson, A. B. (1970): "On the Measurement of Inequality," *Journal of Economic Theory*, 2, 244-263.
- Attanasio, O., A. Barr, J. Cardenas, G. Genicot and C. Meghir (2012): "Risk pooling, risk preferences, and social networks," *AEJ - Applied Economics*, 4, 134-167.
- Austen-Smith, D. and R. Fryer (2005): "The economic analysis of acting white," *Quarterly Journal of Economics*, 120, 551-583.
- Bala, V. and S. Goyal (2000): "A noncooperative model of network formation," *Econometrica*, 68, 1181-1229.
- Banerjee, A., A. Chandrasekhar, E. Duflo and M. Jackson (2013a): "The diffusion of micro-finance," *Science*, 341, No. 6144 1236498
- Banerjee, A., E. Duflo, M. Ghatak, and J. Lafortune (2013b): "Marry for what? Caste and mate selection in modern India." *American Economic Journal: Microeconomics*, 5(2): 33-72.
- Banerjee, A., A. Chandrasekhar, E. Duflo and M. Jackson (2014a): "Gossip: Identifying central individuals in a social network," MIT and Stanford working paper.
- Banerjee, A., A. Chandrasekhar, E. Duflo and M. Jackson (2014b): "Exposure to formal loans: Changes in social networks," MIT and Stanford working paper.
- Banerjee, A., E. Breza, E. Duflo and C. Kinnan (2014c): "Does microfinance foster business growth? The importance of entrepreneurial heterogeneity," mimeo MIT.
- Bennett, E. (1988): "Consistent Bargaining Conjectures in Marriage and Matching," *Journal of Economic Theory*, 45, 392-407.
- Bergstrom, T. and H. Varian (1985): "When do markets have transferable utility?," *Journal of Economic Theory*, 35, 222-233.
- Billand, P., Bravard, C. and S. Sarangi (2012): "Mutual insurance networks and unequal resource sharing in communities," DIW Berlin Discussion Paper.

- Binzel, C., E. Field and R. Pande (2014): "Does the arrival of a formal financial institution alter informal sharing arrangements? Evidence from India," mimeo Harvard University.
- Bloch, F., G. Genicot, and D. Ray (2008): "Informal insurance in social networks," *Journal of Economic Theory*, 143, 36-58.
- Borch, K. (1962): "Equilibrium in a reinsurance market," *Econometrica*, 30, 424-444.
- Bramoullé, Y. and R. Kranton (2007a): "Risk-sharing networks," *Journal of Economic Behavior and Organization*, 64, 275-294.
- Bramoullé, Y. and R. Kranton (2007b): "Risk sharing across communities," *American Economic Review Papers & Proceedings*, 97, 70.
- Burt, Y. (1992): "Structural holes," Academic Press, New York.
- Calvo-Armengol, A. (2001): "Bargaining power in communication networks," *Mathematical Social Sciences*, 41, 69-87.
- Calvo-Armengol, A. (2002): "On bargaining partner selection when communication is restricted," *International Journal of Game Theory*, 30, 503-515.
- Calvo-Armengol, A. and R. Ilkilic (2009): "Pairwise-stability and Nash equilibria in network formation," *International Journal of Game Theory*, 38, 51-79.
- Cameron, A., J. Gelbach and D. Miller (2008): "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 90(3), 414-427.
- A. Chandrasekhar and R. Lewis (2014): "Econometrics of sampled networks," mimeo Stanford University.
- Cole, H., and N. Kocherlakota (2001): "Efficient allocations with hidden income and hidden storage," *The Review of Economic Studies* 68, 523-542.
- Coate, S. and M. Ravallion (1993): "Reciprocity without commitment: Characterization and performance of informal insurance arrangements," *Journal of Development Economics*, 40, 1-24.
- Corominas-Bosch, M. (2004): "Bargaining in a network of buyers and sellers," *Journal of Economic Theory*, 115, 35-77.
- Elliott, M. (2014): "Inefficiencies in networked markets," *AEJ Microeconomics*, forthcoming.
- Fafchamps, M. (1992): "Solidarity networks in preindustrial societies: Rational peasants with a moral economy," *Economic Development and Cultural Change*, 41, 147-74.
- Fafchamps, M. and F. Gubert (2007): "The formation of risk sharing networks," *Journal of Development Economics*, 83, 326-350.
- Feigenberg, B., E. Field and R. Pande (2013): "The economic returns to social interaction: Experimental evidence from micro finance," *Review of Economic Studies*, 80, 1459-1483.
- Fontenay, C. and J. Gans (2013): "Bilateral bargaining with externalities," *Journal of Industrial Economics*, forthcoming.
- Glaeser, E., D. Laibson, and B. Sacerdote (2002): "An economic approach to social capital," *The Economic Journal*, 112, F437-F458.

- Goyal, S. and F. Vega-Redondo (2007): "Structural holes in social networks," *Journal of Economic Theory*, 137, 460–492.
- Hojman, D. and A. Szeidl (2008): "Core and periphery in networks," *Journal of Economic Theory*, 139, 295–309.
- Jackson, M. (2003): "The stability and efficiency of economic and social networks," in *Advances in Economic Design*, edited by Murat Sertel and Semih Koray, Springer-Verlag, Heidelberg.
- Jackson, M. (2005): "Allocation rules for network games," *Games and Economic Behavior*, 51, 128–154.
- Jackson, M. (2010): "Social and economic networks," Princeton University Press.
- 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.
- Jackson, M. and A. Wolinsky (1996): "A strategic model of social and economic networks," *Journal of Economic Theory*, 71, 44–74.
- Karlan, D., M. Mobius, T. Rosenblat and A. Szeidl (2009): "Trust and social collateral," *Quarterly Journal of Economics*, 124, 1307–1361.
- Kets, W., G. Iyengar, R. Sethi and S. Bowles (2011): "Inequality and network structure," *Games and Economic Behavior*, 73, 215–226.
- Kinnan, C. (2011): "Distinguishing barriers to insurance in Thai villages," mimeo Northwestern University.
- Kinnan, C. and R. Townsend (2012): "Kinship and financial networks, formal financial access, and risk reduction," *The American Economic Review*, 102, 289–293.
- Kranton, R. and D. Minehart (2001): "A theory of buyer-seller networks," *American Economic Review*, 91, 485–508.
- Ligon, E., J. Thomas, and T. Worrall (2002): "Informal insurance arrangements with limited commitment: Theory and evidence from village economies," *Review of Economic Studies*, 69, 209–244.
- Manea, M. (2011): "Bargaining in stationary networks," *American Economic Review*, 101, 2042–80.
- Matsuura, K. and C.J. Willmott (2012): "Terrestrial precipitation: 1900–2008 gridded monthly time series," Center for Climatic Research Department of Geography Center for Climatic Research, University of Delaware.
- Mazzocco, M. and S. Saini (2012): "Testing efficient risk sharing with heterogeneous risk preferences," *The American Economic Review*, 102.1, 428–468.
- Mehra, A., Kilduff, M. and D. Bass (2001): "The social networks of high and low self-monitors: implications for workplace performance," *Admin. Sci. Quart.* 46 (2001) 121–146.
- M. Morduch (1991): "Consumption smoothing across space: Tests for village-level responses to risk." Harvard University.

- M. Morduch (1999): "Between the state and the market: Can informal insurance patch the safety net?" *The World Bank Research Observer* 14.2, 187–207.
- Munshi, K. and M. Rosenzweig (2006): "Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy," *The American Economic Review*, 1225-1252.
- Munshi, K. and M. Rosenzweig (2009): "Why is mobility in India so low? Social insurance, inequality, and growth," *National Bureau of Economic Research*, No. w14850.
- Myerson, R. (1977): "Graphs and cooperation in games," *Mathematics of Operations Research*, 2, 225-229.
- Myerson, R. (1980): "Conference structures and fair allocation rules," *International Journal of Game Theory*, 9.3, 169-182.
- Myerson, R. (1991): "Game theory: Analysis of conflict, Harvard University Press, Cambridge.
- Navarro, N. and A. Perea, (2013): "A simple bargaining procedure for the Myerson value," *The BE Journal of Theoretical Economics*, 13, 1-20.
- Pin, P. (2011): "Eight degrees of separation," *Research in Economics*, 65, 259-270.
- Podolny, J., and J. Baron, (1997): "Resources and relationships: Social networks and mobility in the workplace," *American Sociological Review*, 673-693.
- Rosenzweig, M. (1988): "Risk, implicit contracts and the family in rural areas of low-income countries," *Economic Journal*, 98, 1148-70.
- Slikker, M. (2007): "Bidding for surplus in network allocation problems," *Journal of Economic Theory*, 137, 493-511.
- Srinivas, M. (1962): "Caste in modern India and other essays," Bombay, India: Asia Publishing House.
- Stole, L. and J. Zwiebel (1996): "Intra-firm bargaining under non-binding contracts," *Review of Economic Studies*, 63, 375-410.
- Townsend, R. (1994): "Risk and insurance in village India," *Econometrica*, 62, 539-591.
- Udry, C. (1994): "Risk and insurance in a rural credit market: An empirical investigation in Northern Nigeria," *Review of Economic Studies*, 61, 495-526.
- van Lint, J. H., and R. Wilson (2001): "A course in combinatorics," Cambridge university press.
- Walker, T. and J. Ryan (1990): "Village and household economics in India's semi-arid tropics," Johns Hopkins University Press, Baltimore.
- Wilson, R. (1968): "The theory of syndicates," *Econometrica*, 36, 119-132.

TABLE 1. Summary Statistics

	Mean	SD
Households per village	209.27	80.03
Average degree (financial network)	3.28	0.91
Average clustering (financial network)	0.19	0.05
Probability of cross-caste link/ Probability of within-caste link (financial network)	0.35	0.30
Average degree (social network)	3.55	1.15
Average clustering (social network)	0.11	0.04
Probability of cross-caste link/ Probability of within-caste link (social network)	0.62	0.49
Fraction high caste (GM/OBC)	0.67	0.15

Notes: For each variable we present the mean and the cross-village standard deviation. For the Myerson distances we use the approximation algorithm developed in the paper

TABLE 2. Myerson distance vs. income variability and correlation

	Measure I			Measure II		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Incoming Link Deletion Algorithm</i>						
Income Correlation	0.292 (0.137) [0.046]	0.086 (0.106) [0.477]	-0.006 (0.108) [0.935]	0.199 (0.092) [0.034]	0.180 (0.055) [0.000]	0.150 (0.079) [0.044]
Income Variability	0.009 (0.126) [0.96]	-0.217 (0.131) [0.154]	-0.278 (0.097) [0.002]	0.015 (0.158) [0.951]	-0.278 (0.131) [0.072]	-0.315 (0.094) [0.002]
District FE	N	Y	Y	N	Y	Y
Subdistrict FE	N	N	Y	N	N	Y
R-squared	0.1335	0.4368	0.6728	0.0891	0.4637	0.6924
<i>Panel B: Outgoing Link Deletion Algorithm</i>						
Income Correlation	0.296 (0.109) [0.004]	0.109 (0.061) [0.12]	-0.001 (0.064) [0.996]	0.205 (0.088) [0.03]	0.182 (0.069) [0.012]	0.098 (0.063) [0.13]
Income Variability	-0.061 (0.187) [0.82]	0.015 (0.239) [0.928]	-0.281 (0.094) [0.002]	-0.054 (0.206) [0.774]	-0.054 (0.224) [0.794]	-0.304 0.101 [0.002]
District FE	N	Y	Y	N	Y	Y
Subdistrict FE	N	N	Y	N	N	Y
R-squared	0.1066	0.5701	0.7458	0.071	0.5897	0.7525

Notes: Outcome variable is the average Myerson distance, computed by the approximation algorithms as indicated. Panel A uses the incoming link deletion algorithm and Panel B uses the outgoing link deletion algorithm. Outcome variables and regressors scaled by their standard deviations. Columns (1-3) use Measure I of the caste-income correlation, whereas columns (4-6) use Measure II. Specifications include control for caste composition: $p(1-p)$, where p is share in high caste. Wild clustered bootstrap standard errors are presented in (.), using Rademacher weights. The cluster is at the subdistrict level, of which there are 12. 1000 samples are used per bootstrap. p -values from the Wild clustered bootstrap t are presented in [.]

TABLE 3. Myerson distance vs. income variability and correlation, with village FE

	<i>Panel A: Incoming Link Deletion Algorithm</i>		<i>Panel B: Outgoing Link Deletion Algorithm</i>	
	Measure I	Measure II	Measure I	Measure II
	(1)	(2)	(3)	(4)
Income Correlation x $\mathbf{1}\{\text{Financial network}\}$	0.242 (0.104) [0.012]	0.177 (0.115) [0.188]	0.225 (0.084) [0.010]	0.289 (0.115) [0.016]
Income Variability x $\mathbf{1}\{\text{Financial network}\}$	-0.070 (0.107) [0.564]	-0.057 (0.115) [0.712]	-0.151 (0.178) [0.446]	-0.131 (0.171) [0.502]
R-squared	0.8353	0.8297	0.7362	0.7452

Notes: Outcome variable is the average Myerson distance of the graph of a given type (financial or social), computed using the specified algorithm (incoming or outgoing link deletion). Outcome variable, income correlation (either Measure I or Measure II), and income variability all scaled by their standard deviations. With two observations per village, one for each network type, all regressions include village fixed effects (and therefore subsume district and subdistrict fixed effects as well). Specifications include control for caste composition: $p(1-p)$, where p is share in high caste, interacted with network type. Wild clustered bootstrap standard errors are presented in (.), using Rademacher weights. The cluster is at the subdistrict level, of which there are 12. 1000 samples are used per bootstrap. p -values from the Wild clustered bootstrap t are presented in [.]

TABLE 4. Association between average within-caste centrality of nodes with cross-caste links and measures of rainfall variability and within-caste income correlation

	<i>Incoming Link Deletion Algorithm</i>		<i>Outgoing Link Deletion Algorithm</i>	
	Measure I	Measure II	Measure I	Measure II
	(1)	(2)	(3)	(4)
<i>Panel A: Using just the Financial network</i>				
Income Correlation	-0.231 (0.098) [0.018]	-0.196 (0.142) [0.157]	-0.131 (0.061) [0.032]	-0.101 (0.097) [0.28]
Income Variability	0.009 (0.183) [0.961]	-0.019 (0.157) [0.918]	0.086 (0.185) [0.636]	0.093 (0.191) [0.609]
R-squared	0.2265	0.1992	0.127	0.111
<i>Panel B: Difference in Differences using the Financial and Social networks</i>				
Income Correlation x $\mathbf{1}\{\text{Financial network}\}$	-0.191 (0.088) [0.030]	-0.228 (0.150) [0.12]	-0.121 (0.095) [0.200]	-0.263 (0.118) [0.026]
Income Variability x $\mathbf{1}\{\text{Financial network}\}$	0.056 (0.162) [0.727]	0.060 (0.166) [0.717]	0.109 (0.188) [0.564]	0.104 (0.177) [0.556]
R-squared	0.7956	0.7954	0.765	0.772

Notes: Outcome variable is the average within-caste average Myerson centrality, computed using our approximation algorithm (incoming link deletion or outgoing link deletion), of those with a cross-caste link. Income correlation given by Measure I or Measure II as indicated. Outcome variables and regressors are scaled by their standard deviations. Panel A's outcome variable uses only the financial network. Panel B's outcome variable is computed using the financial or social network (two observations per village for the difference-in-difference). All Panel B specifications use village fixed effects, and both panels include controls for caste composition and size (interacted with network type in Panel B). Wild clustered bootstrap standard errors are presented in (.), using Rademacher weights. The cluster is at the subdistrict level in both panels, of which there are 12. 1000 samples are used per bootstrap. p -values from the Wild clustered bootstrap t are presented in [.]

Supplementary Appendix: For Online Publication Only

APPENDIX A. OVER AND UNDER INVESTMENT EXAMPLES

In this Appendix 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 10 implies that Myerson distance of two agents i, j such that $l_{ij} \notin L$ is greater than $1/2$, while the Myerson distance 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 formed. 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 two groups, g, g' both with s members and keep the same parameter values from the previous example. By equation 12 the incentives for form within group links are weakly increased by any across group links. Thus in all stable networks the network structures 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 14 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 the size of both groups s . Considering the Myerson value calculation (equation 5), 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.

APPENDIX B. OMITTED PROOFS

Proof of Proposition 1. To prove the first statement, consider villagers' certainty-equivalent consumption. Let \hat{K} be some constant, and consider the certain transfer K' (made in all states of the world) that i requires to compensate him for keeping a stochastic consumption stream c_i instead of another stochastic consumption stream c'_i :

$$\begin{aligned}
\mathbf{E}[v(c_i + \hat{K})] &= \mathbf{E}[v(c'_i + \hat{K} - K')] \\
-\frac{1}{\lambda}e^{-\lambda\hat{K}}\mathbf{E}[e^{-\lambda c_i}] &= -\frac{1}{\lambda}e^{-\lambda\hat{K}}e^{\lambda K'}\mathbf{E}[e^{-\lambda c'_i}] \\
e^{\lambda K'} &= \frac{\mathbf{E}[e^{-\lambda c_i}]}{\mathbf{E}[e^{-\lambda c'_i}]} \\
K' &= \frac{1}{\lambda} \left(\ln \left(\mathbf{E}[e^{-\lambda c_i}] \right) - \ln \left(\mathbf{E}[e^{-\lambda c'_i}] \right) \right)
\end{aligned}
\tag{20}$$

This shows that the amount K' needed to compensate i for taking the stochastic consumption stream c'_i instead of c_i is independent of \hat{K} . As a villager's certainty-equivalent consumption for a lottery is independent of his consumption level, certainty-equivalent units can be transferred among the villagers without affecting their risk preferences, and expected utility is transferable.

Next, we characterize the set of Pareto efficient risk sharing agreements. Borch (1962) and Wilson (1968) showed that a necessary and sufficient condition for a risk-sharing arrangement between i and j to be Pareto efficient is that in all states of the world $\omega \in \Omega$,

$$\left(\frac{\partial v_i(c_i(\omega))}{\partial c_i(\omega)} \right) \bigg/ \left(\frac{\partial v_j(c_j(\omega))}{\partial c_j(\omega)} \right) = \alpha_{ij}
\tag{21}$$

where α_{ij} is a constant. Substituting in the CARA utility functions, this implies that

$$\begin{aligned}
\frac{e^{-\lambda c_i(\omega)}}{e^{-\lambda c_j(\omega)}} &= \alpha_{ij} \\
c_i(\omega) - c_j(\omega) &= -\frac{\ln(\alpha_{ij})}{\lambda} \\
\mathbf{E}[c_i(\omega)] - \mathbf{E}[c_j(\omega)] &= -\frac{\ln(\alpha_{ij})}{\lambda} \\
c_i(\omega) - c_j(\omega) &= \mathbf{E}[c_i(\omega)] - \mathbf{E}[c_j(\omega)]
\end{aligned}
\tag{22}$$

Letting i and j be neighbors such that $j \in \mathbf{N}(i)$, equation 22 means that when i and j reach any Pareto-efficient risk-sharing arrangement their consumptions will differ by the same constant in all states of the world. Moreover, by induction the same must be true for all pairs of path-connected villagers.

Consider now the problem of splitting the incomes of a set of villagers \mathbf{S} in each state of the world to minimize the sum of their consumption variances:

$$\min_{\mathbf{c}} \sum_{i \in \mathbf{S}} \text{Var}(c_i) \quad \text{subject to} \quad \sum_{i \in \mathbf{S}} y_i(\omega) = \sum_{i \in \mathbf{S}} c_i(\omega) \quad \text{for all } \omega.
\tag{23}$$

Note that,

$$\sum_{i \in \mathbf{S}} \text{Var}(c_i) = \sum_{i \in \mathbf{S}} \sum_{\omega \in \Omega} p(\omega) (c_i(\omega) - \mathbf{E}[c_i])^2,
\tag{24}$$

where $p(\omega)$ is the probability of state ω . As the sum of variances is convex in consumptions and the constraint set is linear, the maximization is a convex program. The first-order conditions of the Lagrangian are that for each $i \in \mathbf{S}$ and each $\omega \in \Omega$, $2(c_i(\omega) - \mathbf{E}[c_i]) = \gamma(\omega)$, where $\gamma(\omega)$ is the Lagrange multiplier for state ω . Thus, $c_i(\omega) - c_j(\omega) = \mathbf{E}[c_i(\omega)] - \mathbf{E}[c_j(\omega)]$ for all $i, j \in \mathbf{S}$. This is exactly the same condition as the necessary and sufficient condition for an ex-ante Pareto efficiency. Hence, a risk-sharing agreement is Pareto efficient if and only if the sum of the consumption variances for all path-connected villagers is minimized.

Using the necessary and sufficient condition for efficient risk sharing, we obtain

$$(25) \quad \sum_{k \in \mathbf{S}} y_k(\omega) = \sum_{k \in \mathbf{S}} c_k(\omega) = |\mathbf{S}| c_i(\omega) - \sum_{k \in \mathbf{S}} (\mathbf{E}[c_i(\omega)] - \mathbf{E}[c_k(\omega)]),$$

which implies that

$$(26) \quad c_i(\omega) = \frac{1}{|\mathbf{S}|} \sum_{k \in \mathbf{S}} y_k(\omega) + \frac{1}{|\mathbf{S}|} \sum_{k \in \mathbf{S}} (\mathbf{E}[c_i(\omega)] - \mathbf{E}[c_k(\omega)]) = \frac{1}{|\mathbf{S}|} \sum_{k \in \mathbf{S}} y_k(\omega) + \tau_i,$$

where $\tau_i = \mathbf{E}[c_i(\omega)] - \mathbf{E}[\sum_{k \in \mathbf{S}} y_k(\omega)]$. □

Proof of Lemma 4. Agent i 's net benefit from forming link l_{ij} is $(MV_i(L) - MV_i(L \setminus \{l_{ij}\}) - \kappa_w)$. We need to show that

$$(27) \quad MV_i(L) - MV_i(L \setminus \{l_{ij}\}) = MV_j(L) - MV_j(L \setminus \{l_{ij}\}) = md(i, j, L)V.$$

Some additional notation will be helpful. Suppose agents arrive in a random order, with a uniform distribution on all possible arrival orders. The random variable $\widehat{\mathbf{S}}_i \subseteq \mathbf{N}$ identifies the set of agents, including i , who arrive weakly before i . For each arrival order, we then have an associate network $\mathcal{L}_L(\widehat{\mathbf{S}}_i)$ that describes the network formed upon i 's arrival (the subnetwork of L induced by agents $\widehat{\mathbf{S}}_i$). Let $q(i, j, L)$ be the probability that i and j are path connected on network $\mathcal{L}_L(\widehat{\mathbf{S}}_i)$.

The certainty-equivalent value of the reduction in variance due to a link l_{ij} in a graph $\mathcal{L}_L(\widehat{\mathbf{S}}_i)$ is V if the link is essential and 0 otherwise. The change in i 's Myerson value, $MV_i(L) - MV_i(L \setminus \{l_{ij}\})$, is then $(q(i, j, L) - q(i, j, L \setminus \{l_{ij}\}))V$. However, $q(i, j, L) = 1/2$. To see this, note that $l_{ij} \in L$ and therefore in every order of arrival in which i arrives after j (which happens with probability $1/2$), i and j are path connected on the network $\mathcal{L}_L(\widehat{\mathbf{S}}_i)$, while i and j are never path connected on $\mathcal{L}_L(\widehat{\mathbf{S}}_i)$ when j arrives after i .

Probability $q(i, j, L \setminus \{l_{ij}\})$ can be computed by the inclusion-exclusion principle, using the fact that the probability of a path connecting i and j existing on network $\mathcal{L}_{L \setminus \{l_{ij}\}}(\widehat{\mathbf{S}}_i)$ is equal to the probability that for some minimal path connecting i and j on $L \setminus \{l_{ij}\}$ all agents on the path are present in $\widehat{\mathbf{S}}_i$. Thus

$$(28) \quad q(i, j, L \setminus l_{ij}) = \sum_{k=1}^{|\mathbf{P}(i, j, L \setminus l_{ij})|} (-1)^{k+1} \left(\sum_{1 \leq i_1 < \dots < i_k \leq |\mathbf{P}(i, j, L)|} \left(\frac{1}{|P_{i_1} \cup \dots \cup P_{i_k}|} \right) \right).$$

We therefore have that

$$(29) \quad MV_i(L) - MV_i(L \setminus l_{ij}) = (1/2 - q(i, j, L \setminus l_{ij}))V = md(i, j, L)V,$$

where the last equality follows from the definition of Myerson distance. \square

Proof of Proposition 5. For there to be underinvestment in a pairwise-stable network L , there must exist a link $l_{ij} \notin L$ for which $TS(L \cup l_{ij}) - TS(L) > 2\kappa_w$. This can only happen if l_{ij} is essential on $L \cup l_{ij}$ as otherwise $TS(L \cup l_{ij}) - TS(L) = 0$. Thus $TS(L \cup l_{ij}) - TS(L) = V$ and so $V > 2\kappa_w$. As $l_{ij} \notin L$ and L is pairwise stable, Lemma 4 implies that $md(i, j, L) \leq \kappa_w/V$. However, as l_{ij} is essential on $L \cup l_{ij}$, $md(i, j, L) = 1/2$. Substituting this into the condition from Lemma 4 we get $V \leq 2\kappa_w$, leading to a contradiction.

For the second part of the proposition, let L be a pairwise-stable network and let $l_{ij} \in L$ be an essential link on L in which there is overinvestment. Thus $TS(L) - TS(L \setminus \{l_{ij}\}) = V < 2\kappa_w$. Since l_{ij} is essential, $md(i, j, L \setminus \{l_{ij}\}) = 1/2$. But Lemma 4 implies that $md(i, j, L \setminus \{l_{ij}\}) \geq \kappa_w/V$. We therefore have that $V \geq 2\kappa_w$, leading to a contradiction. \square

Proof of Proposition 6.

Part (i): By remark 3 and under our regularity condition, all efficient networks are tree networks. By definition, in all tree networks any pair of agents i and j have a unique minimal path between them. Thus, for a tree network L with diameter $d(L)$, there exist agents i and j with a unique minimal path between them of length $d(L)$ and all other pairs of agents have a weakly shorter minimal path between them. Thus by equation 10:

$$(30) \quad md(i, j, L) = \frac{1}{2} - \frac{1}{d(L)} \geq md(k, k', L) \quad \text{for all } k, k' \in \mathbf{N}.$$

By Proposition 5 there is no underinvestment in any stable network. Lemma 4 therefore implies that the efficient network L is stable if and only if $md(k, k', L) \leq \kappa_w/V$ for all k, k' such that $l_{ij} \notin L$. As $md(i, j, L) \geq md(k, k', L)$ and $md(i, j, L) = 1/d(L)$ (see equation 30), this condition simplifies and the efficient network L is stable if and only if

$$(31) \quad \frac{V - 2\kappa_w}{V} \leq \left(\frac{2}{d(L)} \right).$$

As $d(L)$ gets large, the right-hand side converges from above to 0 and so in the limit, the condition for stability becomes $V \leq 2\kappa_w$, which is violated by our regularity condition. Thus, there exists a finite $\bar{d}(L)$ such that the efficient network L is stable if and only if $d(L) \leq \bar{d}(L)$.

Rearranging equation 31, L is stable if and only if

$$(32) \quad d(L) \leq 2 \left(\frac{V}{V - 2\kappa_w} \right).$$

So the key threshold is $\bar{d}(\kappa_w) = \lfloor 2V/(V - 2\kappa_w) \rfloor$.

Fixing the number of agents $|\mathbf{N}|$ in an efficient (tree) network L , the star network is the unique (tree) network (up to a relabeling of players) that minimizes the diameter $d(L)$ while the line network is the unique (tree) network (up to a relabeling of players) that maximizes the diameter $d(L)$. The result now follows immediately.

Part (ii): On any efficient networks all links are essential and generate a net surplus of $V - 2\kappa_w > 0$, where the inequality follows from our regularity condition. As i and j must benefit equally at the margin from the link l_{ij} (see condition (ii) in the definition of agreements that are robust to split the difference renegotiation), agent i 's expected payoff on an efficient network L is

$$(33) \quad u_i(L) = |\mathbf{N}(i; L)|(V/2 - \kappa_w) > 0.$$

Thus i 's net payoff is proportional to his degree.

For any tree network L other than the star network let agent k be one of the agents with the highest degree. Consider a link $l_{ij} \in L$ such that $i, j \neq k$. As L is a tree there is a unique minimal path from i to k and a unique minimal path length from j to k . As we are on a tree network, either the path from j to k passes through i , or else the path from i to k passes through j . Hence either i or j is closer to k and without loss of generality we let i have a longer minimal path to k than j . We now delete the link l_{ij} and replace it with the link l_{ik} . This operation generates a new tree network. Moreover, repeating this operations until there are no links l_{ij} such that $i, j \neq k$, defines an algorithm.

This algorithm terminates at star networks as the operation cannot be applied to this network; There are no links of l_{ij} such that $i, j \neq k$. Moreover the operation can be applied to any other tree network because on all other tree networks there exists an l_{ij} such that $i, j \neq k$. Finally, in each step of the algorithm the degree of k increases and so the algorithm must terminate in a finite number of steps. Moreover, the algorithm must terminate at the star network with k at the center.

By construction, at each step of the above algorithm we decrease the degree of some agent $j \neq k$ and increase the degree of k . Suppose we start with a network L and consider a step of this rewiring where the link l_{ij} is deleted and replaced by the link l_{ik} . Only the expected payoff of agents j and k on L and $L \cup l_{ik} \setminus l_{ij}$ change; The degrees of all other agents remain constant and thus by equation 33 so do their payoffs. Letting $\alpha = (V/2 - \kappa_w)$, we have $u_j(L) = \alpha d_j(L)$, $u_k(L) = \alpha d_k(L)$, $u_j(L \cup l_{ik} \setminus l_{ij}) = \alpha(d_j(L) - 1)$ and $u_k(L \cup l_{ik} \setminus l_{ij}) = \alpha(d_k(L) + 1)$.

It follows that welfare $W(u) = \sum_i f(u_i)$ (see equation 8) decreases through the rewiring in this step if and only if

$$(34) \quad f(\alpha(d_j - 1)) + f(\alpha(d_k + 1)) - f(\alpha d_j) - f(\alpha d_k) < 0,$$

which is equivalent to:

$$(35) \quad f(\alpha(d_k + 1)) - f(\alpha d_k) < f(\alpha d_j) - f(\alpha(d_j - 1))$$

As $f(\cdot)$ is increasing, strictly concave and differentiable $f'(\alpha d_j)\alpha < f(\alpha d_j) - f(\alpha(d_j - 1))$ and $f'(\alpha d_k)\alpha > f(\alpha(d_k + 1)) - f(\alpha d_k)$. Moreover, by concavity $f'(\alpha d_j) \geq f'(\alpha d_k)$ (as $d_k \geq d_j$). Combining these inequalities establishes the claim that $f(\alpha(d_k + 1)) - f(\alpha d_k) < f(\alpha d_j) - f(\alpha(d_j - 1))$.

Thus at each step of the rewiring welfare $W(u)$ decreases. For each network L' reached during the algorithm we can consider the average expected utility $u'(L')$ which if distributed equally would generate the same level of welfare as obtained on L . As aggregate welfare is decreasing at each step of the rewiring $u'(L)$ must be decreasing too. However, the total surplus generated by risk sharing remains constant and so average expected utility \bar{u} remains constant. Recall that Atkinson's inequality measure / index is given by $I(L) = (1 - (u'(L)/\bar{u}))$. Thus at each step of the rewiring the inequality measure $I(L)$ increases. As this rewiring can be used to move from any tree network to the star network, stars network and only star networks maximize inequality among the set of tree networks, which correspond to the set of efficient networks under our regularity condition. As this argument holds for any strictly increasing and differentiable, concave function f it holds for all inequality measures in the Atkinson class.

Consider now an alternative rewiring of a tree network L . Let k be one of the agents with highest degree on L and let j be one of the agents with degree 1 on L . As tree networks contain no cycles, there always exists agents with degree 1 (leaf agents). Pick one of k 's neighbors $i \in \mathbf{N}(k; L)$, remove the link l_{ik} from L and add the link l_{ij} to L . This operation generates a new tree network. Repeating this operation until the highest degree agent has degree 2 defines an algorithm. As the unique tree network with a highest degree of 2 is the line network, the algorithm terminates at line networks and only line networks. At each stage of the rewiring we either reduce the degree of the highest degree agent k or reduce the number of agents who have the highest degree. Thus the algorithm must terminate in a finite number of steps at a line network. Moreover, reversing the argument above, inequality is reduced at each step of the rewiring for any inequality measure in the Atkinson class.

□

Proof of Proposition 7. By definition, underinvestment within group for a network L requires that there exists an $l_{ij} \notin L$ such that $G(i) = G(j)$ and for which $TS(L \cup l_{ij}) - TS(L) > 2\kappa_w$. As

$TS(L \cup l_{ij}) - TS(L) = 0$ for all non-essential links, l_{ij} must be essential on $L \cup \{l_{ij}\}$. Thus l_{ij} is also essential on $\widehat{L} \cup \{l_{ij}\}$ for any $\widehat{L} \subseteq L$. Equation 12 then implies that $TS(\widehat{L} \cup l_{ij}) - TS(\widehat{L}) \geq V$ for any $\widehat{L} \subseteq L$.

Consider any arrival order in which i arrives after j and let \mathbf{S}_i be the agents that arrive (strictly) before i . Agent i 's marginal contribution to total surplus without l_{ij} when i arrives is then $TS(L(\mathbf{S}_i \cup \{i\})) - TS(L(\mathbf{S}_i))$ while with l_{ij} it is $TS(L(\mathbf{S}_i \cup \{i\}) \cup \{l_{ij}\}) - TS(L(\mathbf{S}_i))$. So i 's additional marginal contribution to total surplus when l_{ij} has been formed is $TS(L(\mathbf{S}_i \cup \{i\}) \cup \{l_{ij}\}) - TS(L(\mathbf{S}_i \cup \{i\}))$. As $L(\mathbf{S}_i \cup \{i\}) \subseteq L$, by the above argument $TS(L(\mathbf{S}_i \cup \{i\}) \cup \{l_{ij}\}) - TS(L(\mathbf{S}_i \cup \{i\})) \geq V$. As i arrives after j in half the arrival orders, i 's average additional incremental contribution to total surplus when l_{ij} has been formed is at least $V/2$. Thus $MV_i(L \cup \{l_{ij}\}) - MV_i(L) \geq V/2$. An equivalent argument establishes that $MV_j(L \cup \{l_{ij}\}) - MV_j(L) \geq V/2$. Under our regularity condition $V/2 > \kappa_w$ and so i and j have a profitable deviation to form l_{ij} and the network L is not stable. As L was an arbitrary network within underinvestment within group, there is no stable network with underinvestment within group. \square

Proof of Proposition 8. The proof of the first part of the statement has four steps.

Step 1: Consider any efficient network L that is robust to overinvestment inefficiency within group. This implies that for all path connected agents i, j such that $G(i) = G(j)$ and $l_{ij} \notin L$, either $MV_i(L \cup \{l_{ij}\}) - MV_i(L) \leq \kappa_w$ or $MV_j(L \cup \{l_{ij}\}) - MV_j(L) \leq \kappa_w$. However, by condition (i) in the definition of agreements that are robust to split the difference renegotiation, $MV_i(L \cup \{l_{ij}\}) - MV_i(L) = MV_j(L \cup \{l_{ij}\}) - MV_j(L)$ and so both $MV_i(L \cup \{l_{ij}\}) - MV_i(L) \leq \kappa_w$ and $MV_j(L \cup \{l_{ij}\}) - MV_j(L) \leq \kappa_w$.

Step 2: Let a network $\widehat{L} := \{l_{ij} : G(i) = G(j), l_{ij} \in L\}$ be a network formed from L by deleting all across-group links. Consider any subset of agents $\mathbf{S} \subseteq \mathbf{N}$ such that $i, j \in \mathbf{S}$. As the network L is efficient, it is a tree network that minimizes the number of across-group links conditional on a given set of agents being in a component. This implies that the unique minimal path between i and j cannot contain an across-group link. So, i is path connected to j on the induced subnetwork $L(\mathbf{S})$ if and only if i is path connected to j on the induced subnetwork $\widehat{L}(\mathbf{S})$. Thus, by equation 12, the additional variance reduction that i and j can now achieve by forming a superfluous across-group link on $\widehat{L}(\mathbf{S})$ is weakly lower than on $L(\mathbf{S})$. So, by the Myerson value definition (equation 5), $MV_i(\widehat{L} \cup \{l_{ij}\}) - MV_i(\widehat{L}) \leq MV_i(L \cup \{l_{ij}\}) - MV_i(L)$ and $MV_j(\widehat{L} \cup \{l_{ij}\}) - MV_j(\widehat{L}) \leq MV_j(L \cup \{l_{ij}\}) - MV_j(L)$. This implies that \widehat{L} is robust to overinvestment within group.

Step 3: Let a network \widehat{L}' be a network formed from \widehat{L} by rewiring (alternately deleting then adding a link) each within-group network into a star (for an algorithm that does this, see the part (ii) of the proof of Proposition 6). Consider any two agents i', j' such that $G(i') = G(j')$, $l_{i'j'} \notin \widehat{L}'$. By part (i) of Proposition 6, $MV_{i'}(\widehat{L}' \cup \{l_{i'j'}\}) - MV_{i'}(\widehat{L}') \leq$

$MV_i(\widehat{L} \cup \{l_{ij}\}) - MV_i(\widehat{L})$ and $MV_{j'}(\widehat{L}' \cup \{l_{i'j'}\}) - MV_{j'}(\widehat{L}') \leq MV_j(\widehat{L} \cup \{l_{ij}\}) - MV_j(\widehat{L})$. Thus \widehat{L}' is robust to overinvestment within group.

Step 4: Finally, consider any network $L' \in \mathcal{L}^{CSS}$. This network can be formed by adding a set of across-group links to a network \widehat{L}' such that $\widehat{L}' \subseteq L'$ and if $l_{kk'} \in L' \setminus \widehat{L}'$ then $G(k) \neq G(k')$. Consider any subset of agents $\mathbf{S}' \subseteq \mathbf{N}$ such that $i', j' \in \mathbf{S}'$. Recall that $G(i') = G(j')$ and note that by the construction of L' , $l_{i'j'} \notin L'$. On the induced subnetwork $L'(\mathbf{S}')$, either i' is path connected to j' , in which case $l_{i'j'}$ would be superfluous if added, or else i' and j' are isolated nodes. This is because the within-group network structure for group $G(i')$ is a star. Thus, whenever $l_{i'j'}$ would not be superfluous, the change in i' and j' 's Myerson value if it were added is independent of the across-group links that are present: $MV_{i'}(L' \cup \{l_{i'j'}\}) - MV_{i'}(L') = MV_{i'}(\widehat{L}' \cup \{l_{i'j'}\}) - MV_{i'}(\widehat{L}')$ and $MV_{j'}(L' \cup \{l_{i'j'}\}) - MV_{j'}(L') \leq MV_{j'}(\widehat{L}' \cup \{l_{i'j'}\}) - MV_{j'}(\widehat{L}')$. Thus L' is robust to overinvestment within group.

We turn now to the second part of the result. If $L \notin \mathcal{L}^{CSS}$, then there will be agents i, j such that $G(i) = G(j)$ and $l_{ij} \notin L$ such that either the within-group network structure for $G(i)$ is not a star, or else it is a star but there are across-group links being held by an agent who is not the center agent. In the first case, the inequality in step 3 will be strict by Proposition 6. In the second case, we can without loss of generality let agent i be the non-center agent holding the across-group link. Then, by equation 12, the inequality in step 2 will be strict. Thus for some parameter values L will not be robust to overinvestment within group, but L' will be. \square

Proof of Lemma 9. Denote the set of all possible arrival orders for the set of agents \mathbf{N} , by $\mathcal{A}(\mathbf{N})$. Order this set of $|\mathbf{N}|!$ arrival orders in any way, denoting the k th arrival order by $\widehat{A}_k \in \mathcal{A}(\mathbf{N})$. We will then construct an alternative ordering, in which we denote the k th arrival order by $\widetilde{A}_k \in \mathcal{A}(\mathbf{N})$, such that for arrival order \widetilde{A}_k ,

- (i) i arrives at the same time as agent i' does for the arrival order \widehat{A}_k ;
- (ii) when i arrives he connects to exactly the same set of agents from $\mathbf{N} \setminus \mathbf{S}_0$ that i' connects to upon his arrival for the arrival order \widehat{A}_k ;
- (iii) when i arrives he connects to weakly more agents from \mathbf{S}_0 that i' connects to upon his arrival for the arrival order \widehat{A}_k .

Equation 14 shows that the risk reduction, and hence the marginal contribution made by an agent $k \in \mathbf{S}_0$ from providing the across-group link l_{kj} , is an increasing function of the component size of k 's groups. It then follows that

$$(36) \quad MV(i; L \cup l_{ij}) - MV(i; L) > MV(i'; L \cup l_{i'j}) - MV(i'; L).$$

To construct the alternative ordering of the set $\mathcal{A}(\mathbf{N})$ as claimed we will directly adjust individual arrival orders, but in a way that preserves the set $\mathcal{A}(\mathbf{N})$. First, for each arrival order, we switch the arrival positions of i' and i . This alone is enough to ensure that conditions

(i) and (ii) are satisfied. There are $|\mathbf{S}_0|!$ possible arrival orders for the set of agents \mathbf{S}_0 . Ignoring for now the other agents, we label these arrival orders lexicographically. First we order them, in ascending order, by when i arrives. Next, we order them in ascending order by the number of agents i is connected to upon his arrival. Breaking remaining ties in any way, we have labels $1_i, 2_i, \dots, |\mathbf{S}_0|_i$. We then let every element of $\mathcal{A}(N)$ inherit these labels, so that two arrival orders receive the same label if and only if the agents \mathbf{S}_0 arrive in the same order. We now construct a second set of labels by doing the same exercise for i' , and denote these labels by $1_{i'}, 2_{i'}, \dots, |\mathbf{S}_0|_{i'}$. We are now ready to make our final adjustment to the arrival orders. For each original arrival order \hat{A}_k we find the associated (second) label. Suppose this is $x_{i'}$. We then take the current k th arrival order (given the previous adjustment), and reorder (only) the agents in \mathbf{S}_0 , so that the newly constructed arrival order now has (first) label x_i . Because of the lexicographic construction of the labels, the arrival position of agent i will not change as a result of this reordering of the arrival positions of agents in \mathbf{S}_0 , so conditions (i) and (ii) are still satisfied. In addition, condition (iii) will now be satisfied from the definition of i being more central than i' . The only remaining thing to verify is that the set of arrival orders we are considering has not changed (i.e. that we have, as claimed, constructed an alternative ordering of the set $\mathcal{A}(N)$) and this also holds by construction. \square

Proof of Proposition 10. Let L be an efficient network that is robust to underinvestment across group. This implies that for any across-group link $l_{ij} \in L$ between groups $g = G(i)$ and $\hat{g} = G(j) \neq g$, $MV_i(L) - MV_i(L \setminus \{l_{ij}\}) = MV_j(L) - MV_j(L \setminus \{l_{ij}\}) \geq \kappa_a$, where the inequality follows from condition (i) in the robustness to split the difference renegotiations definition.

We now rewrite L . As the network L is efficient, it is a tree network that minimizes the number of across-group links conditional on a given set of agents being in a component. This implies that the unique minimal path between any two agents from the same group cannot contain an across-group link. We can therefore rewrite the within-group network structures of L to obtain a star by sequentially deleting and then adding within-group links (an algorithm that does this is presented in the proof of part (ii) of Proposition 6). Do this rewiring so that agent i is the agent at the center of the within-group network for group $G(i)$ and let j be the agent at the center of the within-group network for group $G(j)$. Finally, we rewire across-group links so that the same groups remain directly connected, but all across-group links are held by the center agents. Let the network obtained be L' . By construction, $L' \in \mathcal{L}^{CCS}$.

Under our definition of Myerson centrality, it is straightforward to verify that both i and j are weakly more Myerson central within their respective groups on network L' than on network L . An argument almost identical to that in the proof of Lemma 9 then implies that i' and j' have better incentives to keep the link $l_{i'j'}$ than i and j have to keep the link l_{ij} (because the argument is more or less identical we skip it). Hence,

$$(37) \quad MV_{i'}(L) - MV_{i'}(L \setminus \{l_{ij'}\}) \geq MV_i(L) - MV_i(L \setminus \{l_{ij}\})$$

$$(38) \quad MV_{j'}(L) - MV_{j'}(L \setminus \{l_{ij'}\}) \geq MV_j(L) - MV_j(L \setminus \{l_{ij}\})$$

Network L' is therefore robust to underinvestment. Moreover, whenever the within-group networks of i and j on network L are not both stars with i and j at the centers, the inequality is strict because both i and j are strictly more Myerson central within-group on L' than on L . There then exists a range of parameter specifications for which any center-connected star network $L' \in \mathcal{L}^{CCS}$ is robust to underinvestment across group but L is not. \square

Proof of Proposition 11. We simply substitute $s_i = \alpha$ and $\hat{s}_i = \beta$ for $i = 0, \dots, k$ into equation 12. This yields

$$\begin{aligned} \Delta \text{Var}(L \cup l_{ij}, L) &= \left[(1 - \rho_w) + \frac{2(k+1)^3 \beta^2 \alpha^2 - 2(k+1)^3 \alpha^2 \beta^2}{\left(\sum_{i=0}^k s_i\right) \left(\sum_{i=0}^k \hat{s}_i\right) \left(\sum_{i=0}^k s_i + \hat{s}_i\right)} (\rho_w - \rho_a) \right] \sigma^2 \\ &= (1 - \rho_w) \sigma^2, \end{aligned}$$

Multiplying by $\lambda/2$ to get the certainty-equivalent value of the variance reduction completes the proof. \square

Proof of Proposition 13. We will say that agent k is a distance- t neighbor of i if the shortest path from i to k take exactly t steps (and contain $t+1$ agents, including i and k).

Consider the implementation of the Incoming Link Deletion algorithm to find \hat{q}_{ij} . We begin by calculating $x^{1,i} = \mathbf{A}e^i$, where e^i is the i th basis vector. This identifies all agents connected to i . We then set all entries in the i th row from the adjacency matrix \mathbf{A} to 0 and call this new matrix \mathbf{A}_2 . This deletes the inward links to i in the network L . Starting from i 's neighbors, we then find their neighbors on \mathbf{A}_2 . In other words we calculate $x^{2,i} = \mathbf{A}_2 x^{1,i}$. This identifies the distance-2 neighbors of i . We then delete the rows of \mathbf{A}_2 that are indexed by one of i 's neighbors, and so on.

In the t th round the algorithm identifies the distance- t neighbors of i . Thus, for $t < l$, $x_j^{l,i} = 0$; for $t = l$, $x_j^{l,i} = 1$; and for all $t > l$, $x_j^{t,i} = 0$. Deleting incoming links ensures for all $t > l+1$, $x_j^{t,i} = 0$. As L is a tree there, there is no path of length $l+1$ to j and so $x_j^{l+1,i} = 0$. The algorithm therefore finds the unique minimal path from any i to any j and records its length; If the unique path from i to j has length l , $\hat{q}_{ij} = 1/l$. From equation 10 it is also easily checked that $q_{ij} = 1/l$. \square

APPENDIX C. INCOME CORRELATION FROM OCCUPATION CORRELATIONS

Here we present the rationale for our two caste–income correlation measures. Because we lack income data, we must use occupation data in order to construct proxies. In our surveys we have occupation data for all surveyed individuals, coded as small business owner, land-owning farmer, farm laborer, dairy producer/cattle rearer, sericulture owner, sericulture laborer, government official, garment worker, industrial factory worker, industrialist, mason/construction worker, street vendor, artist (e.g., sculptor), and domestic help.

Let $y_{i,g}$ be the income and $o_{i,g} \in O$ the occupation of person i in group $g \in \{A, B\}$. Also, denote the probability that person i is in occupation o and j is in occupation o' by $\Pr(o_{i,A} = o, o_{j,B} = o')$, where $g(i) = A$ and $g(j) = B$. Finally, it will be useful to denote by ϕ_g the proportion of individuals in group g .

In order to operationalize the quantity $\rho_w - \rho_a$ in our empirical exercises, we use the following measures:

- (1) $\widehat{\rho_w - \rho_a}^I = \sum_g \phi_g \Pr(o_{i,g} = o_{j,g}) - \Pr(o_{i,A} = o_{j,B})$
- (2) $\widehat{\rho_w - \rho_a}^{II} = \sum_{k=1}^K \text{corr}(Caste, Occupation_k) \Pr(Occupation_k)$.

The first measure looks at the difference in the (weighted) share of pairs in the same caste who hold the same occupation relative to the share of pairs in different castes who hold the same occupation. We show that this measure is exact when (i) each individual draws an income independently, although the mean can depend on caste and occupation, and (ii) there are occupation level shock, and these shock has the same variance for every occupation.

The second measure presents a score that ranges from 0 to 1. If caste fully explains occupation (where there is therefore scope for maximal within-caste income correlation), the score is 1. However, if caste does not explain occupation at all, the score is 0. This measure is computed as the average caste–occupation correlation, averaged over the occupations.

In what follows, assume the following:

$$(39) \quad y_{i,g} = \mu_{g,o} + \epsilon_o + u_i,$$

where ϵ_o is a mean-0, variance- σ_o^2 iid shock that hits each occupation, u_i is an iid shock with mean zero and variance σ_u^2 that hits each individual, and $\mu_{g,o}$ is a (caste–occupation)-specific mean.

By the law of total covariance, we have

$$(40) \quad \text{cov}(y_{i,A}, y_{j,B}) = \sum_{o \in O} \sum_{o' \in O} \text{cov}(y_{i,A}, y_{j,B} | o_{i,A}, o_{j,B}) \Pr(o_{i,A} = o, o_{j,B} = o') \\ + \text{cov}(E[y_{i,A} | o_{i,A}], E[y_{j,B} | o_{j,B}]),$$

which we will use in our computations below.

C.1. Identical means and variances. We begin with the simple case where all means and variances are identical. This justifies the use of $\widehat{\rho_w - \rho_a}^I$.

Lemma C.1. Assume that $\mu_g = \mu$ for all g and $\sigma_o^2 = \sigma^2$ for all $o \in O$. Then

$$(1) \rho_w \propto \sum_g \phi_g \Pr(o_{i,g} = o_{j,g}) \text{ and}$$

$$(2) \rho_a \propto \Pr(o_{i,g} = o_{j,g'}),$$

both with the same constant of proportionality: $\frac{\sigma^2}{\sigma^2 + \sigma_u^2}$.

Proof. It is immediately clear that

$$\text{cov}(y_{i,A}, y_{j,B} | o_{i,A} = o, o_{j,B} = o') = 0$$

if $o \neq o'$, and if $o = o'$ then

$$\text{cov}(y_{i,A}, y_{j,B} | o_{i,A} = o, o_{j,B} = o) = \mathbb{E}[(\epsilon_o)(\epsilon_o)] = \sigma^2.$$

Also, note that

$$\text{cov}(\mathbb{E}[y_{i,A} | o_{i,A}], \mathbb{E}[y_{j,B} | o_{j,B}]) = 0,$$

implying

$$\text{cov}(y_{i,A}, y_{j,B}) = \Pr(o_{i,A} = o_{j,B}) \sigma^2.$$

Thus the correlation between $y_{i,g}$ and $y_{j,g'}$ is

$$\begin{aligned} \text{corr}(y_{i,g}, y_{j,g'}) &= \frac{\text{cov}(y_{i,g}, y_{j,g'})}{\sqrt{(\sigma^2 + \sigma_u^2)(\sigma^2 + \sigma_u^2)}} \\ &= \Pr(o_{i,g} = o_{j,g'}) \frac{\sigma^2}{\sigma^2 + \sigma_u^2}. \end{aligned}$$

Weighting by population share, we have

$$\rho_w = \phi_A \rho_{w,A} + \phi_B \rho_{w,B},$$

which completes the proof. \square

C.2. Differing average incomes by caste and occupation.

Lemma C.2. Assume that $\mu_{g,o}$ is allowed to vary by caste and occupation. Also, assume that $\sigma_o^2 = \sigma^2$, for all $o \in O$. Then

$$(1) \rho_w \propto \sum_g \phi_g \Pr(o_{i,g} = o_{j,g}) \text{ and}$$

$$(2) \rho_a \propto \Pr(o_{i,g} = o_{j,g'}),$$

both with the same constant of proportionality: $\frac{\sigma^2}{\sigma^2 + \sigma_u^2}$.

Proof. The same argument as in the preceding lemma gives us the result. One can check that the heterogeneity in means does not affect the covariance terms. \square

C.3. Differing variances by occupation.

Suppose now that

$$y_{i,g} = \mu + \epsilon_o + u_i,$$

where ϵ_o is a mean-0, variance- σ_o^2 iid shock that hits each occupation. It follows that

$$\text{cov}(y_{i,A}, y_{j,B} | o_{i,A} = o, o_{j,B} = o') = 0$$

if $o \neq o'$, and if $o = o'$ then

$$\text{cov}(y_{i,A}, y_{j,B} | o_{i,A} = o, o_{j,B} = o') = E[(\epsilon_o)(\epsilon_o)] = \sigma_o^2.$$

We still get $\text{cov}(E[y_{i,A} | o_{i,A}], E[y_{j,B} | o_{j,B}]) = 0$.

Therefore, we have

$$\text{cov}(y_{i,A}, y_{j,B}) = \sum_{o \in O} \Pr(o_{i,A} = o_{j,B} = o) \sigma_o^2.$$

Thus the correlation between $y_{i,A}$ and $y_{j,B}$ is

$$\text{corr}(y_{i,A}, y_{j,B}) = \frac{\sum_{o \in O} \Pr(o_{i,A} = o_{j,B} = o) \sigma_o^2}{(\sum_{o \in O} \Pr(o|A) \sigma_o) (\sum_{o' \in O} \Pr(o'|B) \sigma_{o'})}.$$

For people in different castes, one person from caste A and one from caste B , we then have the following expression for across-caste income correlation:

$$\rho_a = \frac{\sum_{o \in O} \Pr(o|A) \Pr(o|B) \sigma_o^2}{(\sum_{o \in O} \Pr(o|A) \sigma_o) (\sum_{o' \in O} \Pr(o'|B) \sigma_{o'})}$$

For two people in caste A , we have the following expression for within-caste income correlation:

$$\rho_{w,A} = \frac{\sum_{o \in O} \Pr(o|A)^2 \sigma_o^2}{(\sum_{o \in O} \Pr(o|A) \sigma_o)^2}.$$

In sum,

$$\rho_w = \Pr(A) \frac{\sum_{o \in O} \Pr(o|A)^2 \sigma_o^2}{(\sum_{o \in O} \Pr(o|A) \sigma_o)^2} + \Pr(B) \frac{\sum_{o \in O} \Pr(o|B)^2 \sigma_o^2}{(\sum_{o \in O} \Pr(o|B) \sigma_o)^2}.$$

This implies that

$$\begin{aligned} \rho_w - \rho_a &= \Pr(A) \frac{\sum_{o \in O} \Pr(o|A)^2 \sigma_o^2}{(\sum_{o \in O} \Pr(o|A) \sigma_o)^2} + \Pr(B) \frac{\sum_{o \in O} \Pr(o|B)^2 \sigma_o^2}{(\sum_{o \in O} \Pr(o|B) \sigma_o)^2} \\ &\quad - \frac{\sum_{o \in O} \Pr(o|A) \Pr(o|B) \sigma_o^2}{(\sum_{o \in O} \Pr(o|A) \sigma_o) (\sum_{o' \in O} \Pr(o'|B) \sigma_{o'})}. \end{aligned}$$

Note that we do not know the values of σ_o for individual occupations, so we are unable to compute this measure directly from the data as we did in previous subsections. Thus, we take two approaches. First, we proceed as before, using Measure I, which ignores the variance heterogeneity. This is an admittedly imperfect proxy for the desired measure here. Second, we treat caste as a binary random variable and occupation as a multinomial, and therefore take an occupation-share weighted correlation between caste and every occupation. This is also an imperfect measure, but one that captures the intuition that if occupation can be very strongly predicted by caste, then the caste-income correlation should be higher. For this reason we use two proxies for our target quantity in the analysis.