

Altruism and the Topology of Transfer Networks^{*}

Simon Heß[†] and Marcel Fafchamps[‡]

December 2, 2022

Latest version: hessS.org/Gambia_Topology.pdf

Abstract

We examine whether the pattern of transfers among villagers can be explained either by a model of pure altruism or by a model of dynamic incentives, by comparing the predictions of these models to actual transfer patterns in Gambian villages. We find that the topology of transfer networks violates some of the predictions of both models. We then generalize the model of pure altruism to include caps on transfers that represent transaction costs, liquidity constraints, and social norms. Even with a parsimonious parametrization, the generalized model predicts equilibrium networks with topological features that match closely those of intra-village transfer networks in rural Gambia. Using a simulation-based indirect inference approach, we estimate the structural parameters of this model, relying only on income data and the topology of the transfer network for identification. We test for the existence of binding caps on transfers and find that network topologies are consistent with strong caps on transfers. Our structural estimates confirm that kinship is a strong predictor of altruistic preferences. We corroborate our structural estimates using proxies for transaction costs from surveys.

Keywords: private transfers, altruism, social networks, indirect inference

JEL Codes: D85, C15

^{*}We thank Ingela Alger, Renaud Bourlès, Yann Bramoullé, Georg Dürnecker, Daniel Gutknecht, Ludovic Renou, Patrick Schmidt, Matthias Schündeln, participants of the *Annual Meeting of the European Economic Association*, participants of the *12th World Congress of the Econometric Society*, and participants of the Theory Seminar at the Aix Marseille School of Economics for helpful comments and discussions. The data used in this paper were collected by Simon Heß, Dany Jaimovich, and Matthias Schündeln for another research project with support by the DFG (Deutsche Forschungsgemeinschaft) through grant 250842093. Earlier drafts were titled “Altruism, Capacity Constraints and the Topology of Transfer Networks”.

[†]Simon Heß: Goethe University Frankfurt, hess@econ.uni-frankfurt.de

[‡]Marcel Fafchamps: Freeman Spogli Institute, Stanford University: fafchamp@stanford.edu

1 Introduction

Poor households in low-income countries have limited access to formal insurance and credit. They must rely on social networks for informal assistance in the form of intra-household transfers of cash and favors. These transfers help reduce economic inequality (see, e.g., Schechter and Yuskavage, 2011; Asongu, 2015; Beck et al., 2019; Gagnon and Goyal, 2017; Ambrus and Elliott, 2018) and they provide a means to insure against income risk—even if not perfectly (see, e.g., Udry, 1994; Cox and Jimenez, 1998; Fafchamps and Lund, 2003; De Weerd and Dercon, 2006).

Our paper contributes to this line of research by building on two key insights from the theoretical and empirical literature. The first insight is that transfers between households can be formally represented as a graph or network in which households are nodes and transfers are edges (e.g., Genicot and Ray, 2005; Bloch et al., 2008). When transfer flows are examined in this light, the resulting graphs are often found to exhibit complex topological features spanning multiple nodes. Examples include triangles—two transfer partners of a household also having a transfer (e.g., Jackson et al., 2012; Bloch et al., 2008)—and stars—households that receive from or give to many households (e.g., Bala and Goyal, 1998; Jackson and Wolinsky, 1996). Households can also act as intermediaries, receiving from some and giving to others (e.g., Bourlès et al., 2017). These topological features form a central theme of this paper and are used to distinguish between different micro-founded models of transfer networks.

The second insight from the literature relates to the role of incentives in the formation of favor exchange networks. Two classes of theoretical models have been studied, depending on whether transfers are based on incentives or on preferences. The first class allows self-interested individuals to sustain mutually beneficial exchange by relying on the threat of ostracization to deter deviation (e.g., Coate and Ravallion, 1993; Kocherlakota, 1996; Bloch et al., 2008; Karlan et al., 2009; Jackson et al., 2012). In the second class of models, the magnitude of transfers is determined by the strength of altruism, grounded either in sentimental attachment or in personal and social norms (e.g., Bourlès et al., 2017; Bourlès et al., 2020; Alger and Weibull, 2013; Alger and Cox, 2019). Although altruism can be directed towards coreligionists or mankind in general, it tends to be stronger towards relatives and friends (Cox and Fafchamps, 2008; Curry et al., 2013). Our analysis contributes to this second strand of research by structurally estimating a model of network formation with altruistic preferences.

Our analysis builds on the model of network formation of Bourlès et al. (2017). In this model, agents are altruistic and derive utility when others consume. Transfers serve to equalize weighted marginal utilities of consumption between household pairs. The resulting transfer network contains no cycles. We generalize the model to allow caps on feasible transfers, e.g., because of transfer costs or imperfect observation. This addition allows the emergence of undirected cycles in the transfer network—while continuing to rule out directed cycles. In contrast, transfer network models that rely on the threat of ostracization typically imply cycles either between two nodes (i.e., reciprocated transfers) or among multiple nodes (e.g., Bloch et al., 2008; Jackson et al., 2012). These considerations form the starting point of our testing strategy.

Structural models of network formation are notoriously difficult to estimate. As a result, empirical studies of network formation often compare the topological features of empirical networks to model predictions. Existing studies of network formation have followed two main approaches. The first explicitly specifies the topological features of inter-

est and empirically studies their correlates. Exponential random graph models (ERGMs, see Bhamidi et al., 2011; Shalizi and Rinaldo, 2013), statistical ERGMS (Chandrasekhar and Jackson, 2014), and subgraph-generation models (SUGMs, see Chandrasekhar and Jackson, 2018) fall into this category. These reduced-form models offer the advantage that they can be estimated without further modification. But they are limited in the interpretability of their results regarding the underlying network formation process. A second strand of literature starts from models in which topological features emerge endogenously as equilibria from individual or pairwise decision making (e.g., Bloch et al., 2008; Jackson et al., 2012; Vandenbossche and Demuynck, 2013; Bourlès et al., 2017). The advantage of these models is that they offer a clear micro-founded interpretability. However, their structural parameters are difficult to estimate, e.g., due to the multiplicity of equilibria. Furthermore, network data often lack sufficiently detailed information on consumption, transfer amounts, and social structure to pin down model parameters.

The present paper contributes to both strands of the literature. Like the first strand, our estimation relies on the frequency of topographical features. But, like the second strand of literature, our approach to estimation is simulation-based and micro-founded. Our approach to estimation also sidesteps one common problem with transfer data, namely that transfer volumes are difficult to measure and difficult to compare across types of transfers (e.g., gifts and credits; in-kind and cash). This limits the usefulness of transfer volumes as estimation variables. We instead propose an empirical model of transfer flows that can be estimated using only the network topology of transfer flows, combined with data on pre-transfer income.

For estimation purposes, we use detailed transfer data from 56 villages in the Gambia, West Africa (see Heß et al., 2021). What makes these data unusually rich is the fact that we have a complete enumeration of all the households in each village, and of all the pairwise transfers between them. This enables us to sidestep the estimation problems that arise in sampled network data (e.g., Chandrasekhar and Lewis, 2016).

The choice of structural model that we apply to the data is motivated by the characteristics of the transfer flow networks that we observe in our data (see Appendix A for details). First, as in previous studies (e.g., Fafchamps and Gubert, 2007), we find that richer households give to poorer households: an income difference of 1,000 Gambian Dalasis (around US\$25) is associated with a 43 percentage point higher probability of transfer. Second, transfer flows often are indirect: one third of recipients who receive transfers from one household also give to another, thereby acting as intermediaries in transfer flows (see also Beck et al., 2019). Third, households with a richer social network give to households with a poorer social network, controlling for the pair’s income difference: a 1,000 Dalasis difference in the average income of the two household’s kinship networks is associated with a significant 13 percentage point increase in the probability of a transfer between them. Fourth, most networks in our data have a forest-like structure with few cycles; and when there are cycles, these are mostly undirected. This is best illustrated by the low frequency of closed triads in our data: only 13% of all connected triplets form a cycle; and only 2.4% of connected triplets form a directed cycle, that is, have transfers flowing from A to B , from B to C , and from C to A .

The frequency of circular flows is much lower than what would be predicted by repeated game models in which indirect reciprocity is needed to sustain transfers (e.g., Bloch et al., 2008; Jackson et al., 2012), making these models a poor starting point for structural estimation on these data. In contrast, many of the observed features of transfer networks are easy to reconcile with the model of Bourlès et al. (2017) in which households

choose transfers to maximize altruistic utility from their own consumption and that of others. This model, however, is unable to account for the presence of any cycle in the transfer data, whether directed or undirected. To allow for cycles, we generalize the model by adding capacity constraints on transfers and we show that this extension allows for undirected cycles—but still predicts the absence of directed cycles which, as we have just shown, are quite rare in our data.

To estimate the model, we develop an efficient simulation algorithm that finds equilibrium networks for all possible configurations of the model parameters in a short computing time. This enables us to use simulation-based indirect inference (Gourieroux et al., 1993) to estimate the parameters of the structural model. Of particular interest is the inclusion of parameters capturing the existence of caps on transfers and the extent to which kinship predicts altruism. Since our model nests the model by Bourlès et al. (2017), we are able to test their model against ours, and we reject their model for our data. We also find evidence that kinship is an important predictor of altruistic behavior.

This paper makes three contributions. First, we propose a model that generalizes existing work and yields equilibrium networks that are in line with stylized facts about transfer networks. Second, we devise a way to estimate a complex network formation model using only income data and a binary indicator of transfers, without requiring data on consumption or transfer values. Third, our application of this approach to network data from The Gambia contributes to the understanding of how income is distributed via transfers in villages of West Africa.

The remainder of the paper is structured as follows. Section 2 sets up the model and illustrates the properties of equilibrium networks. Section 3 presents and discusses the data, focusing on topological network features that are relevant for the model. Section 4 introduces the estimation strategy and Section 5 presents and discusses the results. In Section 6 we present comparative statics of the fitted model to shed light on the transfers' role for consumption inequality. Finally, Section 7 concludes.

2 Model

The modeling framework used in our empirical investigation is an extension of the model by Bourlès et al. (2017). That model can be summarized as follows. A village, V , consists of a set of households indexed by $i = 1, \dots, n$. Each household earns a stochastic income y_i . Consumption is equal to income plus transfers received minus transfers given. A transfer from household i to household j is denoted as t_{ij} . Transfers have no costs to the sender, other than the foregone own consumption of the transferred income. Consumption c_i is given by:

$$c_i = y_i + \sum_{j \neq i} t_{ji} - \sum_{j \neq i} t_{ij}.$$

Since we only consider transfers within each village, it follows that $\sum_{i \in V} y_i = \sum_{i \in V} c_i$ for each village V .

Households derive concave utility $u(c_i)$ from consuming c_i . Additionally, households are altruistic in the sense that they care about the utility other households derive from consumption. Altruism is allowed to vary across household pairs ij . The matrix $A = \{\alpha_{ij}\}$ with $\alpha_{ij} \in [0, 1)$ for all ij describes the extent to which household i cares about the consumption utility of j . The utility function that household i maximizes is their *social utility* $v_i(c)$, which combines their own consumption utility and the consumption utility

of all the other households that i cares about:

$$v_i(c) = u(c_i) + \sum_{j \neq i} \alpha_{ij} u(c_j).$$

Since the focus of the model is on arbitraging differences in marginal utilities, not on information asymmetries, each household i in village v is assumed to observe the entire vector of consumption quantities for those households j that i cares about, i.e., for which $\alpha_{ij} > 0$. With this assumption, each household i independently chooses transfers t_{ij} to others by maximizing its social utility subject to a budget constraint, the transfer choices of other households, and a non-negativity condition on each transfer. This implies that households cannot extract transfers from others, they can only give. The solution to this utility maximization defines i 's best response vector t_i to other households' transfers, t_{-i} :

$$\begin{aligned} BR_i(t_{-i}|y, \alpha) &= \arg \max_{\{t_{ij}\}_{j \neq i}} u(c_i) + \sum_{j \neq i} \alpha_{ij} u(c_j) \\ \text{s.t.: } & t_{ij} \geq 0 \\ & c_k = y_k + \sum_{j \neq k} t_{jk} - \sum_{j \neq k} t_{kj} \text{ for } k = 1, \dots, n \end{aligned} \tag{1}$$

The maximand is concave in $c = (c_1, \dots, c_n)$, implying that the first order conditions are sufficient to characterize the solution. As transfers linearly reduce own consumption and increase the recipient's consumption, any interior solution to this maximization problem equalizes the marginal utility $u'(c_i)$ that the sending household i derives from consumption, with the weighted marginal utility $\alpha_{ij} u'(c_j)$ that i derives from j 's consumption.

Bourlès et al. (2017) show that a Nash equilibrium of transfers exists, is unique and that a sequence of iteratively updated best responses converges to the Nash equilibrium. Moreover, the undirected network of equilibrium transfers¹ is an acyclic network consisting of a single or multiple disconnected trees, otherwise known as a forest. Another relevant implication is that intermediation—a household receiving funds from one household and passing them to another—can naturally arise even if the original sender is not altruistic towards the final recipient. If household i cares about household j and j cares about k , flows can redistribute income from i to k , even if i and k do not directly care about each other's consumption. We discuss more properties of best responses and the resulting equilibria after introducing caps on transfers as a generalization of the above model.

2.1 Caps on Transfers

Arguably, one of the most central result of the model from Bourlès et al. (2017) regarding network topology is the prediction that the directed transfer network, and generically also the undirected transfer network, has no cycles and thus is a forest. This model is thus easily falsifiable for specific network data by showing the existence of cycles. Rejecting this model, however, does not provide any information on the plausibility of alternative data generating process. To palliate this shortcoming, we extend the model to allow for caps on transfers and we show that this simple extension leads to a less stringent predictions, namely, that the undirected transfer network can contain cycles. In other

¹More precisely, the undirected network $G = [g_{ij}]$ formed by setting $g_{ij} = 1$ if $t_{ij} > 0$ or $t_{ji} > 0$ and setting $g_{ij} = 0$ otherwise.

words, the extended model makes topological predictions that are more general, and includes the original model as a special case.

There are many reasons why caps on transfer may limit flows between pairs of households in a given period, irrespective of how much the sender would like to give. For in-kind transfers, caps may be physical. For instance, transfers of divisible goods such as food may be easy, while the use of assets or housing is more difficult to transfer. A second reason are transactions costs (e.g., transport, coordination, preparation) that put limits on the size of transfers. For financial transfers, such as cash gifts or credit, the sending household may have limited access to cash (e.g., much of agricultural income is in kind) or limited ability to hoard cash for future use (e.g., no bank account). Other caps on transfers may reflect social norms about what transfers are acceptable, e.g., large transfers made outside the family may be resisted by relatives or foster envy among non-recipients. Caps on transfers can be thought of as an approximation of transfer costs, which our model extension captures in a simplified yet straightforward way.

We now extend model characterized by Equation (1) by introducing a fixed cap on transfers which is the same for all household pairs within one village. We illustrate how this affects the best response function. We then show that an equilibrium network exists and that it does not have a strict forest structure. But first we explain the intuition behind the two models.

In the model of Bourlès et al. (2017), equilibrium transfer flows follow the path of least resistance from high-consumption households to low-consumption households until the highest achievable level of marginal utility equalization is reached. Intuitively, this is achieved by forming chains of rich and poor households. In such a chain, transfers can only flow from the highest to the lowest consumption—i.e., like in an irrigation network, with money replacing water. The largest transfers can then be made by channeling funds through a high capacity path, i.e., through household pairs with a high α_{ij} . Without a cap, a single high α_{ij} is capable of channeling any amount of transfers required along the chain from the upstream to the downstream households.

If the chain is long and the income differences large, large amounts of funds may have to be channeled by intermediate households in the middle of the chain. If transfers are capped, these intermediate households along the strongest path may not be able to transfer all the necessary funds, thereby requiring the use of multiple paths. Again the analogy with an irrigation system is telling: if one canal cannot carry all the required water, it may be necessary to use two or more canals to achieve the uncapped equilibrium. If caps on transfers are too limiting, the uncapped equilibrium may not be reachable.

Formally, household i 's utility maximization problem with caps on transfer volumes can be written as:

$$\begin{aligned} BR_i(t_{-i}|y, \alpha, \kappa) = \arg \max_{\{t_{ij}\}_{j \neq i}} \quad & u(c_i) + \sum_{j \neq i} \alpha_{ij} u(c_j) \\ \text{s.t.:} \quad & t_{ij} \in [0, \kappa] \\ & c_k = y_k + \sum_{j \neq k} t_{jk} - \sum_{j \neq k} t_{kj} \text{ for } k = 1, \dots, n, \end{aligned} \tag{2}$$

where κ represents the cap on the transfer volume.

2.2 Equilibrium Analysis

Since the maximands of Equations (1) and (2) are concave in c and t_{ij} , the first order conditions are sufficient to characterize the best responses. A set of transfer flows $\{t_{ij}\}_{i,j \in \{1, \dots, n\}}$ is a Nash equilibrium if the following conditions are met:

$$\forall i, j, \quad t_{ij} \in (0, \kappa) \quad \Rightarrow \quad u'(c_i) = \alpha_{ij} u'(c_j) \quad (3)$$

$$\text{and} \quad t_{ij} = 0 \quad \Rightarrow \quad u'(c_i) \geq \alpha_{ij} u'(c_j) \quad (4)$$

$$\text{and} \quad t_{ij} = \kappa \quad \Rightarrow \quad u'(c_i) \leq \alpha_{ij} u'(c_j). \quad (5)$$

Conditions 3 and 5 also ensure that a transfer occurs only if the recipient consumes less than the sender, meaning that transfers can only flow from the rich to the poor.

The existence of an equilibrium is guaranteed by a variant of the argument in Bourlès et al. (2017). The logic of the proof follows three steps: First, observe that the equilibrium network can have no directed cycles, because in equilibrium agents will only give to agents who consume strictly less, based on Equations (3) and (5). A directed cycle would violate this for at least one transfer. Second, without directed cycles, the set of possible transfers is bounded from above by $\sum_i y_i$, and from below by 0. Moreover, the set of possible transfers is trivially bounded if κ is finite. Third, since the strategy set is closed and bounded and the utility function is concave, the existence of a Nash equilibrium follows from the Kakutani Fixed-Point theorem (Kakutani, 1941; Glicksberg, 1952). Uniqueness of the equilibrium consumption profile is implied by the proof in appendix A in Bourlès et al. (2017): it applies without adaptations to our case since our generalization with $\kappa < \infty$ only further restricts the strategy space. Bourlès et al. (2017, p. 681) show that a sequence of iteratively updated best responses converges to the equilibrium of the model without caps on transfers. The same mechanism works in the case of the model with caps on transfers.


The biggest difference with the case without caps on transfers is that the equilibrium need not produce a forest structure in the undirected transfer network: excess flows may have to be redirected away from the strongest path. It does, however, remain that the *directed* transfer network still cannot have cycles, as argued above. In other words, the presence of cycles in the directed transfer network—e.g., reciprocal gifts—cannot be accounted for by a model in which transfers are the simple product of altruism. To account for such topography, we would have to introduce a mechanism for reciprocal transfers to occur. We revisit this point in the conclusion.

Before turning to the empirical part of the paper, we provide an example to illustrate how caps on transfers affect the structure of equilibrium transfers, and in particular how they lead to undirected cycles.

2.3 Example

Figure 1 offers a graphical presentation of a model simulation with 10 households and consumption utility $u(c_i) = \log(c_i)$. The chosen values for the income vector y and the altruism matrix A are presented in the first and second panel, respectively. In the first panel, each household is shown as a circle, with the radius of the circle representing household income. The value of income is also given next to that circle. The second panel depicts the strength of altruism ties. Denser lines indicate stronger altruism, with the strongest being $\alpha_{56} = 0.62$ and the weakest $\alpha_{19} = 0.02$.

Figure 1: Examples—Income, Altruism and Transfer Networks With and Without Caps on Transfers



Notes: The first two panels show the fixed covariates of the example “village”. Values were chosen to best illustrate the model’s dynamics. The first panel indicates the income of each household. Larger circles correspond to higher incomes. The second panel shows the strengths of altruism ties between households. For illustrative purposes, altruism is symmetric, $\alpha_{ij} = \alpha_{ji}$, which is not generally required for the model. The third panel shows the equilibrium network that would emerge in an uncapped model, i.e., if $\kappa \rightarrow \infty$. The fourth panel shows the resulting equilibrium network that would be observed if $\kappa = 0.5$.

The third panel of Figure 1 presents the equilibrium network without caps on transfers. This equilibrium was computed using iterated best response updating based on Equation (1). In line with predictions: (i) the network of equilibrium transfers is a forest; (ii) transfers strictly flow from high consumption to low consumption nodes; and (iii) indirect transfers (e.g., from 10 to 2) follow the path of strongest altruism.²

The fourth panel of Figure 1 shows the equilibrium network when transfers are capped at $\kappa = 0.5$. Iterative best response updating is used to compute the equilibrium transfer network based on Equation (2). We see that, relative to the uncapped equilibrium, two transfers have been added and one removed. As a result, the network is not a forest anymore although, as predicted, there are no directed cycles in the transfer network—i.e., a sequence of arrows pointing back to its node of origin. But the transfer network has two *undirected* cycles.

The new transfer from 10 to 2 arises because the strongest path from 10 to 2 is capped, which prevents equalizing 10’s marginal utility with 2’s marginal utility weighted by $\alpha_{2,10}$. As a result, 10 has to give to 2 directly in order to maximize 10’s social utility. The new transfer from 9 to 8 arises because the transfer that 10 gives to 8 is reduced, so that 9, who cares about 8, now makes a direct transfer to 8. We also see that the transfer from 1 to 4 that is present in panel 3 has now vanished. In the uncapped case, this flow was made possible by the larger flow from 7 to 1, which is now reduced because of the caps on transfers. Since 7 does not sufficiently care about 4’s consumption and 1, who does, no longer has enough income to justify making a transfer to 4, this transfer disappears.

3 Data on Networks of Informal Transfers in Gambian Villages

To empirically estimate and test our modeling framework, we use detailed data on social and economic networks from 56 Gambian villages. The data were collected as part of a project evaluating the effects of a randomized development project on social and economic networks (Heß et al., 2021). In addition to information on the characteristics of all households in each village, the data contains unusually detailed information on social

²*Strong* refers to the product of all α_{ij} along the path. For details see Bourlès et al. (2017), Theorem 1.

and economic interactions between all households. These interactions were recorded in interviews with all the households in each village, by asking each respondent to list all the households to whom they gave transfers and from whom they received transfers. This is important because, if we only had a sample of households from each village, we would not directly be able to accurately observe the topology of the transfer network (e.g. Chandrasekhar and Lewis, 2016).

To construct our variables of interest, we rely on data regarding household income and kinship ties. Our dependent variable of interest is the inter-household transfers made within each village that can take a variety of forms. We follow Heß et al. (2021), who distinguish consumption and production-related transfers in the data. For our analysis we focus on consumption-related transfers. The dataset also contains dyadic information on production-related transfers, such as transfers of land, labor, and other agricultural inputs, as well as social interactions, such as leisure time spent together. These transactions are more likely to respond to other motives such as commercial exchanges or friendship and hence they fall outside the purview of our model. They are not included in our estimation.


The most basic forms of consumption-related are cash gifts, in-kind gifts and transfers of food. Also informal loans between household are included because they nearly always include a transfer element: they are made at a preferential (e.g., zero) rate of interest, are uncollateralized, and seldom stipulate a strict repayment date. Hence they imply a foregone income as well as a risk for the giver and effect and increase in consumption for the recipient. We combine all these types of transfers into one network. The monetary or utility value of these different types of transfers is difficult to calculate and to compare. To sidestep the issue, our estimation approach relies exclusively on the topology of the transfer flow network. This has the advantage that it does not require knowledge of the value of transferred consumption to the giver and receiver. Of course, estimation efficiency could in theory be increased using information on transfer values, had it been available. But since our data does not contain information about consumption, this point is moot for our purpose.

The full dataset includes 56 villages. All villages in our sample are poor rural communities ranging from 21 to 98 households. The median annual cash income in the sample is approximately US\$625 (25,000 Gambian Dalasis in 2014). Some 76% of the households state farming as their main occupation. 78% of households have kinship ties to other households in the same village and 62% belong to the largest component of their villages kinship network.

The transfer networks exhibit some interesting topological features. Figure 2 shows networks from three different exemplary villages, chosen to illustrate the variety of network topologies. Figure (a) shows a perfect forest, comprised of one big tree and three smaller chains. This is a border case and in most villages some cycles exist. For example in Figure (b), where several edges would need to be removed to obtain a tree structure. Yet in large parts this network remains somewhat tree-like. Networks also vary in density, as seen when comparing Figures (b) and (c). The denser network in Figure (c) exhibits many cycles. Plots of all 56 transfer networks are found in Appendix C.

Table 1 presents descriptive statistics for a number of relevant topographical properties of the 56 transfer networks in our sample. 84% of the households are connected to someone else from the same village via the transfer network, either as a sender or as a recipient (not shown in Table 1). Edge density in the directed transfer networks mostly falls between 1.8% (mostly in larger villages) and 4.6% (mostly in smaller villages). This difference

Figure 2: Three Transfer Networks for Cash Gifts, In-Kind and Food Gifts and Credit



Notes: Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) in three villages from our sample, chosen to illustrate the variation in characteristics describing transfer network typologies. The visual representation of node locations is based on a stress-majorization algorithm to optimally show topological features (Gansner et al., 2004). Numbers next to singleton nodes indicate the number of nodes with a degree of zero.

Table 1: Summary Statistics for Topological Features of Transfer Networks

Statistic	Pctl(25)	Mean	Pctl(75)
Edge Density (%)	1.85	3.48	4.62
Forest Resemblance (% , q^{forest})	50.90	65.52	79.31
Intermediation rate (% , $q^{\text{intermediation}}$)	26.92	37.96	44.44
Dyadic Transfer-Kin Correlation (%)	8.56	12.85	16.58

Notes: The table shows the lower and upper quartiles and the median of the distribution of five indicators for transfer network structure across the 56 village form our sample. These forest resemblance and the intermediation rate are also two of the five auxiliary parameters used for the estimation described in Section 4.2. Edge density measures the fraction of dyads with transfers. Forest resemblance measures whether network components are tree-like. This is quantified by computing the maximal share of edges that can be retained if edges are dropped until all components are trees. The intermediation rate is the share of nodes that have an incoming and an outgoing transfer among households that have any transfer at all (not counting reciprocal links). The dyadic transfer-kinship correlation is the correlation coefficient between an indicator for (undirected) transfer and an indicator for kinship.

mostly reflects variation in village size—the degree distribution remains relatively stable across villages.

Most networks are not perfect forests. The *forest resemblance*-index is the maximal possible proportion of edges that can be retained when dropping edges until all components are trees (the measure is based on the networks' *circuit rank*). For the examples in Figure 2, *forest resemblance* is 100% in panel (a) but only 84% and 44% in panels (b) and (c) respectively. In spite of not having a perfect forest-structure many villages still only show a very limited number of cycles. Half of the village networks could retain two thirds of their links when all edges were dropped until all components are trees. Transfer networks in close to a quarter of our sample villages could be made forests by dropping 6 or fewer edges.³

In line with model predictions (see Section 2), many households act as de-facto intermediaries in the transfer network. On average, 68% of the households that are connected in the transfer network give *and* receive transfers. Transfers are positively correlated with kinship at the dyad level. Further, transfers tend to flow from richer to poorer households. In terms of self-reported annual cash income, the mean difference between sending and

³Based on simple simulations we observe that the networks in the data tend to be less forest-like than Erdős-Rényi-graphs of same size and density.

receiving households is US\$165 (results not shown). This supports the conjecture that one main function of these transfers is to redistribute income.

4 Estimation Strategy

The approach to estimation is parameterized in such a way as to be able to estimate the model parameters of interest using indirect inference (Gourieroux et al., 1993). Indirect inference shares many commonalities with a simulated generalized method of moments (GMM) approach, which can be seen as a special case of indirect inference. The main difference is that elements of the target function, referred to as auxiliary parameters, need not be moments (i.e. means of functions of individual observations in the data), but can be any feature of the observable data. The distinction between the model parameters of interest and the auxiliary parameters is central to indirect inference. Estimates for the model parameters are obtained by finding those values for which auxiliary parameter estimates obtained from simulations most closely resemble the auxiliary parameter estimates computed from the data.

We apply the indirect inference estimation to each village separately, leading to 56 independent sets of estimates. This approach is computationally convenient and has the advantages of allowing for differences across villages, thereby providing more flexibility in the estimation and providing us with a distribution of parameter estimates to examine.

We chose auxiliary parameters such that estimation only requires observing the matrices describing the network of transfers, $T = [\mathbb{1}_{(t_{ij} > 0)}]$, the network of kinship, $K = [\text{kin}_{ij}]$, and the vector of incomes $Y = [y_i]$. The auxiliary parameters are mostly descriptive statistics of the transfer network. For each candidate vector for the structural model parameters, the estimation algorithm simulates draws from the distribution of auxiliary parameters. We then compare these to the auxiliary parameter values for the actual transfer networks, and we iterate until we find the vector of model parameter values that produces the best adequation between the two sets. Thus the estimation algorithm does not try to predict individual links, something we could not do anyway given that we lack data on household consumption and on the value of the different types of transfers.

More precisely, we parameterize the model by positing a relationship between pairwise altruism and kinship:

$$\alpha_{ij} = L(\beta_0 + \beta_1 \text{kin}_{ij} + \varepsilon_{ij}), \quad (6)$$

where $\varepsilon_{ij} \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2)$. $L(\cdot)$ is the logistic function and thus bounded between 0 and 1. Parameter β_1 measures the extent to which altruism is predicted by kinship. The parameter space consists of four parameters: the vector $\beta = \{\beta_0, \beta_1\}$, σ and the village-wide capping parameter κ . The only additional structure we impose is that $u(c) = \log(c)$. The model of Bourlès et al. (2017) is nested as a limit case, $\kappa \rightarrow \infty$, allowing us to test that model's applicability to our data.

Estimation proceeds by evaluating the similarity between the observed networks and the networks implied by any given parameter value $\theta^{(0)} = \{\beta_0, \beta_1, \sigma, \kappa\}^{(0)}$. The four parameters fully specify the conditional distribution of $\alpha_{ij} | \text{kin}_{ij}$. By drawing from this distribution, we generate one possible realization, s , of the altruism matrix $A^{(0,s)} = [\alpha_{ij}^{(0,s)}]$. Given κ , the vector of incomes, and this realization of the altruism matrix, we numerically solve for the equilibrium transfer network through iterative best responses updating. We repeat this process many times, s , for each guess of the parameter vector, n , so as to

calculate a close approximation of the distribution of the vector of auxiliary parameters $q^{(n)}$ corresponding to parameter vector $\theta^{(n)}$.⁴ A optimization algorithm then searches the parameter space of $\theta = \{\beta_0, \beta_1, \sigma, \kappa\}^{\dagger}$ to best match the equivalent sample auxiliary parameters.

Obtaining the equilibrium transfer network via iterative best response updating is computationally burdensome, because it can take many iterations before the process converges to the equilibrium network. To give an overview of the number of steps involved: In each iteration, all households update who they transfer to and how much. Since best response transfers are functions of the vector of predicted consumption, each transfer update potentially affects all other households' best responses.⁵ Thus, solving for equilibrium networks for a given guess of the parameter vector and a given realization of the altruism matrix requires a large number of best response updates. Furthermore, we do this several hundred times for each guess of the parameter vector, each time using a different draw from the conditional distribution of the altruism matrix $[\alpha_{ij}]$. This enables us to calculate the expected value of our objective function for that guess of the parameter vector, by averaging out the random noise stemming from different realizations of the altruism matrix. For the optimization, we then minimize this objective function iteratively through an optimization algorithm that evaluates the objective function several thousand times at different parameter values. Hence estimating the model amounts to computing around 4 million equilibrium networks, i.e., solving for roughly a billion best responses for each of the 56 village network.

4.1 Objective function

We now describe in more detail the GMM-like objective function used for estimation. We consider a vector of parameters from a dyadic regression model and topological features, q , which can be computed for the actual data and for the simulated data. This vector of auxiliary parameters is what is used to identify the parameters of the model. In indirect inference, these features are called auxiliary parameters because they can be understood as parameters of an auxiliary model that is estimable using observable data. The parameters, θ , of the structural model of interest are estimated by choosing them so that the auxiliary parameters of the simulated data and the observed data are as close as possible. The auxiliary model does not need to be correctly specified to ensure consistent estimation, but if it is, indirect inference is equivalent to maximum likelihood (Gourieroux et al., 1993). In our case that implies that the variables used in auxiliary model may be endogenous. The auxiliary model is used as a concise way to summarize the networks. For efficiency, it is desirable to choose auxiliary parameters that are as informative as possible about the estimation parameters of our model.

For a given parameter vector $\theta = (\beta_0, \beta_1, \sigma, \kappa)$, the vector of auxiliary parameters q has a distribution $q \sim Q(\theta)$. Unlike in typical GMM settings, the auxiliary parameters are

⁴The algorithm we use to minimize the target function uses fewer iterations initially, to be fast, but continuously increases the number of iterations to become more precise the closer it gets to finding the minimum of the target function.

⁵In the uncapped model ($\kappa \rightarrow \infty$) transfer quantities can be analytically determined once it is known which transfers are non-zero. For this, we first solve for the equilibrium consumption quantities, using the fact that $t > 0 \Rightarrow u'(c_i) = \alpha_{ij}u'(c_j)$ and that $\sum_i c_i = \sum_i y_i$ is known. Based on the consumption profile and income profile, the transfer quantities can be induced from the structure. For the model with capped transfers this is not possible, because we cannot directly infer from the structure of incomes and altruism whether the $u'(c_i) = \alpha_{ij}u'(c_j)$ holds or whether it would imply a transfer that exceeds the cap.

analytically intractable, so that no moment conditions can be formulated as closed-form expressions that can be solved to derive estimators for the parameters. It is, however, possible to draw from the distribution $Q(\theta)$ and to compute moments of this distribution numerically. By searching the model parameter space for values resulting in similar evaluations of the auxiliary parameters, we obtain estimates of the model parameters.

Specifically, parameters are estimated by minimizing the distance between the auxiliary parameters computed from the real data and the mean auxiliary parameters computed through simulation. This distance is given by:

$$\text{dist}(\theta) = \sum_{s=1}^S (Q^{(s)}(\theta) - q^{\text{data}}) W \sum_{s=1}^S (Q^{(s)}(\theta) - q^{\text{data}}), \quad (7)$$

where q^{data} denotes the value of the auxiliary parameters in the data, S is the number of simulation runs, $Q^{(s)}(\theta)$ indicates the s^{th} simulation-based draw from $Q(\theta)$, and W is a weighting matrix. Hence, $\text{dist}(\theta)$ is evaluated via a large number of simulated draws, S , from the model, given θ . For each draw the statistics in q are computed and the resulting values of $Q^{(s)}(\theta) - q^{\text{data}}$ are averaged.

The indirect inference model parameter estimates are thus defined as:

$$\hat{\theta} = \arg \min_{\theta} \text{dist}(\theta). \quad (8)$$

To find the minimum, we use a combination of two optimization algorithms, namely, a spectral projected gradient (SPG) algorithm (Birgin et al., 2000; Varadhan and Gilbert, 2009) and a particle swarm optimization (PSO) algorithm (El Dor et al., 2012; Zambrano-Bigiarini et al., 2013). This combination has proven useful to achieve a satisfactory optimization performance in the face of two important features of the optimization problem. First, the optimization problem has four dimensions and thus a non-trivial parameter space. Solving for the equilibrium network can take prohibitively long in some parts of the parameter space.⁶ It is thus crucial for performance, that the optimization quickly moves out of areas of the parameter space that correspond to a very low fit to the data. Second, the precision of each evaluation of the target function that is based on a finite number of equilibrium network draws is affected by noise from our simulation-based approximation. In other words, the numeric approximation of the target function is not perfectly smooth. This noise component can be made arbitrary small by increasing the number of equilibrium networks based on which the function is evaluated. This is however computationally costly. The combined algorithm addresses these problems by starting initially with a small number of draws, s , for each function evaluation to quickly move out of unlikely areas of the parameters space and using an increasing number of equilibrium draws the closer it moves to the optimum. The algorithm is described in detail in Appendix E.

Since the data contains transfer networks for many independent villages, we can use the inverse variance-covariance matrix of q^{data} across all villages as weighing matrix. This ensures that highly correlated auxiliary parameters receive less weight and that auxiliary parameters that vary little are matched with higher precision.

⁶This is the case when solving for the equilibrium network requires a large number of best responds updates. This can be the case, e.g., when β_0 and σ have unrealistically extreme values that many α_{ij} are close to zero and one.

4.2 Choice of Auxiliary Parameters

In indirect inference, auxiliary parameters have to be informative about the parameters of the econometric model, defined by Equations (2) and (6), while being calculable using only data observable by the econometrician. In this section, we go through the model parameters, briefly describe how they affect characteristics of the equilibrium transfer network, and discuss auxiliary parameters that capture these characteristics. One benefit of indirect inference is that we can focus on dyadic regression estimates and other topological measures that are easy to compute from the observable data.

Naturally, each auxiliary parameter is likely jointly affected by all model parameters simultaneously. However, some can be linked more directly to particular model parameters than others. Furthermore, some auxiliary parameters may be informative about model parameters only in certain ranges of the parameter space. To account for this, we use multiple auxiliary parameters. Based on simulations, we determined five auxiliary parameters as being most informative. Below we discuss which auxiliary parameters are informative about which of the structural parameters. We start with the parameters that determine the altruism matrix $\alpha_{ij} = L(\beta_0 + \beta_1 \text{kin}_{ij} + \varepsilon_{ij})$, including the variance of ε_{ij} , σ^2 . After this, we discuss the cap parameter κ .

Ideally, the parameters of the altruism matrix would be estimated by regressing altruism, α_{ij} , on kinship and a constant. This is not feasible because α_{ij} is not observable. Instead, we rely on an auxiliary regression, regressing transfers on kinship. Because transfers (unlike altruism) are also determined by the income difference, we additionally control for the dyadic difference in log incomes. The auxiliary dyadic regression model equation is thus:

$$T_{ij} = \gamma_0^{\text{aux}} + \gamma_1^{\text{aux}} \text{kin}_{ij} + \gamma_2^{\text{aux}} \log \text{income diff}_{ij} + \eta_{ij}, \quad (9)$$

where T_{ij} is a binary indicator for whether a transfer from i to j exists.⁷ This model can readily be estimated on the dyadic transfer network from the real data and from networks simulated for specific values for θ . We use the estimates from this auxiliary model together with two additional network topology measures as auxiliary parameters to identify our structural model parameters. Details follow.

Coefficients β_0 and β_1 : Together, the intercept and slope parameters determine the level of altruism and the extent to which altruism depends on kinship. Holding income differences constant, a larger β_0 will lead to higher altruism and thus to more transfers. A larger β_1 will do the same, but only within kin. Equation (9) captures these two effects in γ_0^{aux} and γ_1^{aux} . We thus include these two auxiliary parameters, mainly for the identification of β_0 and β_1 .⁸

Dispersion Parameter σ^2 : Similar to β_0 and β_1 , the variance $\sigma^2 = V(\varepsilon_{ij})$ cannot be directly estimated from observable data, but is related to the variance of the error

⁷For indirect inference estimators to be consistent, the auxiliary model does not need to be correctly specified in the sense of correctly describing the conditional distribution of the dependent variable (Gourieroux et al. (see e.g., 1993)).

⁸Our results are qualitatively robust to replacing these auxiliary parameters with auxiliary parameters that are not obtained from an auxiliary regression. We obtain very similar estimates when using the transfer network density and the dyadic-level correlation between kinship and transfer instead of γ_0^{aux} and γ_1^{aux} .

Table 2: Auxiliary Parameters and Corresponding Model Parameters

Auxiliary Parameter		Model Parameters
Name	Symbol	
Auxiliary Regression Intercept	γ_0^{aux}	β_0
Auxiliary Regression Kin Coefficient	γ_1^{aux}	β_1
Auxiliary Regression Residual Variance	$\sigma^{2,\text{aux}}$	σ^2
Forest Resemblance	q^{forest}	κ
Intermediation Rate	$q^{\text{intermediation}}$	κ

Notes: This table shows an approximate match between auxiliary parameters and model parameters. Details are discussed in Appendix B.

term in the auxiliary regression, $\sigma^{2,\text{aux}} = V(\eta_{ij})$. A larger σ^2 implies that dyads have more extreme values of altruism, i.e. α_{ij} close to zero or one. This in turn implies that kinship and income differences and kinship hold less predictive power for whether or not transactions take place in equilibrium. This is captured in the residual variance of Equation (9), which is why we include the auxiliary model residual variance (an estimate for $\sigma^{2,\text{aux}}$) as an auxiliary parameter mainly to identify σ^2 .

Cap Parameter κ : For the capping parameter, κ , we cannot derive a direct analogue from our dyadic auxiliary regression. There are however two important ways in which this parameter affects observable network structures that we can levy. First, the parameter κ limits the action-space for each individual and is the only parameter of the model that governs whether a forest structure arises in equilibrium. With a sufficiently tight κ networks are no longer guaranteed to be forests. The lower the cap becomes, the more likely cycles emerge. Second, intermediate values of κ that do not break the tree structure may still alter network structures, e.g., by increasing the necessity for indirect transfers to equalize marginal utilities, i.e., transfer chains where i gives to j and j gives to k . This is illustrated in the example in Figure 1. To capture these two topological features as auxiliary parameters, we use (i) the share of links that need to be removed for the network to be a forest, q^{forest} ,⁹ and (ii) the share of connected nodes that have incoming as well as outgoing transfers, $q^{\text{intermediation}}$. The forest resemblance index is important to include because, within the model, it is the strongest indicator for caps on transfers. The intermediation rate is a useful refinement, because caps may be binding, but not strong enough to break the forest structure. The intermediation rate can capture variation in κ where the forest resemblance is still unaffected.

4.3 Summary

Table 2 lists all five auxiliary parameters used to estimate the model parameters and summarizes which model parameters each auxiliary parameter mainly helps identify. Of course, each individual model parameter may affect the realizations of multiple auxiliary parameters. No auxiliary parameter strictly captures a single model parameter. The table serves the purpose of illustrating how each model parameter is related to and approximately identified by at least one auxiliary parameter. Auxiliary parameters may be

⁹This auxiliary parameter can be computed from the circuit rank with a breadth-first algorithm.

correlated, which is why we use the inverse of the empirical variance-covariance matrix of the auxiliary parameters as weighting matrix, W , in the objective function Equation (7).

Indirect inference can be thought of as finding those values of the structural model parameters that yield the networks that maximize the similarity (in terms of the dyadic auxiliary regression results, the forest resemblance and intermediation rate) between simulated data and the real data.

4.4 Inference

To assess estimation uncertainty for our structural estimates, we use a parametric bootstrap approach. We estimate the sampling distribution of parameter estimates by drawing, from the fitted model, $B^V = 10$ alternative realizations of $[a_{ij} | \text{kin}_{ij}]$ per village. Using these draws of altruism and each village’s estimate for κ , we solve for the equilibrium transfer network. Based on these equilibrium networks, we re-estimate all parameters to obtain $B^V \times 56 = 560$ bootstrap draws for all parameter estimates across villages. The variation of these bootstrap estimates around the sample estimates they were obtained from is taken as a measure for estimation uncertainty.


The bootstrap framework has two benefits for our application. First, it provides us with a straightforward way to implement parameter tests, by re-estimating the model with imposed parameter restrictions, thus obtaining a bootstrap sample of (restricted) equilibrium networks. Estimates obtained from these bootstrapped networks provide us with a distribution of parameter estimates under the imposed null hypothesis. Using these distributions, we assess whether our estimates from the real data are consistent with the null hypothesis. Second, the simulation-based estimation inevitably brings about a certain degree of simulation-induced random noise in our estimates. The effect of this on the estimates is small since as our estimation is based on large numbers of independent simulations. It is further reduced because our main focus in the discussion of results are averages of estimates across 56 independent villages, which averages out remaining estimation noise. However the assessment of village-level estimation uncertainty (standard errors) still needs to account for this. Bootstrap-based standard errors, which are computed by repeating the simulation-based estimation procedure, naturally include this source of estimation uncertainty.

5 Results

In this section we present the estimation results from fitting the model to the network data from the 56 Gambian villages. We start by showing how well the fitted structural models replicates the topological features of the transfer networks. A detailed comparison of observed networks with draws from the fitted model is given in Appendix C. Figure 3 shows single random draws from the fitted model for the three examples used in Figure 2. Structural properties are matched remarkably well while—owed to the high degree of interdependence of transfers across dyads—individual links cannot be expected to be predicted with high accuracy. For computational reasons, we estimate monotonous transformations of two parameters instead of the parameters themselves: Instead of σ we estimate $\log(\sigma)$ and instead of κ we estimate $z(\kappa) = \frac{10}{\kappa+1}$.¹⁰

¹⁰The variance parameter σ can, by definition, only take strictly positive values. Our optimization program thus solves for $\log(\sigma)$, which can be any real number. One benefit of this approach is that the

Figure 3: Draws From the Fitted Model for Three Example Villages



Notes: Random draws from the fitted model are shown for the same four villages as in Figure 2. Nodes are not necessarily shown in the same location as in Figure 2. The representation of node locations is based on a stress-majorization algorithm to focuses on topological features (Gansner et al., 2004). Numbers next to singleton nodes indicate the number of nodes with a degree of zero.

Table 3: Average Estimated Structural Parameters

	Mean Estimate (Population Weighted)	Standard Error for Mean Estimates	25 th and 75 th Percentile (Village Estimate)	Standard Error for Village-level Estimate
β_0	-13.4	(0.62)	[-17.9, -7.8]	(6.34)
β_1	6.8	(0.33)	[3.8, 9.1]	(3.38)
$\log(\sigma)$	1.9	(0.06)	[1.6, 2.4]	(0.57)
$z(\kappa)$	6.4	(0.29)	[4.1, 8.8]	(2.83)

Notes: The table shows summary statistics of the coefficient estimates across 56 villages in the sample. The first column shows the mean estimated parameters across 56 independent village estimates. The second columns shows the standard error of that mean, computed through a parametric bootstrap based on repeatedly drawing transfer networks from the fitted model at the estimated parameters, re-estimating parameters based on those draws and computing their means across villages. The reported standard error is obtained as the standard deviation of these bootstrap means. The third column shows the interval spanning the two middle quartiles. The fourth column show estimated standard errors for individual village estimates. These standard errors are also estimated via a parametric bootstrap. The repored numbers are the standard deviations of village-level bootstrap estimates.

In Table 3, we describe the distribution of estimates across the 56 villages. Our discussion of results focuses on the distribution of parameter estimates within our sample, in particular the mean. The dispersion of estimates within our sample, provides a partial indication of the dispersion of the parameter distribution and could be interpreted as the coefficient variance of a random effects model. The variability across villages does however not account for sampling and estimation uncertainty. To gauge estimation uncertainty we use a parametric bootstrap, as described above, to estimate standard errors and to conduct tests for parameters restrictions.

We now focus our discussion on the two main parameters of interest β_1 and κ .

estimation of this parameter follows the logic of an unconstrained optimization problem, like that of our other parameters, β_0 and β_1 . Similarly, by definition κ can only take positive values. Additionally we want to be able to test how well the uncapped model with $\kappa \rightarrow \infty$ fits the data. We thus chose to estimate a monotonous transformation of κ , $z(\kappa) = \frac{10}{\kappa+1}$. This has two immediate implications. First, the nested uncapped model with $\kappa \rightarrow \infty$ is implied by (and testable via) by the parameter constraint $z(\kappa) = 0$. Second, the parameter space for the estimation of $z(\kappa)$ is constrained on both side, $z(\kappa) \in [0, 10]$, and the optimization problem is a standard box-constrained optimization problem in this regard. The scaling factor, 10, as well as the form of this transformation has performance benefits, which we assessed through simulations. Aspects and implications of these transformations are discussed in greater detail in Appendix E.1.

Regarding the relationship between kinship and altruism, β_1 , the results clearly imply that the transfer networks are consistent with a strong positive association between them:¹¹ although estimates for β_1 vary across villages, they are positive in 55 out of 56 villages. To put coefficient size into perspective: using the mean estimates for β and σ , being kin implies an average value of altruism of 0.18, while altruism without kin is 0.03 on average.¹² This indicates that households strongly care about the consumption utility of households with whom they share kinship ties.

Regarding caps on transfers, estimates of $z(\kappa)$ are heterogeneous across villages—a finding that is closely related to the fact that village transfer networks are themselves heterogeneous, in particular with regards to how close they are to being a forest network. As expected, the estimates for $z(\kappa)$ are correlated with the forestness index, implying that the uncapped model with $\kappa = \infty$ fits better to villages with forest-like transfer networks.¹³ The following can help to put the estimated magnitudes into perspective. The village population-weighted mean estimate of $z(\kappa) = 6.44$ implies a cap on transfers of $\kappa = 554$ GMD (US\$14) or approximately 2.2% of the median reported annual per capita income. Put differently, the network typologies are consistent with caps on transfers that amount to the average income of 8 days. A finding that is consistent with the observation that most sample households are subsistence farmers with no significant savings

We conclude that caps on transfers are well below the level of transfers that could equalize marginal altruistic utility from consumption across households. The variability in these estimates across our sample villages deserves closer attention. A small number (7) of the villages have estimates that imply $\kappa > 5,000$ GMD, corresponding to US\$125. It is reasonable to say that, for these villages, caps on transfers can be regarded as effectively non-binding. On the other hand, one half of villages have κ estimates that are at or below US\$9.16, implying significant transfers caps.

Using our parametric bootstrap framework, we conduct a test for the null hypothesis of no caps on transfers. Uncapped transfers imply that $\kappa \rightarrow \infty$ and thus that $z(\kappa) \rightarrow 0$. Our test proceeds as follows. We estimate the restricted model for each village, imposing $\kappa = \infty \Leftrightarrow z(\kappa) = 0$. Using the restricted estimates, we draw from the distribution of altruism and from that obtain the resulting transfer networks under the null. From these bootstrap transfer networks, we estimate the parameters anew to obtain our bootstrap estimates, against which we can compare our original estimates from the sample. If the restricted model is true, the bootstrap transfer networks and thus the estimates we obtain from them, should be similar to those we obtained from our data. If the null is false, they should differ. This is the underlying logic of our test strategy.

In Appendix Figure 21, we plot the distributions of the 56 village estimates against 560 (10 per village) bootstrapped estimates imposing the null hypothesis of no capacity constraint. It is apparent that these distributions are very different. This is the first piece of evidence against the null hypothesis. To construct a p -value for a joint test of the null hypothesis of no caps on transfers across all villages, we use the population-weighted average of the 56 estimates for $z(\kappa)$, as reported in Table 3. We compare this test statistic against the distribution of analogously computed weighted averages

¹¹Since the link function in Equation (6) is strictly increasing, the sign of the marginal effect of kinship on altruism corresponds to the sign of β_1 .

¹²These estimate are based on 1,000,000 draws from Equation (6) with parameters fixed at the mean estimates shown in Table 3.

¹³The median estimate of $z(\kappa)$ among all villages is 7.32. The median estimate for $z(\kappa)$ is 3.4 for villages in the top quartile of the forestness distribution, while it is 6.7 for villages in the bottom quartile.

of 56 bootstrap estimates. One minus the relative rank of the sample estimates in the distribution of restricted bootstrap estimates can be interpreted as a p -value for the null hypothesis. A p -value of 0 in this case indicates that none of the bootstrap replications resulted in an average estimated cap as low as in our data. We find that all bootstrap realizations of the weighted average fall into the range of $[0.6, 2.9]$, they are smaller than the sample estimate. This implies a p -value of 0. We thus reject the null of no binding caps on transfers in this model for our data.

Finally, to corroborate our estimates for the caps, we use separate survey data from the same villages that was not used in the estimation. We construct proxy variables for four factors that we consider to be determinants of liquidity constraints, transaction costs, social norms that relate to caps on transfers. These data come from two survey modules. First, Heß et al. (2021) conducted detailed household interviews with a small random sample of ten households from each village. From these data we construct a variable measuring the fraction of respondents that have access to a bank account. This variable measures liquidity constraints and ranges from 0 to 0.7 and has a mean of 0.15. Second, in the survey in which the network data were collected, a survey module was administered to all villagers with questions on asset ownership. Two assets are relevant indicators for physical transaction costs: carts and motorbikes. We construct a variable indicating the share of villagers owning either a cart or a motorbike. This continuous variable ranges from 0.41 to 1 and has a mean of 0.7. Third, if distance is a determinant of transaction costs, then population density should be correlated with caps on transfers. We therefore construct a measure for population density by dividing the number of households in a village by the village’s area in hectares—where area is approximated by the square of the longest distance between two households in the village. This variable ranges from 0.6 to 14.2 with mean of 4.1. Lastly, we use ethno-linguistic fractionalization (ELF). This variable ranges from 0 to 0.8 with a mean of 0.3 and is assumed to proxy for social norms that can represent caps on transfers (across ethnicities).

We estimate a log-linear specification, regressing the capacity constraint on the these proxy variables. Our sample is geographically diverse, not only covering parts of the country that are close to the urban areas around Serrekunda but also parts of the country that are several hours away from the capital. Since rurality and other spacial properties can be argued to also be correlated with transaction costs, we additionally present results accounting for location in the form of fixed effects for 6 Local Government Areas (LGAs).

Table 4: Relationships Between Estimated Caps and Village-Level Proxies for Caps on Transfers from other Data Sources

	(1) $\log(\kappa + 1)$	(2) $\log(\kappa + 1)$	(3) $\log(\kappa + 1)$	(4) $\log(\kappa + 1)$	(5) $\log(\kappa + 1)$	(6) $\log(\kappa + 1)$	(7) $\log(\kappa + 1)$	(8) $\log(\kappa + 1)$
Villagers with access to bank account	0.48 (0.66)	0.59 (0.65)						
Villagers with carts or motorbikes			0.57 (0.57)	0.65 (0.59)				
Hausholds per hectare					0.02 (0.03)	0.04 (0.03)		
Ethno-linguistic fractionalization							-0.38 (0.42)	-0.42 (0.49)
LGA fixed effects		✓		✓		✓		✓
Observations	56	56	56	56	56	56	56	56

Notes: Robust standard errors in parentheses.

Regression estimates are shown in Table 4. They should not be considered as causal, but as a means to ascertain whether the estimated κ s correlate with variables we would normally associate with constraints on transfers. The power of this analysis is significantly limited by the small sample size of 56 and the use of constructed (and fairly imprecise) proxies and an estimate as the dependent variable. But, with these caveats in mind, the estimates are reassuring. We find that access to banking, availability of means of transport, and population density are correlated with looser caps on transfers, while ethno-linguistic fractionalization is correlated with tighter caps on transfers. Columns 1 and 2 indicate that networks from villages where households have access to bank accounts have on average larger (i.e., less binding) estimated cap on transfers in our model. The point estimate in column 2 suggests that villages with a ten percentage point higher rate of banked households is associated with an approximately 5.9% looser cap on transfers. Columns 3 and 4 indicate that villages with a higher share of households owning carts or motorbikes also have a larger estimated cap. The estimates in column 4 imply that villages with a ten percentage point higher rate of households with means of transportation have on average a 6.5% larger cap on transfers. Columns 5 and 6 show rather small correlations for the relationship between caps and population density. Based on the estimates in column 6, villages with one more household per hectare have on average 0.04% larger caps on transfer. Lastly, columns 7 and 8 show that ethno-linguistically fractionalized villages are estimated to have lower (i.e. tighter) caps on transfers. The point estimate in column 8 implies that villages that are 10 percentage points more fractionalized have on average a 4.2% lower estimate for κ . None of these estimates is significant at a 10%, which is not surprising given the variables used here. The purpose of this part of the analysis is to investigate whether the model-based capacity constraints, estimated from the observed network structure, correlate with more conventional proxies for caps on transfers. Overall these results offer some reassurance that estimates for κ may be interpreted as a cap on transfers.

6 Comparative Statics and Consumption Inequality

While our data only consists of information about income and binary transfer indicators, the fitted model also predicts transfers volumes. Based on predicted transfers we can compute the implied consumption levels and consumption inequality. This allows us to study how consumption inequality is shaped by transfers and caps on transfers.

Table 5: Income and Consumption Inequality

	N	Sample Mean	25 th and 75 th Percentile
Income Gini	56	0.37	[0.33, 0.4]
Consumption Gini $ \widehat{\text{Transfers}}$	56	0.34	[0.29, 0.38]
Consumption Gini $ \widehat{\text{Transfers}} \text{ w } \kappa \rightarrow \infty$	56	0.11	[0.05, 0.16]

Notes: Estimates for consumption inequality measures are averaged over 100 draws of equilibrium transfer networks per village using the parameter estimates described in Section 5.

To this end, we can draw repeatedly from $[\alpha_{ij} | \text{kin}_{ij}, \hat{\beta}, \hat{\sigma}^2]$ for each village, solve for the equilibrium transfer network, obtain the resulting consumption after transfers, and compute the Gini coefficient of consumption. Table 5 shows the result of this process for three different scenarios. The first row shows consumption inequality without transfers,

i.e., income inequality. The second row shows consumption inequality for transfers resulting from the actual parameter estimates. Finally, the third row shows consumption inequality, assuming the actual parameter estimates for altruism (i.e., β_0 , β_1 , and σ) but removing the cap on transfers (i.e., $\kappa \rightarrow \infty$). The Gini coefficients are computed for 100 draws of $[\alpha_{ij}|\text{kin}_{ij}, \hat{\theta}]$ for each village in each scenario and averaged. The table reports the mean and inter-quartile range of this average across our 56 villages.

A first observation is that there is substantial income inequality in the villages, with an average Gini coefficient of 0.37. A second observation is that the transfer volumes predicted by our fitted model reduce this inequality to an average consumption-based Gini coefficient of 0.34. This suggests that informal transfers play an important role in reducing consumption inequality in these villages.

An interesting counterfactual scenario is the consumption inequality that would result from uncapped transfers, i.e., if $\kappa \rightarrow \infty$ in all villages. Consumption inequality with uncapped transfers is predicted to be much lower, with an average Gini coefficient of 0.11 across the 56 villages. This implies that, while the villagers’ altruistic preferences allow them to significantly reduce within-village consumption inequality, the villagers’ ability to achieve this objective is severely limited by transfer caps.

7 Conclusion and Discussion

Transfers of income in the form of gifts and credit play an important role in the lives of people in low-income countries. They serve to redistribute income and are key to understanding the dynamics of inequality and vulnerability to income risk. Theoretical research by Bourlès et al. (2017) examines how altruism shapes transfers and shows that, in the absence of caps on transfers and without alternative motives for transfers, transfer networks have a forest topology. Taking this idea as starting point, we propose a model of capped transfers between households with heterogeneous income and altruistic preferences. The model captures two key properties of observed networks in 56 Gambian villages well, namely: the existence of undirected cycles; and the presence of transfer intermediaries.

We use indirect inference to estimate the structural parameters of the model, including a parameter that captures the correlation between altruism and kinship. Our findings are consistent with the existence of caps on the size of transfer between households. Within the model we study, these caps can explain why transfer networks do not strictly adhere to a forest structure, i.e., why they exhibit undirected cycles. We also find evidence that kinship is a strong correlate of altruistic preferences.

The effectiveness of inequality reduction policies is suspected to be affected by crowding out, that is, the reduction in informal transfer flows that they cause (Cox and Fafchamps, 2008). Our results show that, because of the presence of caps on transfers, transfer networks are limited in their ability to redistribute income. This reduces crowding out and suggests that redistributive policies may be more effective than what would be predicted, e.g., by the model in Bourlès et al. (2017). However, by showing that kinship is an important correlate of altruism, our results highlight one possible concern with income redistribution interventions: In the presence of transfer caps, redistribution across or within kinship groups may affect the ability of existing transfer networks to reduce inequality—a form of crowding out that depends on the topography of the network.

Our results should be seen as one step in empirically studying the determinants of informal inter-household transfers. This paper focuses on modeling transfers that result from preferences rather than from dynamic incentives to sustain cooperation (e.g., Jackson et al., 2012; Bloch et al., 2008). These two mechanisms of network formation are not mutually exclusive and our analysis should be seen as complementary to analyses of mechanisms in which incentives for transfers result from anticipated direct or indirect reciprocal help. For our particular data the focus on altruistic preferences is justified by the low incidence of directed cycles, which we would expect to be more pronounced in networks that are formed through mechanisms based on a dynamic incentive for (indirect) reciprocity. As discussed above, in transfer networks in our data that have closed triads at all, on average one in seven closed triads also have a directed flow. Bilateral reciprocity is equally relatively rare. In the average village, less than one in seven transfers is reciprocated and in three quarters of the villages, less than one in five links are reciprocated. These numbers do not rule out that reciprocal agreements and dynamic incentives play a role. But they suggest that their role may be secondary in our particular context.

Future empirical research should focus on extending our research to other cultural and economic contexts and on generalizing the structural estimation. In particular, estimating models that combine network formation based on altruistic preferences and on dynamic incentives seems relevant. In even richer data than ours, the estimation could be improved upon by accounting for a greater set of dyadic characteristics and by allowing caps to be pair-specific or household-specific. These heterogeneous caps may still be identified through local deviations from tree-like structures (such as triadic closure). In sum, the determinants of transfer networks is a field of research where future developments are likely to be fruitful, as detailed data become increasingly available for structural estimation.

References

- Alger, I. and Cox, D. (2019). “Evolution of the Family: Theory and Implications for Economics”. *Oxford Research Encyclopedia of Economics and Finance*.
- Alger, I. and Weibull, J. W. (2013). “Homo moralis—preference evolution under incomplete information and assortative matching”. *Econometrica* 81 (6), pp. 2269–2302.
- Ambrus, A. and Elliott, M. (2018). *Investments in Social Ties, Risk Sharing and Inequality*. Working Paper 179. Economic Research Initiatives at Duke (ERID).
- Asongu, S. (2015). “The Impact of Mobile Phone Penetration on African Inequality”. *International Journal of Social Economics* 42 (8), pp. 706–716.
- Bala, V. and Goyal, S. (1998). “Learning from neighbours”. *The review of economic studies* 65 (3), pp. 595–621.
- Beck, U., Bjerre, B., and Fafchamps, M. (2019). “The Role of Social Ties in Factor Allocation”. *World Bank Economic Review* 33 (3), pp. 598–621.
- Bhamidi, S., Bresler, G., and Sly, A. M. (2011). “Mixing Time of Exponential Random Graphs”. *Annals of Applied Probability* 21 (6), pp. 2146–2170.
- Birgin, E. G., Martínez, J. M., and Raydan, M. (2000). “Nonmonotone spectral projected gradient methods on convex sets”. *SIAM Journal on Optimization* 10 (4), pp. 1196–1211.
- Bloch, F., Genicot, G., and Ray, D. (2008). “Informal Insurance in Social Networks”. *Journal of Economic Theory* 143 (1), pp. 36–58.
- Bourlès, R., Bramoullé, Y., and Perez-Richet, E. (2017). “Altruism in Networks”. *Econometrica* 85 (2), pp. 675–689.
- (2020). “Altruism and Risk Sharing in Networks”. *Journal of the European Economic Association*.
- Chandrasekhar, A. and Lewis, R. (2016). *Econometrics of Sampled Networks*. Working Paper. Stanford University.
- Chandrasekhar, A. G. and Jackson, M. O. (2014). *Tractable and Consistent Random Graph Models*. Working Paper. National Bureau of Economic Research.
- (2018). *A Network Formation Model Based on Subgraphs*. Working Paper. Stanford University.
- Coate, S. and Ravallion, M. (1993). “Reciprocity without commitment: Characterization and performance of informal insurance arrangements”. *Journal of Development Economics* 40 (1), pp. 1–24.
- Comola, M. and Fafchamps, M. (2014). “Testing Unilateral and Bilateral Link Formation”. *Economic Journal* 124 (579), pp. 954–976.
- Cox, D. and Fafchamps, M. (2008). “Extended Family and Kinship Networks: Economic Insights and Evolutionary Directions”. *Handbook of Development Economics* 4, pp. 3711–3784.
- Cox, D. and Jimenez, E. (1998). “Risk Sharing and Private Transfers: What about Urban Households?” *Economic Development and Cultural Change* 46 (3), pp. 621–637.
- Curry, O., Roberts, S. G., and Dunbar, R. I. (2013). “Altruism in Social Networks: Evidence for a ‘kinship Premium’”. *British Journal of Psychology* 104 (2), pp. 283–295.
- De Weerdt, J. and Dercon, S. (2006). “Risk-Sharing Networks and Insurance against Illness”. *Journal of Development Economics* 81 (2), pp. 337–356.
- El Dor, A., Clerc, M., and Siarry, P. (2012). “A multi-swarm PSO using charged particles in a partitioned search space for continuous optimization”. *Computational Optimization and Applications* 53 (1), pp. 271–295.

- Fafchamps, M. and Gubert, F. (2007). “The Formation of Risk Sharing Networks”. *Journal of Development Economics* 83 (2), pp. 326–350.
- Fafchamps, M. and Lund, S. (2003). “Risk-Sharing Networks in Rural Philippines”. *Journal of Development Economics* 71 (2), pp. 261–287.
- Gagnon, J. and Goyal, S. (2017). “Networks, Markets, and Inequality”. *American Economic Review* 107 (1), pp. 1–30.
- Gansner, E. R., Koren, Y., and North, S. (2004). “Graph drawing by stress majorization”. *International Symposium on Graph Drawing*. Springer, pp. 239–250.
- Genicot, G. and Ray, D. (2005). “Informal insurance, enforcement constraints, and group formation”. *Group Formation in Economics; Networks, Clubs and Coalitions*, pp. 430–46.
- Glicksberg, I. L. (1952). “A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points”. *Proceedings of the American Mathematical Society* 3 (1), pp. 170–174.
- Gourieroux, C., Monfort, A., and Renault, E. (1993). “Indirect Inference”. *Journal of Applied Econometrics* 8 (S1), S85–S118.
- Heß, S., Jaimovich, D., and Schündeln, M. (2021). “Development Projects and Economic Networks: Lessons from Rural Gambia”. *The Review of Economic Studies* 88 (3), pp. 1347–1384.
- Jackson, M. O., Rodriguez-Barraquer, T., and Tan, X. (2012). “Social Capital and Social Quilts: Network Patterns of Favor Exchange”. *American Economic Review* 102 (5), pp. 1857–1897.
- Jackson, M. O. and Wolinsky, A. (1996). “A strategic model of social and economic networks”. *Journal of economic theory* 71 (1), pp. 44–74.
- Kakutani, S. (1941). “A Generalization of Brouwer’s Fixed Point Theorem”. *Duke Mathematical Journal* 8 (3), pp. 457–459.
- Karlan, D., Mobius, M., Rosenblat, T., and Szeidl, A. (2009). “Trust and social collateral”. *The Quarterly Journal of Economics* 124 (3), pp. 1307–1361.
- Kocherlakota, N. R. (1996). “Implications of efficient risk sharing without commitment”. *The Review of Economic Studies* 63 (4), pp. 595–609.
- Schechter, L. and Yuskavage, A. (2011). “Inequality, Reciprocity, and Credit in Social Networks”. *American Journal of Agricultural Economics* 94 (2), pp. 402–410.
- Shalizi, C. R. and Rinaldo, A. (2013). “Consistency under Sampling of Exponential Random Graph Models”. *Annals of Statistics* 41 (2), pp. 508–535.
- Udry, C. (1994). “Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria”. *Review of Economic Studies* 61 (3), pp. 495–526.
- Vandenbossche, J. and Demuynck, T. (2013). “Network Formation with Heterogeneous Agents and Absolute Friction”. *Computational Economics* 42 (1), pp. 23–45.
- Varadhan, R. and Gilbert, P. (2009). “BB: An R package for solving a large system of nonlinear equations and for optimizing a high-dimensional nonlinear objective function”. *Journal of Statistical Software* 32 (1), pp. 1–26.
- Zambrano-Bigiarini, M., Rojas, and R. (2013). “A model-independent Particle Swarm Optimisation software for model calibration”. *Environmental Modelling & Software* 43, pp. 5–25.

A Differences in Two Households' Peer's Incomes Predict Transfers

To show that differences in two households' social networks' incomes hold predictive power for transfers, even after controlling for both households' incomes, we use dyadic linear probability regressions in a sample of 153,862 directed dyads between 2,805 households in 56 Gambian villages. Section 3 discusses descriptives of the sample. For these dyadic regressions, we follow Comola and Fafchamps (2014) and run the regression on a sample in which each directed dyad appears in two measurements—the reported transfer from the household end and the reported transfer from the giving household, to avoid implicitly making assumptions about misreporting in the case of discordant responses.

Table 6: Transfers are More Likely to Occur Between Households with a Larger Difference in Incomes and Between Households with a Larger Difference in the Incomes of Their Social Networks.

	(1) Food & Gifts & Credit $i \rightarrow j$	(2) Food & Gifts & Credit $i \rightarrow j$
difference in own income ($i-j$)	0.43*** (0.09)	0.43*** (0.09)
difference in kin's income	0.13*** (0.04)	
difference in kin and neighbor's incomes		0.06*** (0.02)
Controls	✓	✓
Dyads	307724	307724
Mean transfer probability (%)	1.89	1.89

Notes: This table shows points estimates from a linear regression with a binary dependent variable, taking the value 0 if household i made no transfer to household j and the value 100 if they did. Coefficient estimates can thus be interpreted as percentage points. The independent variables measure the difference between household i 's and household j 's incomes, or the difference in their social ties incomes. In column 1, the income of the social ties is measured as the sum of all incomes of all households with kin relationship, excluding the other household in this dyad. In column 2 this sum also include the income of neighbors. For dyads in which one or both households have no social ties, 0 is imputed for the difference. Control variables are indicators for whether the household head was interviewed or another member, enumerator fixed effects, ethnicity fixed effects. The annual income is measured in USD 1000. Standard errors in parentheses allow for clustering at the village level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The regression results in Table 6 indicate that on average 2% of all dyads have a transfer and that this number is larger in dyads with larger income differences between the sending and the receiving household. However, not only the income of the two households in a dyad matters. Transfers are also more likely to occur if the difference between incomes of other households connected to the sending and to the receiving household respectively are larger. These regressions provide an indication that the structure and income in a wider neighborhood of the village network are drivers of transfers, at least at the extensive margin.

B Example: Effects of Model Parameters

In this section we use an example to illustrate how each parameter affects the equilibrium transfer network. Recall that altruism is parametrized as $\alpha_{ij} = L(\beta_0 + \beta_1 \text{kin}_{ij} + \varepsilon_{ij})$ and household i maximizes utility

$$\begin{aligned} \max_{\{t_{ij}\}_{j \neq i}} \quad & u(c_i) + \sum_{j \neq i} \alpha_{ij} u(c_j) \\ \text{s.t.} \quad & t_{ij} \in [0, \kappa] \\ & c_k = y_k + \sum_{j \neq k} t_{jk} - \sum_{j \neq k} t_{kj} \text{ for } k = 1, \dots, n, \end{aligned}$$

with $u(x) = \log(x)$. For illustrative purposes, here we assume symmetric altruism, $\alpha_{ij} = \alpha_{ji}$, which not generally assumed in the paper.

All examples below use the same kinship network and the same income allocation. The noise in the altruism matrix is simulated using a single draw that is re-scaled in some examples, to reflect varying levels of σ . The left panel always shows the fixed values for income and kinship. The mid panel shows the resulting altruism, which is constructed from kinship and the parameters in β and an error term ε with variance σ^2 . Thicker, darker lines indicate more altruism. The right panel shows the resulting network of equilibrium transfers. The first figures (Figures 4 to 6, 8 and 10) illustrate the role of the parameters of the altruism matrix, while the last two (Figures 11 and 12) focus on the caps on transfers.




Figure 4: Base-example: $\beta_0 = -1.8, \beta_1 = 2, \sigma = 1.3, \kappa = \infty$

It can be seen that some dyads (e.g., 5-9) have strong altruism ties in spite of not being kin, while others ties are weak in spite of kinship (2-6). The transfers result from three factors: differences in income, kinship ties, and other transfers to and from each household. It can be clearly seen how transfers flow from rich to poor, following the paths of strongest altruism. Household 10, the richest, is one of the two household (besides 7) who only gives in equilibrium. Households 2, 3, and 8 only receive and households 5 and 9 act as intermediaries and give and receive at the same time. Four households are disconnected from the main largest component, because they have only weak altruism ties to other households who have enough income to share. One household has neither ingoing no outgoing transfers. Two of the seven transfers occur within the kinship network.

Figure 5 shows the same households with lower β_0 , which significantly reduces altruism in dyads. As a result the equilibrium network only retains two links ($10 \rightarrow 8$ and $7 \rightarrow 1$),

which are in dyads not characterized by strong altruism, but mainly by a large income gap.




Figure 5: Less altruism: $\beta_0 = -2.8, \beta_1 = 2, \sigma = 1.3, \kappa = \infty$

The opposite is illustrated in Figure 6, where β_0 is much higher and altruism is thus stronger everywhere. As a result more transactions appear, so that now all households are part of the transaction network. Three interesting changes occur. First, compared to the base example transfers increase on average in number as well as in volume. Second, this only holds on average. Some transfer volumes are reduced (e.g., $10 \rightarrow 8$), because households who give now care more about other households. Third, intermediaries can give more than they receive (household 9). Household 10 could give to 5 directly and thus reduce the distance the transfer has to travel. The reason this does not happen is that 10 is not directly altruistic towards 5 and transfers (and detours) are not costly in the model.




Figure 6: More altruism: $\beta_0 = 0.2, \beta_1 = 2, \sigma = 1.3, \kappa = \infty$

Figures 7 and 8 illustrate the effect of the parameter β_1 , which determines how kinship affects altruism. In Figure 7 where kinship is a stronger determinant of altruism ($\beta_1 = 3$), three of the four dyads linked by kinship have a transaction. In Figure 8 ($\beta_1 = 1$) not a single transaction occurs within kin.




Figure 7: Altruism depends more on kinship: $\beta_0 = -1.8, \underline{\beta_1 = 3}, \sigma = 1.3, \kappa = \infty$




Figure 8: Altruism depends less on kinship: $\beta_0 = -1.8, \underline{\beta_1 = 1}, \sigma = 1.3, \kappa = \infty$

Figures 9 and 10 show the effect of σ . In Figure 9, noise plays almost no role and kinship is the main determinant of altruism. As a result, transactions are determined by income differences and kinship. Even if altruism was stronger (a larger β_0 , not shown), transactions would still be, largely, determined by income differences. In Figure 10 on the other hand, where the noise is increased ($\sigma = 3$), altruism is less dependent on kinship and in some dyads (e.g., 8-9) extremely strong in spite of no kinship link. As a result, the transfer and the kinship network are almost uncorrelated. Neither are income differences a strong predictor of transfer existence anymore. For very strong altruistic links, transfers occur even with only minor differences in income.




Figure 9: Altruism less noisy: $\beta_0 = -1.8, \beta_1 = 2, \underline{\sigma = 0.3}, \kappa = \infty$




Figure 10: Altruism more noisy: $\beta_0 = -1.8, \beta_1 = 2, \underline{\sigma} = 3, \kappa = \infty$

In all of the above examples, the equilibrium transfer networks were acyclic forests, because with $\kappa = \infty$ the model is identical to the one discussed in Bourlès et al. (2017). The following two figures show parameter combinations that deviate from this. Figure 11 introduces a slight restriction so that no transfers may exceed a volume of $\kappa = 1$. As can be seen from the numbers next to the edges in the right panel, the transfer volumes are reduced compared to Figure 4, even for dyads where the restriction is not binding (see $5 \rightarrow 2$). Aside from the volumes, the transfer network is unchanged (technically it is also possible that links vanish).




Figure 11: Mild caps on transfers: $\beta_0 = -1.8, \beta_1 = 2, \sigma = 1.3, \underline{\kappa} = 1$

Figure 12 shows the transfer network resulting from a stronger cap of $\kappa = 0.5$. Here, new transactions emerge because, after the original transfers (those in Figure 4) are served up to the cap, the remaining income differences between households that would not have income transfers in the base example, are still big enough so that households, e.g., 9 and 10, prefer forming new transactions rather than consuming themselves.

Therefore, the cap on transfers affects the density of the network and also thresholds in disposable income that result in new links. Most clearly, however, caps on transfers can be identified from deviations from the forest structure.




Figure 12: Severe caps on transfers: $\beta_0 = -1.8, \beta_1 = 2, \sigma = 1.3, \kappa = 0.5$

C Topologies of Predicted Networks Versus Actual Networks




Figure 13

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.




Figure 14

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.




Figure 15

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.




Figure 16

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.




Figure 17

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.




Figure 18

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.




Figure 19

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.





Figure 20

Notes: “Real” indicates that the figure shows network as in the data. Transfers of gifts (solid blue), credit (dashed red) and food (dotted green) are shown. “Predicted” indicates that the figure shows a single draw from distribution of equilibrium networks at the estimated parameter values.

D Bootstrapping Results

Figure 21: Distribution of Estimates for $z(\kappa)$ from the 56 Villages Against the Distribution of Parametric Bootstrap Estimates After Imposing the null of $z(\kappa) = 0$



E Numeric Optimization

For optimization we use a combination of two optimization algorithms, namely, a spectral projected gradient (SPG) algorithm Birgin et al. (2000) and Varadhan and Gilbert (2009) and a particle swarm optimization (PSO) algorithm (El Dor et al., 2012; Zambrano-Bigiarini et al., 2013).

The goal of combining the two algorithms is twofold. First, we need to ensure that the parameter space is fully explored. Second, evaluating the target function is extremely slow for large M (which is required for high precision) in areas of the parameter space that are very unlikely. Thus, we do the following:

1. Start 15 PSO searches with a low precision of $M = 4$, initially box-constraining the parameter space of $\{\beta_0, \beta_1, \log(\sigma), z(\kappa)\}$ between $\{6, 15, 4, 10\}$ and $\{-30, -3, -2, 0\}$ and iteratively restricting it further to the lowest and highest values obtained by any particle in the swarm from the previous iteration.
2. At the parameter estimate each final optimum from (1), start an SPG optimization with $M = 200$.
3. Evaluate the 15 parameter estimates from (2) with $M = 4000$ and continue with the best.
4. Start a final SPG optimization routine at that value with $M = 5000$.

This process was gauged in simulations and yields numerically stable results across runs with different starting seeds.

E.1 Parameter Transformations

The optimization algorithm uses monotonous transformations of parameters. The variance parameter, σ , is estimated through estimating $\log(\sigma)$ and the transfer cap, κ , is estimated through estimating $z(\kappa) = \frac{10}{\kappa+1}$. To obtain the actual parameter estimates we thus have to apply the inverse transformations to the estimation output. As a consequence of this, the parameter space is transformed which simplifies and improves the estimation in two ways.

First, the transformations transform the bounds of the parameter space into a more convenient space. While the estimation of σ has to be bounded away from zero, $\log(\sigma)$ behaves like an unconstrained parameter and may take on all real values. Conversely κ would be also bounded away from zero, whereby the special case we want to test corresponds to $\kappa \rightarrow \infty$. Applying the transformation $z(\kappa) = \frac{10}{\kappa+1}$ implies that the special case we want to test corresponds to a standard null hypothesis of $0 = \frac{10}{\kappa+1}$.

A second important benefit of the transformation relates to numeric precision and performance. The algorithm approximates the gradient in steps of homogeneous size. The transformations stretch and compress different areas of the parameter space and can thus increase and decrease precision of this approximation in these areas. For σ , we can clearly see that values $\sigma \ll 10$ should be treated with higher precision, as variation within higher values are uninformative given the logistic transformation (almost all resulting values for α would be very close to either zero or one). The logarithmization of σ thus increases the estimation precision in the areas of the parameters space where variation is meaningful. A similar argument applies to κ where, by design, the variation in κ in exceedingly high ranges does not matter (e.g., a cap that is larger than the villages' total income can never effectively contain the transfer network). Thus, the transformation again ensures that the estimation focuses on variation in relevant ranges.

The gained precision has additional benefits on computational performance. Finding the equilibrium networks for given parameters takes longer in certain parameter ranges, which typically are those ranges where the data fit is very bad. For example, if the altruism matrix has many entries that are numerically very close to 1, the sequential updating of best responses tends to converge to the equilibrium network much slower. Thus, to speed up the optimization algorithm, it is helpful to move on quickly from parameter ranges which provide a bad fit but take longer to evaluate. This also is improved by the stretching and compressing of areas of the parameter space.