This version, July, 2018

Altruism and Insurance In Costly Solidarity Networks[1]

Abstract. We model limited commitment informal insurance networks among altruistic individuals where the impurely altruistic gains to giving to others diminish with the number of transfers one makes; giving is costly, and stochastic income has both publicly observable and unobservable components. Contrary to the canonical informal insurance model, in which bigger networks and observable income are preferable, our model predicts that unobservable income shocks may facilitate altruistic giving that better targets the least well off within one's network and that too large a network can overwhelm even an altruistic agent to cease giving. We test the empirical salience of the model using a unique data set from southern Ghana. We analyze transfer flows among households by coupling observations of gift-giving networks with experimental cash windfall gains - randomized between private and publicly observable payouts - repeated every other month for a year, and find evidence supporting the model predictions.

Chris B. Barrett
Department of Economics and
Charles H. Dyson School of Applied Economics and Management
Cornell University

Vesall Nourani
Department of Economics
Cornell University

Eleonora Patacchini
Department of Economics
Cornell University

Thomas Walker
World Bank

[1]Corresponding Author: vnourani@mit.edu. Seminar participants at Cornell University, Einaudi Institute for Economics and Finance, London School of Economics, and University of Copenhagen provided excellent feedback during early presentations of this paper. Special thanks to residents in the four villages used in this study for opening their homes to share their experiences with me. Data used in this study were collected under USAID's AMA-CRSP Program (P686140). The corresponding author acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1144153 and the African Development Bank Structural Transformation of African Agriculture and Rural Spaces project supported by the Korean government through the Korea-Africa Economic Cooperation (KOAFEC) Trust Fund.

1 Introduction

Social solidarity networks have long been understood to play a central role in village economies. There can be both altruistic and self-interested drivers behind such networks' functioning (Ligon and Schechter (2012)). Although the possibility of altruism has been accommodated in some work within that literature (notably Foster and Rosenzweig (2001)), at least since Popkin (1979) and Posner (1980), the dominant framework for social scientists' understanding of transfers within social networks has rested on self-interested dynamic behavior, commonly framed as self-enforcing informal insurance contracts (Fafchamps, 1992; Coate and Ravallion, 1993; Townsend, 1994). In this tradition, larger networks expand one's social insurance pool, thereby stabilizing consumption, provided that income realizations are publicly observable so as to ensure enforceability of the informal insurance contract (Ambrus et al., 2014). A nice implication of this framework for public policy is that social networks should (at least partially) correct targeting errors in publicly observable transfer programs, as non-recipients who have suffered adverse shocks will approach recipients within their network to share their windfall gains (Angelucci and De Giorgi, 2009).

A related but distinct literature emphasizes the dark side of sharing within social networks, as social pressures can place significant demands on those who enjoy income growth, discouraging investment and potentially even trapping households in poverty (Platteau, 2000; Sen and Hoff, 2006; Jakiela and Ozier, 2016; Squires, 2017). This contrary perspective raises important questions about prospective limits to the value of extensive social networks and of full transparency of individual outcomes.

In this paper we integrate these two streams of thought on the roles social solidarity networks play in village economies. We follow Foster and Rosenzweig (2001) in modeling limited commitment risk pooling allowing for altruistic preferences. Our model includes two key refinements, however, reflecting how our research subjects in rural Ghana describe to us the operation of sharing arrangements within their social networks. First, we model an impure, 'warm glow' component to altruistic preferences (following Andreoni (1990)) that diminishes with the more gifts one gives within one's network. While individuals might vary in the extent of their altruism, everyone faces some limit to the pleasure they derive from beneficence. If giving is costly and the returns to giving diminish, there then emerges some point at which even altruistic individuals cease giving because of the social taxation pressures they face. We term this the 'shutdown hypothesis'. Second, altruistic individuals would like to target their giving toward the neediest members of their social network. But when stochastic income realizations are publicly observable, the demands of less needy members of the solidarity network to share in a windfall can crowd out giving to those with greater need. These two key refinements thereby overturn two key implications of the canonical model of informal insurance - that a larger network and observable income are better - and reflect findings from the literature on social taxation. Nonetheless, within the links that endogenously remain, full risk pooling should obtain, so the social solidarity network retains its key function in distributing stochastic income realization across a network, thereby smoothing consumption.

We take these findings to unique data from southern Ghana, where over the course of a year we randomized private and public bimonthly cash payments to subjects whose gift networks we had previously mapped. These exogenous income shocks enable us to test the predictions of our analytical model. We rely primarily on regressions of giving within subjects' social networks as function of exogenous (randomized) private and public winnings. We corroborate those findings with regressions of how subjects' consumption varies with winnings within one's network and with dyadic regressions reflecting the flows between any two subjects.

Several striking empirical findings emerge, each consistent with the predictions of the analytical model but not with the canonical model of informal insurance. First, the average size of gifts one gives within one's network are larger for private than for public windfall gains. This indicates more targeted giving when altruistic behavior dominates because the unobservability of one's winnings attenuates network demand to honor the informal insurance contract. Furthermore, this confirms the existence of altruistic motives in social solidarity networks. In the absence of altruistic preferences and observability of the income shock, one would never share private winnings. Second, and relatedly, those with unobservable income gains target their giving to the neediest households within their networks. Private, altruistic giving is more sensitive to correcting maldistribution than is public sharing. Third, the number of gifts given is larger for public than for private winnings, consistent with greater network demand for transfers. Fourth, the shutdown hypothesis appears to hold. Winners of publicly revealed cash prizes cease making transfers at all when they have too large a network. Finally, we show that, within these gift networks, we can reject the null hypothesis of full risk pooling, indicating that gift-giving is less likely to follow from insurance motives relative to altruism.

Consistent with the literature on social taxation (Platteau, 2000; Sen and Hoff, 2006; Jakiela and Ozier, 2016), these results highlight the limits to social networks as channels for managing income shocks as well as the trade-offs inherent to transparency in transfer programs. Although observability of income is essential in informal insurance arrangements among purely self-interested agents, observability may impede altruistic agents' ability to focus their giving on the most needy as they are compelled to respond to demands for assistance from the less needy within their network.

Our findings have practical policy implications, especially for cash transfer programs which have, over the past decade or two, become the foundation for social protection programs throughout the developing world. For example, if networks are sufficiently well-connected and populations are motivated by the well-being of others in the network, then transparency may limit the efficiency of redistributive behaviors within networks. Angelucci et al. (n.d.) show that Progresa transfers are pooled by family networks to finance consumption and investment and Advani (2017) shows using experimental data from Pakistan that poverty traps can exist at the network level. Simons (2016) shows that community targeting of a social safety net program is pro-poor relative to centralized targeting. These results suggest that communities in many parts of the world have intimate knowledge of their members' needs and can potentially allocate resources more efficiently than state institutions (Alderman, 2002; Bowles and Gintis, 2002). However, Vanderpuye-Orgle and Barrett (2009) warn that within-community transfers may not benefit "socially invisible" community members — this evidence may, however, be taken with the caveat that exclusion from risk-sharing social networks may also be driven by reputation or punishment for past infringements within the community.

Combined with the above studies, our evidence suggests that governments should tread a careful path when considering the transparency of social safety net transfers. Transparent cash transfers can decrease the opportunity cost of default from potentially efficient risk-sharing networks while also providing a means of triggering social taxation that may deter investment. At the very least, governments should not treat communities as a "black box" and should make efforts to understand and measure the quality of social connections and degree of participation in social networks.

This paper proceeds by first discussing the risk-sharing model that generates the predictions we test in the data in section 2. Then, we discuss the data in section 3. In section 4 we present results from our preferred specification and conduct analysis of dyadic data in section 5. Section 6 concludes by summarizing our findings and providing pathways for future research.

2 The Model

Environment. We introduce 2 agents, $i = \{1, 2\}$ receiving stochastic incomes, $y_i(s_t) \geq 0$ that depend on the state, s_t, realized in period t — a sequence of the state history is characterized by $h_t = \{s_1, s_2, ..., s_t\}$.[2] We model history-dependent transfers from household 1 to household 2, $\tau(h_t)$, when both households have gift-networks with $g_1 = g_2 \geq 1$ other households. Depending on the realization of a particular state, households will receive $g_i p_i(s_t)$ different gift-requests from their network, where $0 \leq p_i(s_t) \leq 1$ reflects the unconditional probability that a given household in one's network will request a transfer in period t. To focus attention on transfers between households 1 and 2, we assume that net transfers with all other households in one's network equals zero for now. Thus, net income for household 1 is $y_1(s_t) - \tau(h_t)$ and net income for household 2 is $y_2(s_t) + \tau(h_t)$. If $\tau(h_t) > 0$, then household 1 (2) is a net sender (receiver) of transfers. Otherwise, if $\tau(h_t) < 0$ household 1 (2) is a net receiver (sender) of transfers.

Preferences. Following Foster and Rosenzweig (2001), we assume households hold altruistic preferences towards each others' single-period utilities. We introduce individual i's altruistic preferences by assuming that household single-period utility is separable in own and other household consumption. Single-period utility for household 1 is reflected in the following equation:

$$\begin{aligned} u_1(c^1) + \gamma_1(g_1, s_t) u_2(c^2) \\ \text{such that } 0 \leq \gamma_1(g_1, s_t) \leq 0.5 \end{aligned} \tag{1}$$

and single-period utility for household 2 can be written in symmetric fashion. $u_1()$ and $u_2()$ are increasing and concave $\gamma_1(g_1, s_t)$ represents the altruism weight household 1 holds towards 2.

We diverge from others in that we characterize altruistic preferences as a function of a household's "altruism stock" and their transfer-network size. The altruism weight diminishes as a household's period-specific gift-requests increase, which in turn rely on a household's gift-giving network size, g_i, and the probability that it will be requested to provide transfers to other households, reflected in $p_i(s_t)$. Specifically, altruism weights consist of a fixed, or "pure," component, $\bar{\gamma}_1^F \geq 0$, and a warm-glow (Andreoni, 1990), or "impure," component $\bar{\gamma}_1^W \geq 0$. Again for household 1, we represent these components of altruism in the following

[2] The assumption of stochastic exogenous income is reasonable in our empirical context since we distribute cash prizes randomly across the sample.

manner:

$$\gamma_1(g_1, s_t) = \min\{\bar{\gamma}_1^F + \frac{\bar{\gamma}_1^W}{g_1 \cdot p_1(s_t)} \mathbb{1}(\tau(h_t) \neq 0), \overline{\gamma}_1\} \tag{2}$$

where $\mathbb{1}(\cdot)$ is an indicator function equal to one when there is a transfer between households 1 and 2, and $\overline{\gamma}_1$ places an upper bound on household 1's altruism weight towards household 2 so that altruism does not rise to arbitrarily large levels when $p_1(s_t)$ is small.

Explicitly stated we assume here that the amount of warm-glow altruism household 1 holds towards household 2 is a decreasing function of the total number of household 1's period t gift-obligations, $g_1 \cdot p_1(s_t)$. This reflects the idea that warm-glow is diminishing in the number of discrete transfers each household participates in — intuitively, the novelty of warm-glow wears off as transfers become more common-place. Without loss of generality, we will set $\bar{\gamma}_1^F = 0$ and focus our analysis around warm-glow altruism — thus, when we speak of altruism moving forward, we are no longer referring to "pure" altruism. Intuitively, and taken together, each household is altruistic towards others but is not unlimitedly so. Households can vary in the "stock" of altruism (or altruistic capital as in Ashraf and Bandiera (2017)) they possess, but will be limited in the degree of altruism they exercise towards other households.

Dynamic Payoffs and Transfer Choices. At period t, households seek to maximize their expected lifetime utility, which requires agreeing upon a history-contingent transfer contract that is preferable to zero transfers across all states. Thus, we assume that households compare payoffs from the dynamic contract to payoffs from a no-transfer rule.[3] To set up the household's problem, we define $U_1(h_t)$ as 1's expected discounted utility gain from the risk-sharing contract with 2 relative to a no-transfer rule after history h_t:

$$\begin{aligned} U_1(h_t) = \quad & u_1(y_1(s_t) - \tau(h_t)) - u_1(y_1(s_t)) \\ & + \gamma_1(g_1, s_t) u_2(y_2(s_t) + \tau(h_t)) - \gamma_1(g_1, s_t) u_2(y_2(s_t)) \\ & + \mathbb{E} \sum_{k=t+1}^{\infty} \delta^{k-t} \left\{ \begin{array}{l} u_1(y_1(s_k) - \tau(h_k)) - u_1(y_1(s_k)) \\ +\gamma_1(g_1, h_t) u_2(y_2(s_k) - \tau(h_k)) - \gamma_1(g_1, h_t) u_2(y_2(s_k)) \end{array} \right\} \\ & - \alpha_1(g_1) \end{aligned} \tag{3}$$

where δ represents the dynamic discounting factor. $\alpha_1(g_1)$ represents a second way in which

[3] Households in Foster and Rosenzweig (2001) revert to a sequence of history-dependent Nash equilibria (SHDNE) in which transfers are maintained even when a household defaults from the contract. Such an environment is not crucial for the type of analysis we conduct in our study. Nevertheless, appendix section A shows how one can adapt our own model to reflect such SHDNE default transfers.

our model diverges from others' — it is the incremental cost to household 1 of maintaining a gift-giving link with household 2 given network size g_1. We assume that the cost of maintaining such a link is convex in network size and can be thought of as the effort required to maintain a social bond and, for example, awareness of household 2's realized income. The contract is enforced if the expected discounted utility surplus is nonnegative. The contract requires an implementability constraint that states that gains from the contract be at least as high as the no-transfer rule: $U_1(h_t) \geq 0$ and $U_2(h_t) \geq 0$. Together, the economic environment, payoffs and transfer decision represent a simultaneous game in which agents seek to find a contract that can be implemented in the presence of limited commitment and no external enforcement mechanism.

Limited Commitment Contract Solution. Following Foster and Rosenzweig (2001) and Ligon et al. (2002), the solution to the utility maximization problem will be a dynamic program in which the current state is given by s out of the set of all states ($s \in \{1, 2, ..., S\}$), and targeted discounted utility gain for household 2, U_2^s, is given.[4] Choice variables in the programming problem will be consumption assignments c_1, c_2 and the continuation utilities U_1^r and U_2^r for each possible state r, resembling the next period. This enables us to write the value function for household 1 as dependent on current target utilities and collective resources: $U_2^s, \{y_1(s) + y_2(s)\}$. Formally, we write the dynamic programming problem as

$$\begin{aligned} U_1^s(U_2^s) = \quad \max_{\tau_s,(U_1^r,U_2^r)_{r=1}^S} \quad & u_1(y_1(s) - \tau_s) - u_1(y_1(s)) \\ & + \gamma_1(g_1(s))u_2(y_2(s) + \tau_s) - \gamma_1(g_1(s))u_2(y_2(s)) \\ & - \alpha_1(g_1) + \delta \sum \pi_{sr} U_1^r(U_2^r) \end{aligned} \tag{4}$$

[4] U_2^s is defined by equation 19 when all subscripts with 1 are replaced with a 2 and vice versa.

subject to

$$
\begin{aligned}
\lambda: \quad & u_2(y_2(s) + \tau_s) - u_2(y_2(s)) \\
& + \gamma_2(g_2(s))u_1(y_1(s) - \tau_s) - \gamma_2(g_2(s))u_1(y_1(s)) \\
& - \alpha_2(g_2) + \delta \sum_{r=1}^{S} \pi_{sr} U_2^r \geq U_2^s
\end{aligned} \tag{5}
$$

$$\delta\pi_{sr}\mu_r: \quad U_1^r(U_2^r) \geq \underline{U}_1^r = 0 \quad \forall r \in S \tag{6}$$

$$\delta\pi_r\phi_r: \quad U_2^r \geq \underline{U}_2^r = 0 \qquad \forall r \in S \tag{7}$$

$$\psi_1: \quad y_1(s) - \tau_s \geq 0 \tag{8}$$

$$\psi_2: \quad y_2(s) + \tau_s \geq 0, \tag{9}$$

where π_{sr} represents the probability of state r occuring. Equation 5 says that transfer and future utility allocations will satisfy the promise-keeping constraint. Equations 6 and 7 state that allocated utility in any state r will be at least as high as the lower bound utility household 1 and, respectively, 2 can receive via defaulting to the no-transfer arrangement. Equations 8 and 9 place non-negativity constraints on consumption allocations in period s. The actual contract can be computed recursively, starting with an initial value for U_2^s.

The concavity of the dynamic programming problem renders the first-order conditions both necessary and sufficient to obtain a solution. Thus, the evolution of the ratio of marginal utility (re-inserting t subscript), together with the envelope condition, characterizes the optimal contract:

$$\frac{u_1'(y_1(s_t) - \tau(h_t)) + \gamma_1(g_1(h_t))u_2'(y_2(s_t) + \tau(h_t))}{u_2'(y_2(s_t) + \tau(h_t)) + \gamma_2(g_2(h_t))u_1'(y_1(s_t) - \tau(h_t))} = \lambda + \frac{\psi_2 - \psi_1}{u_2'(y_2(s_t) - \tau(h_t)} \tag{10}$$

$$-U_1^{r\prime}(U_2^r) = \frac{\lambda + \phi_r}{1 + \mu_r}, \quad \forall r \in S \tag{11}$$

$$\lambda = -U_1^{s\prime}(U_2^s). \tag{12}$$

Taken together, these three conditions imply that a constrained-efficient contract can be characterized in terms of the evolution over time of λ, where $-\lambda$ is the slope of the Pareto frontier.[5] For each state s, there is a history independent interval $[\underline{\lambda}_s, \overline{\lambda}_s]$ that constitute the set of implementable contracts in state s. The lower bound value is the point at which household 1 is indifferent between participating in a risk-sharing contract and default — the upper bound reflects the symmetric position for household 2. The exact value of $\lambda(h_{t+1})$ is history dependent and evolves according to the value of $\lambda(h_t)$ in the following manner

$$\lambda(h_{t+1}) = \begin{cases} \underline{\lambda}_s \text{ if } \lambda(h_t) < \underline{\lambda}_s \\ \lambda(h_t) \text{ if } \underline{\lambda}_s \leq \lambda(h_t) \leq \overline{\lambda}_s \\ \overline{\lambda}_s \text{ if } \lambda(h_t) > \overline{\lambda}_s. \end{cases} \tag{13}$$

Given this contract structure and assumptions on utility parameters and income values, numerical solutions for all interval endpoints can be obtained by solving an $S \times 2$ dimensional non-linear system of equations.

Figure 1 describes the intuition behind this contract structure using a stylized example. Suppose that in an initial period, t, a state is realized in which household 1 receives income $y_1(s_t) = 2$ and household 2 receives $y_2(s_t) = 1$.[6] If the two households follow the contract structure in equation 13, then each household will weigh participation in risk-sharing against the payoff received when they default from such a contract. Household 2 will only consider this contract if $\lambda(h_t)$ is greater than $\underline{\lambda}_{zv}$ — the point at which household 2 is indifferent between defaulting and participating in the risk-sharing contract (discounted utility surplus equal to zero). Household 1 will have a similar payoff structure when $\lambda(h_t) = \overline{\lambda}_{zv}$. Both households will prefer risk-sharing if they can settle on a dynamic contract between these two numbers. Suppose the realized state in period $t+1$ is zz, where $y_1(zz) = y_2(zz) = 1$. If altruistic preferences (and discount rates) are such that the contract intervals for the realized state in $t+1$ does not overlap with the state in t (left panel in figure 1), the surplus will be divided according to $\lambda(h_{t+1}) = \underline{\lambda}_{zz}$. If the contract intervals do overlap, then , $\lambda(h_{t+1}) = \overline{\lambda}_{zv}$. Notice that this results in a division of the surplus in which both households strictly benefit relative to default (i.e., closer to full consumption smoothing).

Income shocks. We now add more structure to the model to study the importance of the transparency of cash transfers. Let us define two types of exogenous income shocks: 1) privately revealed cash prizes (denoted by v) and 2) publicly revealed cash prizes (b). Households that do not receive cash prizes experience zero exogenous income shocks (z). Thus, there are potentially nine different states that can be realized, though we limit our

[5]For a formal proof, see Ligon et al. (2002) and Thomas and Worrall (1988). The extension to the case with altruistic preferences is straightforward as noted by Foster and Rosenzweig (2001).

[6]In later simulations, this income combination will be referred to as state zv

analysis to states in which only up to one household receives a prize of any type: neither 1 nor 2 receive a prize (zz), 2 receives a private prize (vz), 2 receives a public prize (bz), 1 receives a private prize (zv), and 1 receives a public prize (zb).[7] Explicitly, here we are assuming that the prize-winning household receives a higher income than the non-prize winning household and the prizes are equal in value:

Assumption 1 (Prize-winners Have Higher Incomes)

$$y_1(zv) = y_1(zb) = y_2(vz) = y_2(bz) > y_1(zz) = y_1(vz) = y_1(bz) = y_2(zz) = y_2(zv) = y_2(zb)$$

Let us assume that the probability of receiving a transfer request, $p_i(s_t)$, is highest when a household wins a publicly revealed prize. In other words,

Assumption 2 (Observability of Income)

$$p_1(zb) > p_1(s') \textit{ for all } s' \neq \{zb\} \textit{ and } p_2(bz) > p_2(s'') \textit{ for all } s'' \neq \{bz\}.$$

We argue that households who receive easily observable positive income shocks are more likely to be approached by others to uphold their end of an informal gift-giving obligation. This assumption is supported by evidence in similar contexts (e.g., Jakiela and Ozier (2016) and Squires (2017)) in which participants in behavioral experiments willingly spend part of their payoff to allow winfall income gains to be hidden from their peer group. This assumption implies that the warm-glow altruism weight household 1 holds towards household 2, for example, decreases when household 1 wins a publicly revealed lottery.

2.1 Model Simulations

Given the complexity of the state-space, it is not possible to analytical explore solutions to this model. We are, however, fundamentally interested in how the risk-contract depends on the size of the gift giving network g_1 and the public or private nature of the prize in the realized state — thus, we explore numeric solutions using set values for model parameters while allowing network size to vary. We find that as network size increases, the marginal utility of participating in a risk-sharing contract is decreasing in network size, but is decreasing at a faster rate in the state when a household wins a public prize. This, combined with the cost

[7] There are four additional combinations that can occur in principle: bb, vv, bv, and vb. We are primarily interested in analyzing the transfer behaviors of lottery winners to those who did not win a lottery, thus we exclude these four states from our analysis to preserve simplicity.

of maintaining a gift-giving link will result in a gift-giving "shutdown" — beyond a certain network size threshhold, if requests for gifts are too large, then the household will not give any gifts. Additionally gift-transfers will in most cases be **larger** when a household wins a privately-revealed prize.

For the purposes of the simulation, we use log-utility for both household 1 and 2's single-period utility over consumption and use the following values for the model parameters. When a household wins a prize their income is equal to 2, e.g., $y_1(zv) = 2$, otherwise income is equal to 1, e.g., $y_1(zz) = 1$. Warm-glow altruistic capacity is set at $\overline{\gamma}_1^W = \overline{\gamma}_2^W = 2.5$ for both households. Transition probabilities are $\pi_{zz} = 0.3$, $\pi_{zv} = \pi_{zb} = \pi_{vz} = \pi_{bz} = 0.175$, which reflect that the most probable outcome is the case in which neither household wins a prize (zz) — all other states transpire with equal probability. When a household receives a publicly revealed prize, it will receive gift requests from all network members, i.e., $p_1(zb) = p_2(bz) = 1$. Otherwise, the probability that any given gift-network household requests a gift is $p_1(zz) = p_2(zz) = p_1(bz) = p_1(vz) = p_1(zv) = p_2(zb) = p_2(zv) = p_2(vz) = 0.2$. Finally, the discount rate is set to $\delta = 0.65$ for both households.

Without loss of generality, we focus our analysis on household 1's behavior. Figure 2 shows the evolution of the optimal (log) contract intervals as network size increases. At low network-size values, less than 4, the contract intervals overlap and are unchanging — they are unchanging because we limit warm-glow altruism towards household 2 to a maximum of 0.5. Once network size increases beyond 4, the influence of warm-glow altruism decreases in the state in which household 1 wins a publicly revealed lottery — zb. The lower- and upper-bound intervals start to increase until they no longer overlap with state zz and then with state zv. In our example, the contract intervals in state zz and zv overlap over the entire domain in figure 2.

Figure 3 shows the resulting discounted lifetime expected utility of such a contract when the initial state is either zv or bz and when household 1 extracts all the possible surplus — in other words, in the initial state, we select $\lambda(h_1) = \underline{\lambda}(s_1)$ since household 1 extracts the highest surplus when household 2's surplus is set to zero. Here, we see that discounted utility in state zb is less than discounted utility in state zv throughout the domain — this is due to the lower warm-glow altruism one experiences when encumbered with a higher number of gift-requests. Additionally, discounted utility decreases at a faster rate in the zb state until the zz and zb contract intervals cease to overlap — at this point, there is a slight jump in discounted utility in the zb state. However, after this jump, utility in the zb state continues to decrease at a faster rate. Figure 3 also includes a plot of the cost of maintaining one's gift-giving ties, $\alpha(g_1)$. Once discounted utility falls beneath this line, household 1 will shut down all giving to other households when state zb is realized.

2.2 Model Implications

These features of the model lead to our first empirical prediction:

Prediction 1 (The Shut-down Hypothesis) *Households with large gift-giving networks that experience positive publicly-revealed income shocks are more likely to "shut down," resulting in lower levels of transfers to others.*

Figure 4 uses gift transfers between households 1 and 2 to show the empirical implications of the shut-down hypothesis. Notice that at low values for gift-networks, household 1 transfers the same amount to household 2 regardless of being in state zv or zb. However, as the network size increases, transfer amounts start to decrease until they are equal to zero at the shutdown threshold and beyond. This relationship leads to two additional empirical implications:

Prediction 2 (Privately Revealed Prize = Higher Average Transfer Value) *The average gift value is higher in households that win privately revealed prizes than households that receive publicly revealed cash prizes.*

Prediction 3 (Publicly Revealed Prize = Higher Number of Gifts Given) *The average number of gifts given is higher in households that win publicly revealed prizes prior to passing the shutdown threshold.*

The above two predictions also imply that the total value of gifts out of households who win publicly revealed prizes are higher than the total value of gifts given from other households prior to the household reaching its shut-down threshold. This is easily shown by multiplying the average transfer value by the number of gift-obligations in period t (see appendix figure C.1 for a graphical representation). The prediction can be stated as:

Prediction 4 (Prior to shut-down = Larger Volume of Transfers After Public Prize) *Prior to reaching their shut-down threshold, the volume of gifts given by households who win publicly revealed income will be larger than the volume of gifts given by households who win privately revealed income.*

So far we have discussed how the model generates predictions regarding the gift-transfer behavior of household 1. Naturally, if household 2 receives gifts from household 1, we should be able to symmetrically identify changes in household 2's consumption as a function of household 1's lottery winnings. This implies that household 2's consumption levels will be higher on average when their gift-giving network wins a cash prize. However, since transfers are predicted to be higher when the peer household wins a private lottery, it is likely that the effect will only be observed in such a state. Furthermore, since household 1's marginal utility is decreasing in household 2's consumption, we should see stronger patterns of gift-giving through the private lottery when the income gap between households 1 and 2 is large. It is straightforward to show via simulation that average transfer sizes increase as the gap between 1 and 2's per-period income increases.[8] This leads to the final prediction:

Prediction 5 (Consumption Increasing in Others' Winnings) *A household's per-capita consumption is an increasing function of its peer-network's average private lottery winnings. It may be an increasing function of its peer-network's public lottery winnings if its peers do not experience a shutdown in giving (i.e., peers have small gift-giving networks).*

3 Data and Descriptive Evidence

We combine a field experiment with household surveys to construct the data used in the analysis. The field experiments were conducted between March and October 2009 in conjunction with a year-long household survey in four communities in Akwapim South district of Ghana's Eastern Region. This district lies some 40 miles north of the nation's capital, Accra, but is sufficiently far away that only a handful of respondents commute to Accra for work. The sample consists of approximately 70 households from each of the four communities.[9] Individuals in the sample include the household head and his spouse.[10] There

[8] Similarly, one could add one more income-realization possibility to the state space — negative income shock — to generate relevant predictions. This would likely over-complicate the model for our purposes so we have left such simulations out of this paper.

[9] The survey was part of a three-wave panel, the first two waves having been conducted in 1997-98 (e.g., in Conley and Udry (2010)) and 2004 (Vanderpuye-Orgle and Barrett, 2009). Slightly more than half of the 70 households were part of the initial 1997-98 sample, and the rest were recruited in January 2009 using stratified random sampling by the age of the household head: 18-29, 30-64, 64+. the shares of households whose head was in each of these age categories corresponded to the community's population shares. In the original sample, and in the 2009 re-sampling, we selected only from the pool of households headed by a resident married couple. However, we retained households from the 1997-98 sample even if only one of the spouses remained.

[10] Some men in the sample have two or three wives, all of whom were included. However, for the sake of simplicity we refer to households throughout the text as having two spouses.

are between 7 and 12 sampled 'single-headed households' in each community. In total the sample used in our study includes 606 individuals comprising 325 households in each of the four communities.

Survey. Each respondent was interviewed five times during 2009, once every two months between February and November.[11] Each survey round took approximately 3 weeks to complete, with the two survey teams each alternating between two villages. The survey covered a wide range of subjects including personal income, farming and non-farm business activities, gifts, transfers and loans, and household consumption expenditures. Each round, both the husband and wife heading each household were interviewed separately on all of these topics. Our data set is assembled mainly using information contained in the expenditure, gift and social network modules of the survey.

The expenditure module asked detailed information on the quantities and values purchased of a long list of items with broad categories including home produced food consumption, purchased food consumption, school-related expenditures (fees and complementary goods such as uniforms), medical expenditures (medicine and health fees), among others. Referring to the month prior to the interview, we asked each spouse about his or her own expenditures, those of their partner, and about expenditures of the household as a whole. Appendix table B.1 collects summary statistics. Of note in this table is the observation that there is within-household specialization in food expenditures: household heads (mostly males) are more responsible for procuring food produced on the household's farm while the spouse (mostly females) are responsible for purchasing food to supplement home-produced food. Given that the household head and spouse seem to coordinate around total household food consumption, we aggregate variables at the household level.[12]

In the gifts module, respondents were asked to report any gifts (in cash or in kind) given and received during the past two months, obtaining information on the counterparty's location and relationship to the respondent. The value of the gift given and an estimated value for in-kind gifts were also recorded. We pool monetary and in-kind gifts into a single measure and drop incidents within-household transfers — i.e., gifts transferred to one's spouse. Since we are primarily interested in gifts received from others who are potential winners of lottery prizes, we drop observations of gifts received from others who do not

[11] For details regarding interview timing and survey instruments, see Walker (2011).

[12] For food expenditures, this involves summing the household head and spouse's "own food" consumption. Each individual provides his or her own list of gifts given/received and is not asked to report spouse's gift information, so household aggregation is a straightforward sum of these lists for gift-related variables. See Castilla and Walker (2013) for a study analyzing how information asymmetry influences spending decisions within the household using the same data.

reside within the village.

After selecting the sample but before collecting baseline data a detailed survey of respondents' social contacts was conducted. Each respondent was asked in turn (and in random order) about every other respondent in the sample from his or her community. More specifically, the social network module of the survey asked whether they knew the person, by name or personally, how often they saw them, whether they were related, what they perceived the strength of the friendship to be, whether they had ever given or received a gift to or from the person, and whether they would trust the person to look after a valuable item for them. Due to the nature of the data, we can confirm the bi-directional nature of reciprocal gift-networks. We do this by merging individual i's response regarding j's gift-giving behavior with individual j's response of i's gift-giving behavior. We establish a reciprocal gift-link between any two individuals when both state that they have ever received and given gifts to the other individual. This substantially reduces any concerns regarding the measurement error of the network. We consider that two households are linked to each other in a reciprocal gift-giving relationship if at least one pair of the potential combinations engages in mutual gift-giving. For example, there are four total combinations between household A and B when both households have one male (M) and female (F) head/spouse: AM-BM, AM-BF, AF-BM, AF-BF. If a reciprocal gift-giving links exists between at least one of these pairs, then we state that household A and B have a reciprocal gift-giving relationship. Otherwise, no such link exists. Household-aggregated measures that form the basis of our analysis are represented in table 1. On average, each household has roughly five members and has reciprocal gift-giving relations with 11.5 other households. Roughly 13% of the households in do not have in reciprocal gift-giving links with any other household in the sample. Across the five rounds of data, households give and receive 1.58 and 0.58 gifts respectively to any other household over the course of two months. The average total value of the gifts given (received) is 20.02 (12.58) GH¢. Household per-capita food consumption is reported in the third panel of table B.1. The total household per-capita food consumption in our sample is 26.64 GH¢, 68% of which is purchased food.[13]

Experimental Data. The first round of the survey was designed as a baseline, therefore no lottery took place in that round. One week before each subsequent round we visited each village to distribute prizes to selected respondents. Twenty prizes were allocated in each community, in each of the four lottery rounds, so that in all 320 prizes were given. Over the four lotteries, approximately 42 percent of individuals and 62 percent of households won at least one prize. Ten of the prizes were divisible in the form of cash, whereas the other ten

[13]Seasonal conditions account for inter-temporal variation in the amount/value of purchased food in the household.

were in the form of livestock. Of these, five each were allocated publicly by lottery, and the other five (identical in type) were allocated in private, by lucky dip. The values of the prizes varied from GH¢10 to GH¢70 as described in figure 5.[14] The prizes were of a substantial size - the largest prize is equivalent to a month's worth of food consumption for an average household with five members. In aggregate, each community received GH¢370 of cash in each round to use however they would like.

The lotteries took place one week before the commencement of the survey interviews. We took great care to make clear to participants that the allocation of prizes was random, and that each individual had an equal chance of winning in each round. A village meeting was held in a central area of the community, and all respondents were invited to attend. A small amount of free food and drink was provided as an incentive to come. Attendance at the meetings was generally around 100 people; roughly half of the respondents appeared for each meeting.[15] There were usually a number of non-respondents at these meetings as well, including many children. At each gathering we thanked the participants for their continued support. We explained that respondents had a chance to win one of 20 prizes that day, framing the prizes as a gratuity for their participation in the survey.[16] We then proceeded to draw winners for the ten public prizes (without replacement) from a bucket containing the names of the survey respondents. A village member not in the sample was chosen by the villagers to do the draw, in order to emphasize that the outcomes were random. Each winner was announced to the group, and asked to come forward to receive their prize. The prizes were announced and displayed clearly before being awarded. Respondents who were absent at the time of drawing were called to pick up their prize in person, if possible. Unclaimed prizes were delivered in person to the winner after the lottery. After the public lottery prizes were distributed, we conducted a second round in private. Respondents were asked to identify themselves to a survey worker, who took their thumbprint or signature and issued

[14] In this paper, we are primarily interested in transfers of divisible windfall gains, thus we focus our attention on cash lottery winnings. The livestock were purchased in Accra on the morning of the lottery and transported to the community. The value of the price differed according to the type of livestock: Chickens (10¢), two chickens (20¢), small goat (35¢), medium goat (50¢), and large goat (70¢). Different households may face different transaction costs, so the value of livestock, as opposed to cash, is heterogeneous across households, which further complicates the use of livestock in the analysis. Additionally, in this study context, it is more difficult to 'privately' grant lottery winners a large goat than it is to privately grant them the same amount in cash.

[15] Around 125 of the 150 respondents in each community appeared for the privately revealed lottery, some of them arriving before or after the public meeting.

[16] Respondents signed an informed consent form at the start of the survey, explaining how they would be remunerated for their participation in the survey. Entry in the lottery and lucky dip was part of this remuneration. In addition to the chance of winning a prize, each respondent was given a small amount of cash for their participation, which varied across rounds. This gift was used as an endowment in a public goods experiment as part of a separate study.

them with a ticket displaying their name and identification number. They then waited to enter a closed school room, one at a time, where an enumerator invited them to draw a bottle cap without replacement from a bag. There was one bottle cap for each of the N respondents in the community. Of these, N - 10 were non winning tokens (red colored), and ten were winning tokens, marked distinctively to indicate one of the ten prizes listed in figure 5.[17] Those who drew winning tokens were informed immediately that they had won a prize, which was identified to them, and were told that they did not have to tell anyone else that they had won. We emphasized that the survey team would not divulge the identities of winners who won in private. Cash prizes were given to the winners immediately and winners often hid their prizes in their clothes before leaving the room. The survey interviews in each round commenced one week after the lottery, deliberately delayed to allow winners to receive their prize and do something with it. The interviews took place in no specified order throughout the following three weeks, so that some winners were interviewed a week after receiving their prize, and others up to four weeks afterward.

Appendix table B.3 presents balance tests conducted on variables collected at baseline according to whether one member of the household won any of the public or private lottery at any point over the course of the year. 119 of the households in the study are thus in our "treatment" group while the remaining 190 did not win a cash prize. We also separate the test according the households that won the privately revealed vs. publicly revealed lottery. The table suggests that randomization was successful — of the 21 tests along which we seek to reject balance, one is significant at the 5% level and another is significant at the 10% level. For the others, balance cannot be rejected at the 10% level.

To calculate gift-network lottery winnings, we simply take the average cash winnings (private vs. public) of each household's gift-network. In other words, for every household i out of N, private (replaceable with public) network lottery winnings are

$$\overline{\text{Private}}_{it} = \sum_{j=1}^{N} \frac{\text{Private}_j \times \mathbb{1}(g_{ij} == 1)}{\sum_{j=1}^{N} \mathbb{1}(g_{ij} == 1)},$$

where $g_{ij} = 1$ if there is a reciprocal gift-giving relationship between households i and j (0, otherwise), $\text{Private}_j \in \{0, 10, 20, 35, 50, 70\}$ are the values of cash prizes household j can win

[17] Care was taken to shuffle the bottle caps after each draw, and to prevent respondents from seeing into the bag. If a respondent drew more than one bottle cap, those caps were shuffled and the respondent was asked to blindly select one of them. Respondents were shown a sheet relating the tokens to the prizes (See Walker (2011)). At the conclusion of the day, tokens that had not been drawn were counted and the remaining prizes allocated randomly among the non-attending respondents using a computer. There were usually 25-30 non-attendees and less than three prizes remaining.

and $\mathbb{1}$ represents the indicator function.[18] The bottom two panels of table 1 present the average (log) value of own and network cash winnings and show that average prize winnings roughly represent the expected value of the cash prize of all households in the village sample.

4 Empirical Investigation

The unique features of our experimental design allows us to bring the model predictions to the data in a straightforward manner. Let y_{it} be the outcome of interest: either the amount of round t gifts distributed or the number of round t gifts distributed by household i. The shutdown hypothesis (Prediction 1 in Section 2) can be investigated using the following regression:

$$\begin{aligned} y_{it} = \alpha &+ \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} \\ &+ \beta_{vg} \text{Private}_{it} \times \text{Net-size}_i + \beta_{bg} \text{Public}_{it} \times \text{Net-size}_i \\ &+ \text{hh}_i + \text{r}_t + \epsilon_{it}, \end{aligned} \tag{14}$$

where β_v captures the extent to which round t gift-behavior is influenced by round t privately revealed lottery winnings and β_b captures the influence of publicly revealed lottery winnings. Net-size_i is household i's reciprocal gift-network size. hh_i captures household fixed effects and r_t captures round fixed effects. Importantly, notice that household fixed-effects control for all time-constant household factors including the size of its gift-network. Given the distribution of the outcome variables, the specific estimator will place restrictions on the error term, ϵ_{it}. Specifically, when the outcome variable is the (log) amount of gifts given, we use the tobit estimator where we integrate out censored observations equal to zero. The number of gifts given follows a poisson distribution, so we use a poisson estimator to estimate the coefficients of interest under this dependent variable. Predictions 2 and 3, that do not depend on heterogeneity in network size, can simply be tested by setting the interaction terms equal to zero.

An empirical investigation of the model implication in terms of consumption (Prediction 5, Section 2) requests instead to relate household i's consumption behavior against the average lottery winnings of i's gift network — i.e.,the average network treatment effect on

[18] In one round, both the household head and spouse within the same household won the lowest of two public lotteries. Hence, for this household, the prize winnings amounted to 30.

consumption (defined in Section 3). Thus, we estimate the following equation:

$$y_{it} = \alpha + \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} + \beta_{vn} \overline{\text{Private}}_{it} + \beta_{bn} \overline{\text{Public}}_{it} + \text{hh}_i + \text{r}_t + \epsilon_{it}, \tag{15}$$

where y_{it} is log per-capita household food consumption. We do not expect food consumption to be an increasing linear function in network lottery winnings. However, we do expect that households with lower levels of period-specific food-consumption will receive more support from their network. Therefore, we opt to use a quantile regression estimator to examine effects at different locations along the consumption distribution. Table 2 summarizes how coefficients in each of the estimation equations link to predictions from our theoretical model.

Table 3 contains the estimation results of model 14 with three different outcome variables, with and without interaction terms. The negative coefficient in the fourth row (β_{bg}) of columns 4-6 indicates that individuals winning the public lottery are associated with lower levels of transfers the larger is their gift network size. This is in line with the shut-down hypothesis predicted by our model (prediction 1). Notice that in column 3, the first row coefficient (β_v) is larger than the second row coefficient (β_b) — this confirms the hypothesis that each individual gift is, on average, larger for individuals winning the private lottery (prediction 2). Notice the difference in the second row in columns 2 and 5. In column 5, β_b is positive and large but is insignificant from zero in column 2. This indicates that without including the interaction effect, we underestimate the number of gifts given for someone with a relatively small gift-network after winning the public lottery. This suggests that prediction 3 is confirmed. Furthermore, the second row in columns 1 and 4 show a similar story — furthermore, β_b is larger than β_v in column (4), which is consistent with the pattern we observe in figure C.1 (prediction 4). Figure 7 provides a non-parametric test of prediction 4 — it is consistent with figure C.1 generated by model simulations. Finally, figure 6 estimates equation 14 as a fourth-order polynomial (take powers 0-4 on the interaction coefficient and sum them together) and shows a shut-down network size of roughly 15 individuals in the mutual gift-giving network (taking point estimates as given).

Turning to the model implications in terms of consumption, we depict graphically the results of the quantile estimation of equation (15) in Figure 8. We use observations from the first three rounds of data — the hungry season in Ghana when unlucky households are most likely to require help from others. The lower the per-capita food consumption, the more likely one is to increase consumption when their friends win the private lottery winning — the coefficient on private average network lottery winnings is positive and greater than zero for analyzed quantiles less than the 50th percentile. The same cannot be said about the public lottery winnings of one's gift network. Here, the coefficient is only positive for households whose per-capita food consumption falls below the 10th percentile. Furthermore, consumption increases by a smaller margin at the lowest percentiles when gift network members win public relative to private lotteries. These results are all in line with prediction 5 and likely flow from the fact that households who win publicly revealed lotteries are subject to social taxation and are unable to assist connected households who exhibit great need.

5 Robustness and Extentions

The extremely detailed micro-structure of our data offers an alternative strategy to test the model predictions and look further into underlying mechanisms. Let g_{ij} be a dyadic variable taking value 1 if household i has an established reciprocal gift-giving link with household j, equation 14 takes the following form:

$$\begin{aligned} y_{ijtv} = \alpha &+ \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} \\ &+ \beta_{vg} \text{Private}_{it} \times \text{Net-size}_i + \beta_{bg} \text{Public}_{it} \times \text{Net-size}_i \\ &+ \gamma \text{Net-size}_i + \text{village}_v + \text{r}_t + \epsilon_{ijt} \end{aligned} \tag{16}$$

where the outcome variable measures (log) gifts amounts and numbers given from i to j and village fixed effects are included (instead of household fixed effects). Although transfers among dyads within our sample constitute a small share of total transfers reported in the survey's gift-module, this model specification provides a robustness check and, more importantly, allows us to examine the extent to which households target each other for gifts. Further, we can differentiate between dyads who have a mutual gift-giving relationships and those that do not.

Recall, altruistic preferences imply that household i's marginal utility is an increasing function of the relative suffering of household j — in other words, the lower household j's consumption levels relative to i, household i is more incentivized to transfer resources to household j under altruistic preferences. To examine this prediction in our data, we can estimate the following equation via dyadic analysis:

$$\begin{aligned} y_{ijtv} = \alpha &+ \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} \\ &+ \beta_{v\chi} \text{Private}_{it} \times (\hat{\chi}_i - \hat{\chi}_j) + \beta_{b\chi} \text{Public}_{it} \times (\hat{\chi}_i - \hat{\chi}_j) \\ &+ \gamma (\hat{\chi}_i - \hat{\chi}_j) + \text{village}_v + \text{r}_t + \epsilon_{ijt} \end{aligned} \tag{17}$$

where we have replaced the "Net-size$_i$" interaction term with the difference in household i

and j's food shocks, $\hat{\chi}_i - \hat{\chi}_j$. The "more positive" i's food consumption shock is relative to j's, the higher the difference. If i holds altruistic preferences over j, then $\beta_{\{v,b\}\chi}$ will be positive. We measure food shocks in the following way. Given the panel nature of our data, we measure deviations from round-adjusted average household per-capita food consumption by recovering the OLS residual error term from the equation:

$$\text{Per-Capita Food}_{it} = hh_i + r_t + \chi_{it}$$

where hh_i and r_t are respectively household and round fixed effects. The resulting variable, $\hat{\chi}$ measures household-round specific shocks to per-capita food consumption.

Table 4 and 5 reports the OLS estimation results of model 16 without and with interaction terms, respectively. Each of these tables splits the sample into those dyads who have a mutual gift-giving relationships and those that do not. Overall, evidence is consistent with the model and suggests that private lottery winnings are transferred to individuals who are already in one's gift-giving network. Public lottery winnings are more likely to go to individuals who are not in one's mutual gift-giving network, suggesting that individuals outside of this network will solicit the lottery winner for help. The shut-down hypothesis is again confirmed, although standard errors are large. More specifically, column 2 of table 6 shows that while all individuals in one's network receive more gifts when individual i wins the privately revealed lottery, the largest gifts are reserved for those households j whose food-shocks are largest. In other words, household i gives more to households whose realized food consumption is much lower than household i's realized food consumption. This same pattern does not take place when household i wins the publicly revealed lottery. Instead, when household i wins a publicly revealed lottery, he gives to households who are not in its gift network only when they experience a negative food shock relative to household i. This is totally in line with model prediction 5. Table B.6 estimates a triple-interaction that pools all observations and shows evidence consistent with table 6.

We conclude this section with an exploration of the channels motivating solidarity network. In particular, we would like to investigate whether we find evidence consistent with altruism being an important driver, besides insurance. If social solidarity networks indeed smooth members' consumption by distributing income shocks across the network, the familiar prediction, following Townsend (1994), is that the inter-temporal change in one member's consumption should track one-for-one the average consumption change over the same period within the rest of one's network. Within our model, the testable prediction is the null that the coefficient relating a survey respondent's period-on-period change in log consumption to the contemporaneous change in network average consumption equals one. Within our model, across the full social solidarity network we expect to reject the null in favor of the one-sided alternate hypothesis that the coefficient is less than one but also to reject the null that change in consumption is uncorrelated, in favor of the one-sided alternate hypothesis that they are positively correlated. This occurs because of the shut down hypothesis and because private winnings will not get shared with networks members who do not exhibit great material need.

Table 7 reports results of those hypothesis tests. We show that limited risk pooling occurs within the full network. The point estimate of 0.23 is statistically significantly greater than 0 only at the 10 percent level and one can easily reject the null that it equals 1.00. Meanwhile, the respondent's own winnings, whether private or public, and the average winnings within one's network are statistically insignificantly related to a respondent's consumption volatility once one controls for consumption volatility within one's network, consistent with the altruism in networks model of Bourlés et al. (2017). From this result, we can conclude that there are multifaceted drivers of gift-giving in this gift-network that may include limited degrees of risk pooling, but likely involve solidarity among network ties. In summary, our evidence points toward th efact that it is hard to argue that the solidarity network is motivated mainly by insurance. Combined with the significant giving from private winnings, it certainly appears that altruism and taxation are more compelling explanations. Insurance may play a role, but it hardly seems a primary role.

6 Conclusion

We analyse altruistic preferences in networks by examining a model of risk sharing under imperfect commitment where the impurely altruistic gains to giving to others diminish with the number of transfers one makes. Giving is costly, and stochastic income has both publicly observable and unobservable components. Contrary to the canonical informal insurance model, in which bigger networks and observable income are preferable, our model predicts that unobservable income shocks may facilitate altruistic giving that better targets the least well off within one's network and that too large a network can overwhelm even an altruistic agent to cease giving. Full risk-pooling is maintained within the network that remains so long as an agent does not exist the arrangement. We take these predictions to a unique data set from southern Ghana. We couple observations of gift-giving networks with experimental cash windfall gains - randomized between private and publicly observable payouts - repeated every other month for a year to analyze transfer flows among households. We find four striking results. First, on average, more gifts are given out of private cash winnings than public cash winnings, signaling that altruistic preferences - not just self-interested behavior within an endogenously enforceable insurance scheme - must be a significant driver of inter-household transfers. Second, winners of privately revealed prizes target giving to the neediest households within their networks, indicating greater social welfare gains from altruistic transfers than from insurance transfers. Third, winners of publicly revealed cash prizes do not make transfers when they have large networks; they break the informal contract due to network size. Fourth, conditional on transfers flowing within one's network, we cannot reject the null of full risk pooling. These results highlight the limits to social networks as channels for managing income shocks as well as the trade-offs inherent to transparency in transfer programs. Although observability of income is essential in informal insurance arrangements among purely self-interested agents, observability may impede altruistic agents' ability to focus their giving on the most needy as they are compelled to respond to demands for assistance from the less needy within their network.

References

Advani, Arun. (2017). 'Insurance Networks and Poverty Traps', *Working Paper* .

Alderman, Harold. (2002). 'Do local officials know something we dont? Decentralization of targeted transfers in Albania', *Journal of Public Economics* 83(3), 375 – 404.
URL: *http://www.sciencedirect.com/science/article/pii/S0047272700001456*

Ambrus, Attila, Gao, Wayne Y and Milán, Pau. (2017). 'Informal risk sharing with local information', *Unpublished Manuscript* .

Ambrus, Attila, Mobius, Markus and Szeidl, Adam. (2014). 'Consumption Risk-Sharing in Social Networks', *American Economic Review* 104(1), 149–82.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.104.1.149*

Andreoni, James. (1990). 'Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving', *The Economic Journal* 100(401), 464–477.
URL: *http://www.jstor.org/stable/2234133*

Angelucci, Manuela and De Giorgi, Giacomo. (2009). 'Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?', *American Economic Review* 99(1), 486–508.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.99.1.486*

Angelucci, Manuela, De Giorgi, Giacomo and Rasul, Imran. (n.d.). 'Consumption and Investment in Resource Pooling Family Networks', *The Economic Journal* 0(0).
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/ecoj.12534*

Ashraf, Nava and Bandiera, Oriana. (2017). 'Altruistic Capital', *American Economic Review* 107(5), 70–75.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.p20171097*

Bourlés, Renaud, Bramoullé, Yann and PerezRichet, Eduardo. (2017). 'Altruism in Networks', *Econometrica* 85(2), 675–689.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13533*

Bowles, Samuel and Gintis, Herbert. (2002). 'Social Capital And Community Governance', *The Economic Journal* 112(483), F419–F436.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0297.00077*

Castilla, Carolina and Walker, Thomas. (2013). 'Is Ignorance Bliss? The Effect of Asymmetric Information between Spouses on Intra-household Allocations', *American Economic Review* 103(3), 263–68.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.103.3.263*

Coate, Stephen and Ravallion, Martin. (1993). 'Reciprocity without commitment: Characterization and performance of informal insurance arrangements', *Journal of Development Economics* 40(1), 1 – 24.
URL: *http://www.sciencedirect.com/science/article/pii/030438789390102S*

Conley, Timothy G. and Udry, Christopher R. (2010). 'Learning about a New Technology: Pineapple in Ghana', *American Economic Review* 100(1), 35–69.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.100.1.35*

Fafchamps, Marcel. (1992). 'Solidarity Networks in Preindustrial Societies: Rational Peasants with a Moral Economy', *Economic Development and Cultural Change* 41(1), 147–174.
URL: *https://doi.org/10.1086/452001*

Foster, Andrew D. and Rosenzweig, Mark R. (2001). 'Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas', *The Review of Economics and Statistics* 83(3), 389–407.
URL: *https://doi.org/10.1162/00346530152480054*

Jakiela, Pamela and Ozier, Owen. (2016). 'Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies', *The Review of Economic Studies* 83(1), 231–268.
URL: *http://dx.doi.org/10.1093/restud/rdv033*

Ligon, Ethan and Schechter, Laura. (2012). 'Motives for sharing in social networks', *Journal of Development Economics* 99(1), 13 – 26.
URL: *http://www.sciencedirect.com/science/article/pii/S0304387811001179*

Ligon, Ethan, Thomas, Jonathan P. and Worrall, Tim. (2002). 'Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies', *The Review of Economic Studies* 69(1), 209–244.
URL: *http://dx.doi.org/10.1111/1467-937X.00204*

Mazzocco, Maurizio and Saini, Shiv. (2012). 'Testing Efficient Risk Sharing with Heterogeneous Risk Preferences', *American Economic Review* 102(1), 428–68.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.102.1.428*

Platteau, J.P. (2000), *Institutions, Social Norms, and Economic Development*, Fundamentals of Development Economics, Harwood Academic.
URL: *https://books.google.com/books?id=IPJRCagTB4QC*

Popkin, Samuel L. (1979), *The rational peasant : the political economy of rural society in Vietnam / Samuel L. Popkin*, University of California Press Berkeley.

Sen, Arijit and Hoff, Karla. (2006), The kin system as a poverty trap?, *in* **Samuel Bowles, Steven Durlauf and Karla Hoff.**, eds, 'Poverty traps', Princeton University Press, Princeton, chapter 4, pp. 95–115.

Simons, Andrew M. (2016), What is the optimal locus of control for social assistance programs? Evidence from the productive safety net programme in Ethiopia, ESSP working papers 86, International Food Policy Research Institute (IFPRI).
URL: *https://ideas.repec.org/p/fpr/esspwp/86.html*

Squires, Munir. (2017). 'Kinship Taxation as an Impediment to Growth: Experimental Evidence from Kenyan Microenterprises', *Working Paper* .

Thomas, Jonathan and Worrall, Tim. (1988). 'Self-Enforcing Wage Contracts', *The Review of Economic Studies* 55(4), 541–554.
URL: *http://dx.doi.org/10.2307/2297404*

Tables

Townsend, Robert M. (1994). 'Risk and Insurance in Village India', *Econometrica* pp. 539–591.

Vanderpuye-Orgle, Jacqueline and Barrett, Christopher B. (2009). 'Risk Management and Social Visibility in Ghana*', *African Development Review* 21(1), 5–35.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8268.2009.00201.x*

Walker, Thomas. (2011). 'Appendix to Thomas Walker's Dissertation', *Unpublished Manuscript* .
URL: *http://barrett.dyson.cornell.edu/research/datasets/ghana.html*

Table 1: Household Summary Statistics

	N	Mean	Sd	5 p-tile	95 p-tile
Fixed Over Time:					
HH size	309	4.94	2.19	2	8
N of HH in Solidarity Network	309	11.56	10.11	0	32
Gifts (last 2 months, GH¢):					
N Gifts Given	1,525	1.60	2.16	0	6
N Gifts Received	1,525	1.45	1.95	0	5
Total Value of all Gifts Given	879	26.77	88.94	1	90
Total Value of all Gifts Received	893	46.38	130.77	1.50	158
Food Consumption (last month, GH¢):					
PC Food Consumption	1,525	26.64	21.70	5.05	65.10
PC Purchased Food	1,525	18.12	19.14	0	48.93
PC Home-produced Food	1,525	8.51	8.53	0	23.45
Own Lottery Winnings (GH¢):					
Cash - Private	1,525	1.84	9.42	0	0
Cash - Public	1,525	1.87	9.55	0	0
Solidarity Network Average Lottery Winnings (GH¢):					
Solidarity Network Cash - Private	1,525	1.81	4.52	0	8.33
Solidarity Network Cash - Public	1,525	1.74	4.29	0	8.33

Gift Network data missing for a subset of observations. N of gifts given/received equal zero if none given/received. Value of gifts contingent on having received at least one. Household food consumption (total) sums the head of households and spouse's response. Solidarity network lottery winnings multiply the vector of lottery winners by the row-normalized network adjacency matrix (result is average networks' lottery winnings). Network values represent log transformations of original winnings/averages.

Table 2: Linking Empirical Analysis to Theoretical Predictions

Equation	Predictions				
	1 Shut-Down	2 Average Gift Value	3 Gift Number	4 Total Value	5 Consumption*
14|Net-size $= 0$		$\beta_v > \beta_b$	$\beta_b = \beta_v$	$\beta_b = \beta_v$	
14	$\beta_b > 0, \quad \beta_{bg} < 0$		$\beta_b > \beta_v$	$\beta_b > \beta_v$	
15					Quantile Regression**

Notes: * Tests referring to prediction 5 suggest that those households who have lower levels of food consumption relative to network contacts who win the lottery will be more likely to receive transfers and will receive higher values per transfers. ** Indicates coefficient on β_{vn} should be larger at lower quantiles.

Table 3: Prize Winnings Influence Gift-Giving

	No Interaction			Shut-down Hypothesis		
Dep. Var.: Gifts-Given	Value	Number	$\frac{\text{Value}}{\text{Number}}$	Value	Number	$\frac{\text{Value}}{\text{Number}}$
Private Cash Winnings	0.010*	0.017**	0.007**	0.010*	0.016**	0.007*
	(0.005)	(0.007)	(0.004)	(0.005)	(0.007)	(0.004)
Public Cash Winnings	0.002	0.005	-0.001	0.012	0.026***	0.004
	(0.005)	(0.007)	(0.003)	(0.007)	(0.010)	(0.005)
Public Cash Winnings × N Mutual Gifts				-0.001*	-0.002***	-0.001
				(0.001)	(0.001)	(0.000)
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1602	1602	1602	1602	1602	1602

*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Dependent Variable equals log value of gifts given in columns 1 and 4, number of gifts given in column 2 and 5, and log value per gift ($\frac{\log(\text{Total Value})}{\text{Total Number}}$) in columns 3 and 6. Household and Round Fixed Effects Included in Every Specification. Coefficients estimated using Tobit estimator with a lower bound of zero (no upper bound). Log transformation of variables adds one to original value so that zero values are preserved under log transformation.

TABLE 4: Dyadic Regressions

	Mutual		All Other Sample Links	
	$\text{Log}(\text{Amount}_{ijt})$	Number_{ijt}	$\text{Log}(\text{Amount}_{ijt})$	Number_{ijt}
$\text{Lottery-Private}_{it}$	0.322*	0.296***	-0.462	-0.241*
	(0.190)	(0.081)	(0.289)	(0.145)
$\text{Lottery-Public}_{it}$	0.168	-0.013	0.136	0.086
	(0.185)	(0.095)	(0.202)	(0.101)
Village FE	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
N	19330	19308	114645	111453

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable in even columns equals amount of actual gift given from i to j in any given round t; Odd columns equals number of gifts. 'Network Size_i' indicates household i's gift-network size (any type of gift-relation). Odd columns are tobits with lower bound of zero. Even columns are poisson regressions. Columns 1-2 only include links (i and j) with mutual gift-relations at baseline (178 actual gifts given during the 5 rounds). Columns 3-4 include all oher links (180 actual gifts given).

Table 5: Dyadic Regressions - Shutdown

	Mutual		All Other Sample Links	
	Log(Amount$_{ijt}$)	Number$_{ijt}$	Log(Amount$_{ijt}$)	Number$_{ijt}$
Network Size$_i$	-0.017	-0.010	0.052	0.023
	(0.023)	(0.013)	(0.036)	(0.015)
Lottery-Private$_{it}$	0.607**	0.330***	0.239	0.083
	(0.301)	(0.125)	(0.340)	(0.156)
Lottery-Public$_{it}$	0.495	0.131	0.502**	0.298**
	(0.365)	(0.182)	(0.247)	(0.125)
Lottery-Private$_{it}$ $\times$ Network Size$_i$	-0.016	-0.002	-0.112***	-0.052***
	(0.014)	(0.006)	(0.043)	(0.019)
Lottery-Public$_{it}$ $\times$ Network Size$_i$	-0.019	-0.009	-0.040*	-0.026**
	(0.020)	(0.011)	(0.023)	(0.013)
Village FE	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
N	19330	19330	114645	114645

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable in even columns equals amount of actual gift given from i to j in any given round t; Odd columns equals number of gifts. 'Network Size$_i$' indicates household i's gift-network size (any type of gift-relation). Odd columns are tobits with lower bound of zero. Even columns are poisson regressions. Columns 1-2 only include links (i and j) with mutual gift-relations at baseline (178 actual gifts given during the 5 rounds). Columns 3-4 include all oher links (180 actual gifts given).

Table 6: Dyadic Regressions - Food Shocks

	Mutual Gift		All Other Sample Links	
	Log(Amount$_{ijt}$)	Number$_{ijt}$	Log(Amount$_{ijt}$)	Number$_{ijt}$
Lottery-Private$_{it}$	0.259	0.295***	-0.560*	-0.285**
	(0.205)	(0.092)	(0.302)	(0.144)
Lottery-Public$_{it}$	0.118	-0.038	0.007	0.032
	(0.190)	(0.096)	(0.229)	(0.113)
Food-Shock$_{ijt}$	-0.157	-0.093	-0.061	-0.039
	(0.210)	(0.101)	(0.282)	(0.136)
Lottery-Private$_{it}$ $\times$ Food-Shock$_{ijt}$	0.462***	0.118*	-0.574*	-0.256*
	(0.169)	(0.064)	(0.322)	(0.132)
Lottery-Public$_{it}$ $\times$ Food-Shock$_{ijt}$	-0.265	-0.110	0.618*	0.254*
	(0.304)	(0.153)	(0.336)	(0.140)
Village FE	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
N	18374	18374	92347	92347

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable in even columns equals amount of actual gift given from i to j in any given round t; Odd columns equals number of gifts. 'Food-Shock$_{ijt}$' indicates the difference between i and j estimated food consumption residual ($\hat{\chi}_i - \hat{\chi}_j$ — household and round fixed effects). Odd columns are tobits with lower bound of zero. Even columns are poisson regressions. Columns 1-2 only include links (i and j) with mutual gift-relations at baseline (169 actual gifts given during the 5 rounds). Columns 3-4 include all oher links in the sample (156 actual gifts given).

Table 7: Tests of Full Risk-Sharing

	$\Delta\log(\text{food})_{ivt}$
$\Delta\overline{\log(\text{food}\vert\text{NET})}_{ivt}$	0.232*
	(0.138)
Private Cash Winnings	0.062
	(0.057)
$\overline{\text{Private Cash Network Winnings}}$	0.016
	(0.046)
Public Cash Winnings	-0.038
	(0.035)
$\overline{\text{Public Cash Network Winnings}}$	-0.018
	(0.039)
HH FE	Yes
Round FE	Yes
N	1279

${}^{*}p < 0.1$, ${}^{**}p < 0.05$, ${}^{***}p < 0.01$. Dependent Variable equals first difference of log per-capita household food consumption. Coefficients represent OLS estimators. Variables with line above the variable name indicate household i network's average for each variable.

Figures

Contract Intuition

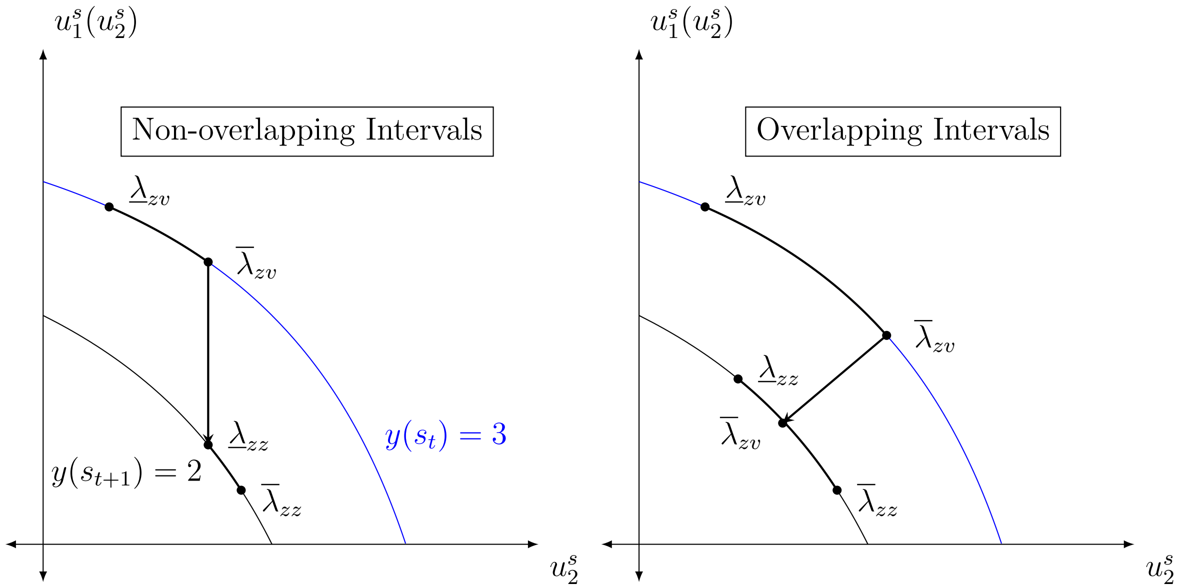

Note: This figure shows how contract intervals relate to the pareto frontier when 1) intervals overlap and 2) when they do not. Values along the x-axis represent household 2's single-period utility and y-axis represents household 1's single-period utility. In state $s_t = zv$, household 1 receives an income of $y_1(zv) = 2$ and household 2 receives an income of $y_2(zv) = 1$ (aggregate income, $y(zv)$, equals 3). In state $s_{t+1} = zz$, both households receive an income of 1 ($y(zz) = 2$). We assume that in period t contracts are such that household 2 receives the entire discounted utility surplus ($\lambda(h_t) = \overline{\lambda}_{zv}$). In period $t+1$, the resulting division of surplus depends on whether or not the contract intervals overlap. When there is no overlap (left-hand side), $\lambda(h_{t+1}) = \underline{\lambda}_{zz}$. When there is overlap, $\lambda(h_{t+1}) = \lambda(h_t) = \overline{\lambda}_{zv}$. Overlapping contracts allow for higher degrees of consumption smoothing over periods.

FIGURE 1: figure

Figure 2: Contract Intervals

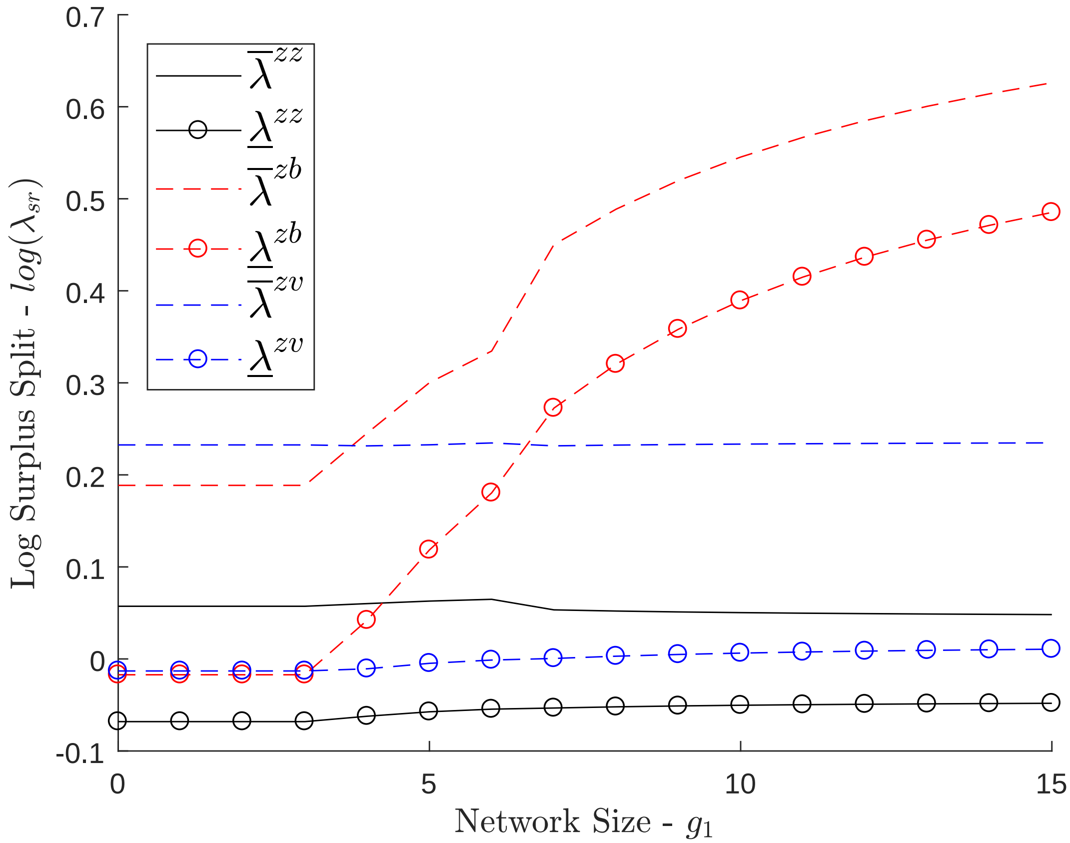

Note: Contract interval solutions as a function of network size with log utility (i.e., $u_1() = u_2() = ln()$. Logged values of λ on the y-axis and network size on x-axis. Contract intervals in state zb increase when $g_1 > 3$ and no longer overlaps with zz when $g_1 > 4$. Furthermore, it is non-overlapping with zv when $g_1 > 6$. The first-best contract (stationary share of aggregate output) is only available when network size is less than three.

Figure 3: Discounted Lifetime Expected Utility

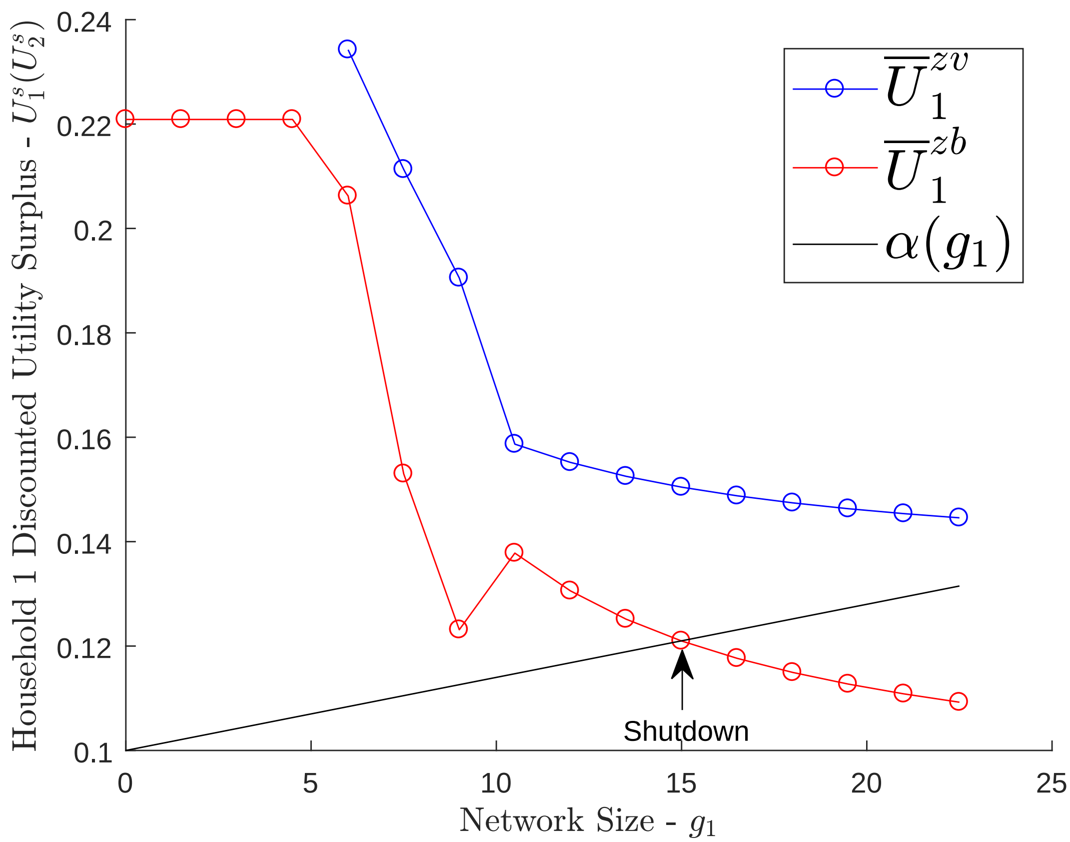

Note: Discounted lifetime expected utility for household 1 when the initial state iz zv vs. zb and when household 1 takes all available surplus from the transfer arrangement (hence, the underline in $\underline{U}_1^s$). Utility values are universally smaller in state zb and decrease at faster rates than state zv throughout. Utility spikes for a single period ($10 < g_1 < 11$), which coincides with the zb contract interval no longer overlapping with zv (see figure 2). The cost of maintaining each network tie, arbitrarily set to $\alpha(g_1) = .1 + .001g_1^{1.2}$ is increasing in network size and intersects with $\overline{U}_1^{zb}$ at a threshold of $g_1 = 15$. Beyond this point, household 1 shuts down all gift transactions when it reaches the zb state. We plot $\overline{U}_1^{zb}$ without the possibility of shutdown; however, utility is $\overline{U}_1^{zb} = 0$ whenever $g_1 > 15$.

FIGURE 4: Amount of Transfer by Network Size

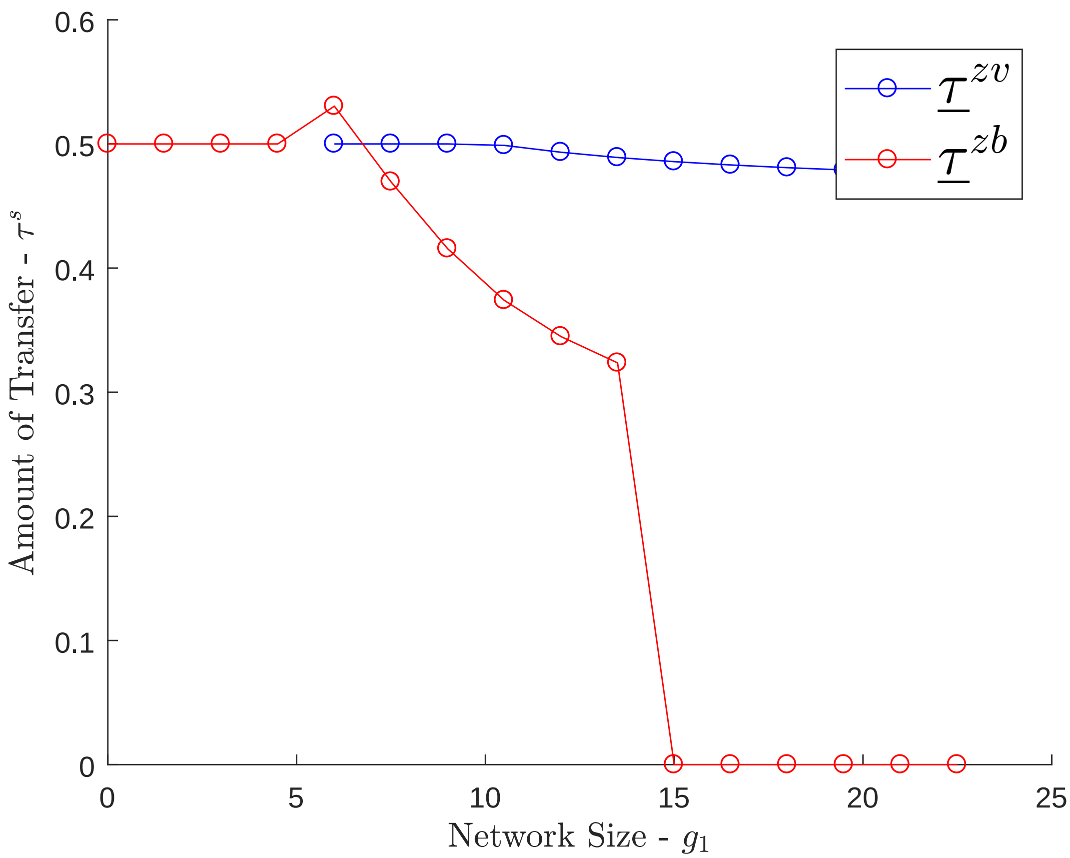

Note: This figure represents transfer amounts $\underline{\tau}^s$ from household 1 to household 2 when household 2 takes the entire share of the surplus (U_1^s is set to zero) and when household 1 wins a cash prize. Thus, it also represents the average transfer amount from household 1 to any other household in its gift network when it wins a cash prize. The average transfer amount is generally smaller when household 1 wins the publicly revealed prize (zb) relative to when it wins the privately revealed prize (zv). Transfers are reduced to zero beyond household 1's shutdown point ($g_1 = 15$).

Figure 5: Experimental Data: Lottery Payouts

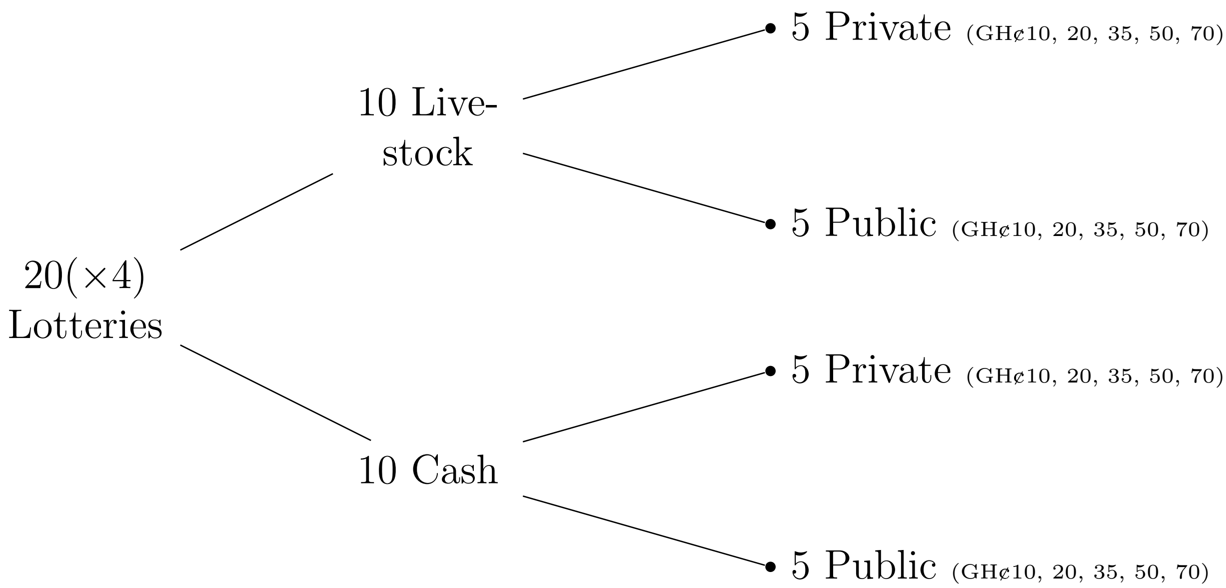

Figure 6: Shut-down Hypothesis on Number of Gifts Given

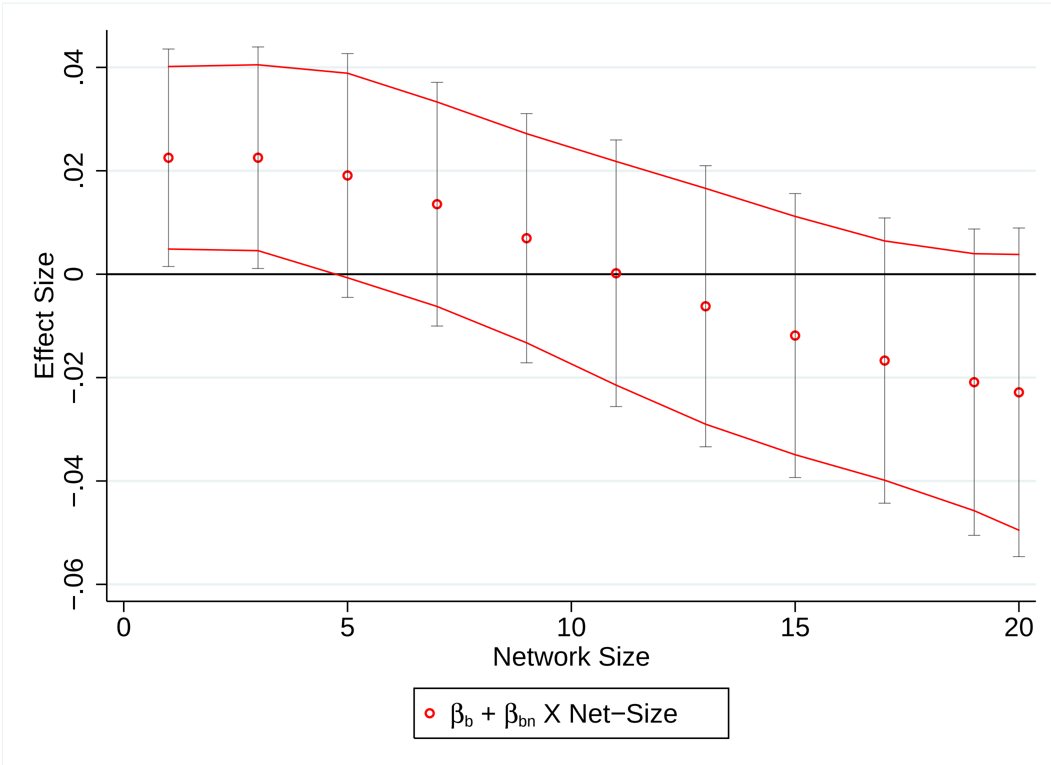

Note: Dependent variable equals number of gifts given. Estimation of equation 14 with the inclusion of 2nd, 3rd, and 4th order polynomial interactions on network-size variable. Results indicate that the coefficient is equal to zero when solidarity network consists of 11 households (within-sample).

Figure 7: Nonparametric Analysis of Shut-down Hypothesis on Total Value of Gifts

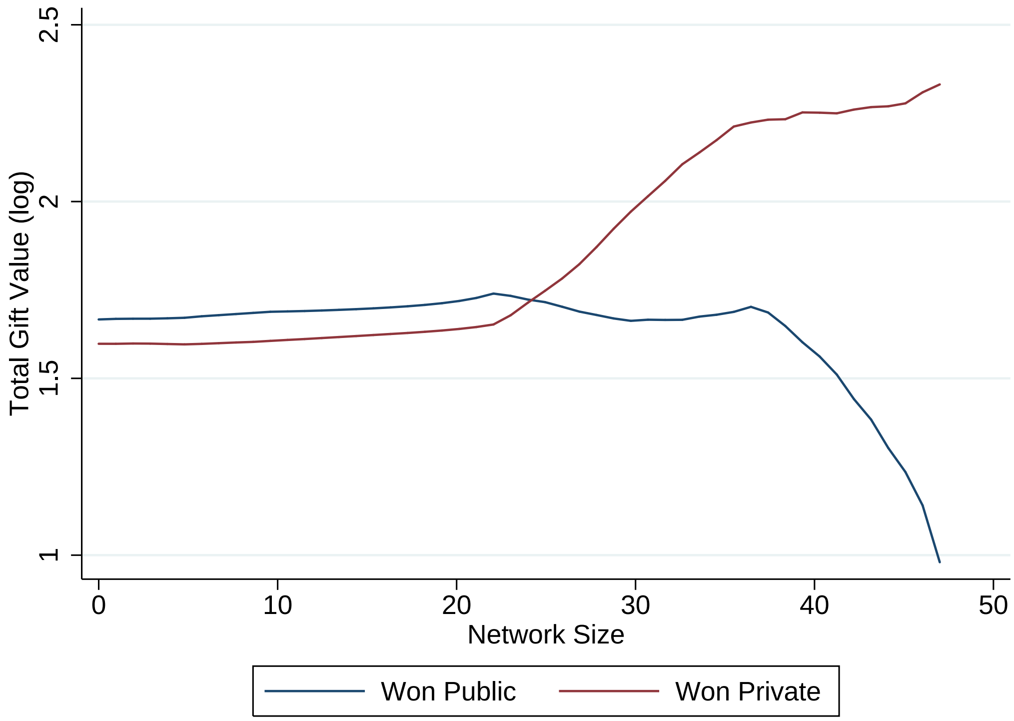

Note: Locally smoothed polynomial of total gift-giving as a function of network size. Variable on the Y-axis is total value of gifts given minus household average across rounds. Red-line only includes sample of individuals who won privately revealed lotteries. Blue line only includes sample of individuals who won publicly revealed lotteries. At small network-size values, total gifts given by public lottery winners is slightly higher than private lottery winners. This relationship remains flat and is inverted at a network size of roughly 20 at which point total gifts given by publicly revealed lottery winners starts to decrease — indicating behavior consistent with the shut-down hypothesis.

FIGURE 8: Effects of Network Lottery Winnings on Food Consumption

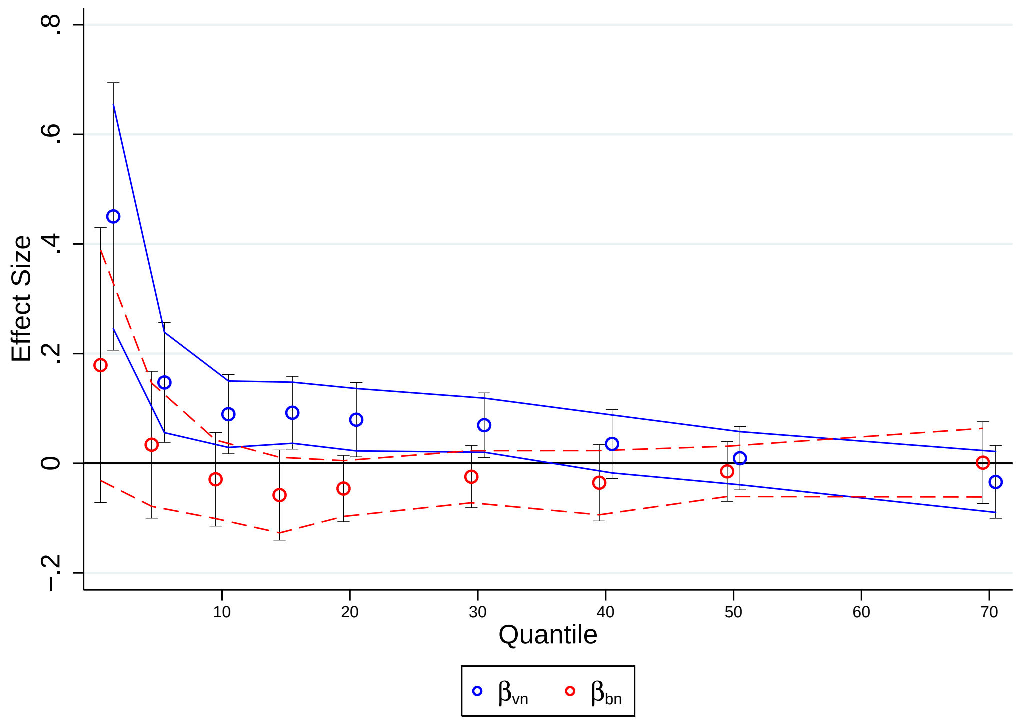

Note: Figure represents coefficient estimates of quantile regression of equation 15 (HH fixed effects replaced by village fixed effects due to dimensionality constraints). Dependent variable equals log per-capita food consumption. Simultaneous quantile regression estimator with 100 bootstrap repetitions. Only first three rounds of data used — these periods coincide with Southern Ghana's hungry season. Blue represents average network treatment effect of privately revealed lottery winnings and red represents publicly revealed lottery winnings. Evaluated at the 1%, 5%, 10%, 15%, 20%, 30%, 40%, 50%, and 70%-tiles. Per-capita food consumption more likely to increase for lowest quantiles following solidarity network's private lottery winnings relative to public lottery winnings.

Appendix Materials

A Adding a Sequence of History-dependent Nash Equilibria (SHDNE) Transfers to Our Model

Households can default to an SHDNE (instead of a no-transfer equilibria) and transfer amounts in such settings will depend on the level of altruism between household 1 and 2 and the number of household 1's outstanding gift-commitments. The SHDNE transfer, $\tau^D(h_t)$, given history h_t is

$$\tau^D(h_t) = \begin{cases} r \text{ s.t. } u_1'(y_1(s_t) - r)/u_2'(y_2(s_t) + r) = \gamma_1(g_1(h_t)) \\ \qquad\qquad \text{if } u_1'(y_1(s_t))/u_2'(y_2(s_t)) < \gamma_1(g_1(h_t)) \\ r \text{ s.t. } u_1'(y_1(s_t) - r)/u_2'(y_2(s_t) + r) = 1/\gamma_2(g_1(h_t)) \\ \qquad\qquad \text{if } u_1'(y_1(s_t))/u_2'(y_2(s_t)) > 1/\gamma_2(g_1(h_t)) \\ 0 \text{ otherwise.} \end{cases} \tag{18}$$

In other words, 1 will transfer to 2 when 2's marginal utility of consumption at his state-specific income level is high enough relative to individual 1's history-dependent gift-network size. Similarly 2's transfers to 1 will depend on 2's history-dependent gift-network size. In either case, the SHDNE transfer is voluntary and not contingent on any requirement for the recipient party to reciprocate in a future period.

To set up the household's problem with default to SHDNE transfers after history h_t, $U_1(h_t)$ can be re-written in the following manner:

$$\begin{aligned} U_1(h_t) = \quad & u_1(y_1(s_t) - \tau(h_t)) - u_1(y_1(s_t) - \tau^D(h_t)) \\ & + \gamma_1(g_1(h_t))u_2(y_2(s_t) + \tau(h_t)) - \gamma_1(g_1(h_t))u_2(y_2(s_t) + \tau^D(h_t)) \\ & + \mathbb{E} \sum_{k=t+1}^{\infty} \delta^{k-t} \left\{ \begin{array}{l} u_1(y_1(s_k) - \tau(h_k)) - u_1(y_1(s_k) - \tau^D(h_t)) \\ +\gamma_1(g_1(h_t))u_2(y_2(s_k) - \tau(h_k)) - \gamma_1(g(h_t))u_2(y_2(s_k) - \tau^D(h_t)) \end{array} \right\} \\ & - \alpha_1(g_1^D(h_t)) \end{aligned} \tag{19}$$

where instead of only receiving income $y_1(s_t)$ in each period after h_t, household 1 will subtract net SHDNE transfers as well. The rest of the maximization problem is straightforward to compute once a functional form for utility is identified.

B Appendix Tables

Table B.1: Individual Summary Statistics

	N	Mean	Sd	5 p-tile	95 p-tile
Fixed Over Time:					
HH size	606	5.09	2.23	2	9
Gift Network Size	597	9.94	10.10	0	31
Gifts and Loans (last 2 months):					
N Gifts Given	2,983	0.82	1.37	0	4
N Gifts Received	2,983	0.30	0.80	0	2
N Loans Given	2,983	0.16	0.51	0	1
N Loans Taken	2,983	0.07	0.29	0	1
Total Value of all Gifts Given	1,175	20.02	75.25	1	66
Total Value of all Gifts Received	542	12.58	35.75	1	35
Total Value of all Loans Given	362	57.00	113.25	5	200
Total Value of all Loans Taken	191	54.90	133.46	3	220
Food Consumption (last month):	**HH Head**	**Spouse**	**P-value**	**Total**	**SD**
PC Food Consumption	10.43	16.71	0	26.45	20.77
PC Purchased Food	3.11	15.86	0	19.42	18.83
PC Home-produced Food	7.78	1.60	0	8.63	7.98

Gift Network data missing for a subset of observations. N of loans/gifts given equal zero if none given/received. Value of gifts/loans contingent on having received at least one. Gift/loan data excludes within-household transfers and "Gifts Receives" and "Loans Taken" exclude all gifts or loans that originate outside of the study village. Household food consumption (total) sums the head of households and spouse's response. P-value is t-test significance of difference in category of food spending between HH head and spouse.

Table B.2: Lottery Winnings

Variables:	N	Mean	Sd
Own Lottery Winnings:			
Cash - Private	1,288	1.21	5.70
Cash - Public	1,288	1.12	5.52
Gift-Giving Network Average Lottery Winnings:			
Friends Cash-Private	1,184	1.17	1.64
Friends Cash-Public	1,184	1.16	1.45

Friend lottery winnings multiply the vector of lottery winners by the row-normalized gift network adjacency matrix (result is average friends' lottery winnings).

Table B.3: Household Summary Statistics

	N Winners		Win-at-all		Win-Private		Win-Public	
	N-No	N-Win	Diff	P-Value	Diff	P-Value	Diff	P-Value
Fixed Over Time:								
HH size	190	119	0.31	0.22	0.56	0.06*	-0.04	0.91
N Mutual Gifts	190	119	-0.52	0.66	-0.78	0.57	0.24	0.86
Gifts and Loans (last 2 months):								
N Gifts Given	190	119	0.14	0.64	0.07	0.85	0.34	0.32
N Gifts Received	190	119	0.11	0.46	0.18	0.28	0.11	0.52
N Loans Given	190	119	0.08	0.45	-0.09	0.49	0.12	0.36
N Loans Taken	190	119	-0.05	0.32	-0.06	0.26	-0.00	0.95
Food Consumption (last month):								
PC Food Consumption	187	117	-0.75	0.79	-4.40	0.18	2.49	0.44
PC Purchased Food	187	117	-0.02	0.99	-1.97	0.49	2.04	0.47
PC Home-produced Food	187	117	-0.73	0.49	-2.43	0.05**	0.45	0.71

Balance test of round 1 observations. categorizes households according to those who won either lottery at any point during the course of the year (N-No, number of HH that did not win; N-Win, number of HH that did win). represents either lottery. represents those who only won private (public) lotteries. All P-Values represent two-tailed hypothesis tests (t-statistics).

Table B.4: Gift Network = MMG

	No Interaction			Shut-down Hypothesis		
	(1)	(2)	(3)	(4)	(5)	(6)
N Mutual Gifts = 0	-0.488**	0.118	-0.396**	-0.491**	0.102	-0.398**
	(0.215)	(0.197)	(0.163)	(0.215)	(0.189)	(0.161)
N Mutual Gifts	0.050***	0.021***	0.020***	0.055***	0.023***	0.022***
	(0.008)	(0.006)	(0.006)	(0.008)	(0.006)	(0.006)
Private Cash Winnings	0.085	0.038	0.062	0.047	-0.003	0.040
	(0.086)	(0.035)	(0.040)	(0.127)	(0.047)	(0.063)
Public Cash Winnings	0.179**	0.072	0.052	0.499***	0.200***	0.205**
	(0.086)	(0.046)	(0.051)	(0.127)	(0.066)	(0.080)
N Mutual Gifts × Private Cash Winnings				0.003	0.003	0.002
				(0.008)	(0.003)	(0.004)
N Mutual Gifts × Public Cash Winnings				-0.031***	-0.013***	-0.015***
				(0.009)	(0.005)	(0.005)
Village FE	Yes	Yes	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1645	1602	1645	1645	1602	1645

*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Dependent Variable equals log of gifts given in columns 1 and 4, number of gifts given in column 2 and 5, and value per gift in columns 3 and 6. Village and Round Fixed Effects Included in Every Specification. Tobit regression in columns 1, 3, 4 and 6 with a lower bound of zero (no upper bound); poisson regression in columns 2 and 5.

Table B.5: Dyadic Regressions - Pooled Base Test

	Pooled - HH FE		Pooled - TGT HH FE	
	Log(Amount$_{ijt}$)	Number$_{ijt}$	Log(Amount$_{ijt}$)	Number$_{ijt}$
main				
Network Size$_i$			0.030	0.014
			(0.028)	(0.015)
Mutual Gift$_{ij}$			4.488***	2.380***
			(0.601)	(0.306)
Network Size$_i$ × Mutual Gift$_{ij}$			-0.047	-0.020
			(0.036)	(0.020)
Lottery-Private$_{it}$	0.486	0.255	0.211	0.065
	(0.434)	(0.238)	(0.284)	(0.156)
Lottery-Public$_{it}$	0.134	0.094	0.439**	0.297**
	(0.184)	(0.089)	(0.200)	(0.123)
Lottery-Private$_{it}$	-0.151**	-0.087**	-0.099***	-0.050**
× Network Size$_i$	(0.060)	(0.034)	(0.037)	(0.020)
Lottery-Public$_{it}$	-0.027	-0.019*	-0.033*	-0.026*
× Network Size$_i$	(0.018)	(0.010)	(0.019)	(0.014)
Lottery-Private$_{it}$	1.158**	0.462	0.525	0.306
× Mutual Gift$_{ij}$	(0.560)	(0.297)	(0.434)	(0.201)
Lottery-Public$_{it}$	1.279***	0.425*	0.093	-0.137
× Mutual Gift$_{ij}$	(0.420)	(0.244)	(0.374)	(0.194)
Lottery-Private$_{it}$	0.106*	0.072**	0.077*	0.045**
× Mutual Gift$_{ij}$ × Network Size$_i$	(0.062)	(0.035)	(0.039)	(0.020)
Lottery-Public$_{it}$	-0.004	0.008	0.011	0.016
× Mutual Gift$_{ij}$ × Network Size$_i$	(0.028)	(0.016)	(0.028)	(0.017)
TGT HH FE	No	No	Yes	Yes
HH FE	Yes	Yes	No	No
Village FE	No	No	No	No
Round FE	Yes	Yes	Yes	Yes
N	133975	130761	130761	130761

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable in odd columns equals amount of actual gift given from i to j in any given round t; Even columns equals number of gifts. 'Network Size$_i$' indicates household i's gift-network size (any type of gift-relation). Odd columns are tobits with lower bound of zero. Even columns are poisson regressions. Columns 1-2 only include links (i and j) with mutual gift-relations at baseline (actual gifts given during the 5 rounds). Columns 3-4 include all oher links (actual gifts given).

Table B.6: Dyadic Regressions - Food Shocks

	Log(Amount$_{ijt}$)	Log(Amount$_{ijt}$)	Number$_{ijt}$	Number$_{ijt}$
Mutual Gift$_{ij}$	3.840***	3.603***	2.085***	1.900***
	(0.332)	(0.322)	(0.197)	(0.180)
Lottery-Private$_{it}$	-0.449*	-0.473*	-0.278*	-0.287*
	(0.248)	(0.260)	(0.145)	(0.148)
Lottery-Public$_{it}$	0.061	0.012	0.041	0.016
	(0.180)	(0.188)	(0.114)	(0.100)
Food-Shock$_{ijt}$	-0.063	-0.098	-0.055	-0.043
	(0.215)	(0.187)	(0.119)	(0.101)
Food-Shock$_{ijt}$	-0.155	-0.105	-0.054	-0.049
× Mutual Gift$_{ij}$	(0.313)	(0.262)	(0.161)	(0.125)
Gift-Network Interaction				
Lottery-Private$_{it}$	0.711**	0.675**	0.559***	0.486***
× Mutual Gift$_{ij}$	(0.335)	(0.318)	(0.170)	(0.161)
Lottery-Public$_{it}$	0.034	0.279	-0.071	0.045
× Mutual Gift$_{ij}$	(0.269)	(0.271)	(0.151)	(0.139)
Food-Shock Interaction				
Lottery-Private$_{it}$	-0.467*	-0.560*	-0.258**	-0.295*
× Food-Shock$_{ijt}$	(0.258)	(0.301)	(0.129)	(0.151)
Lottery-Public$_{it}$	0.494**	0.623**	0.247*	0.309**
× Food-Shock$_{ijt}$	(0.252)	(0.265)	(0.141)	(0.122)
Triple Interaction				
Lottery-Private$_{it}$	1.043***	1.014***	0.394***	0.337*
× Food-Shock$_{ijt}$ × Mutual Gift$_{ij}$	(0.312)	(0.381)	(0.138)	(0.190)
Lottery-Public$_{it}$	-0.792**	-0.894*	-0.337*	-0.414*
× Food-Shock$_{ijt}$ × Mutual Gift$_{ij}$	(0.350)	(0.476)	(0.177)	(0.236)
HHN FE	No	Yes	No	Yes
TGT HHN FE	Yes	No	Yes	No
Round FE	Yes	Yes	Yes	Yes
N	110721	110721	110721	110721

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable in columns 1-2 equals amount of actual gift given from i to j in any given round t; in columns 3-4 it equals number of gifts. 'Food-Shock$_{ijt}$' indicates the difference between i and j estimated food consumption residual ($\hat{\chi}_i - \hat{\chi}_j$ — household and round fixed effects). Columns 1-2 estimated using tobit estimator with lower bound on dependent variable of zero. Columns 3-4 use poisson estimator. 'Mutual Gift$_{ij}$' refers to the existence of reciprocal gift-relationships between i and j.

C Appendix Figures

FIGURE C.1: Amount of Total Transfers by Network Size

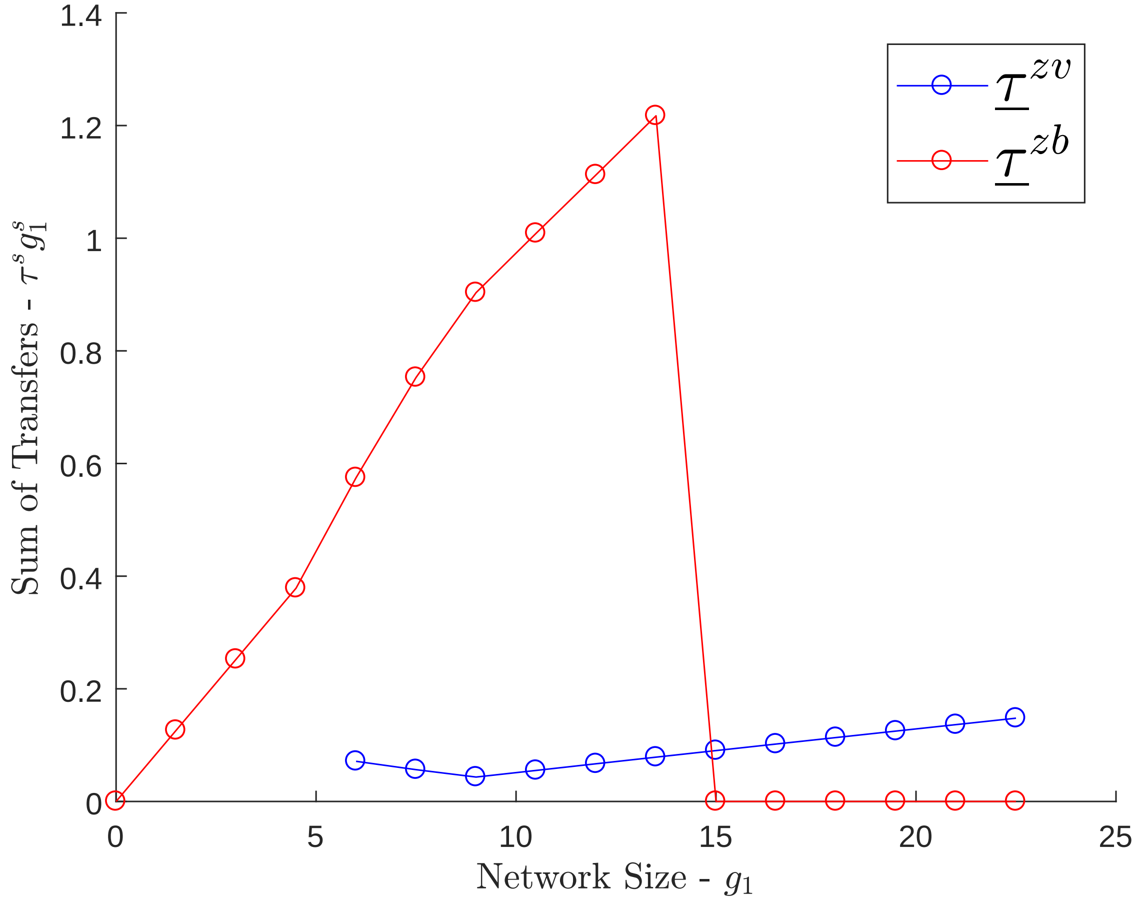

Note:

Figure C.2: Local Polynomial Smoothing - Shut-down Hypothesis

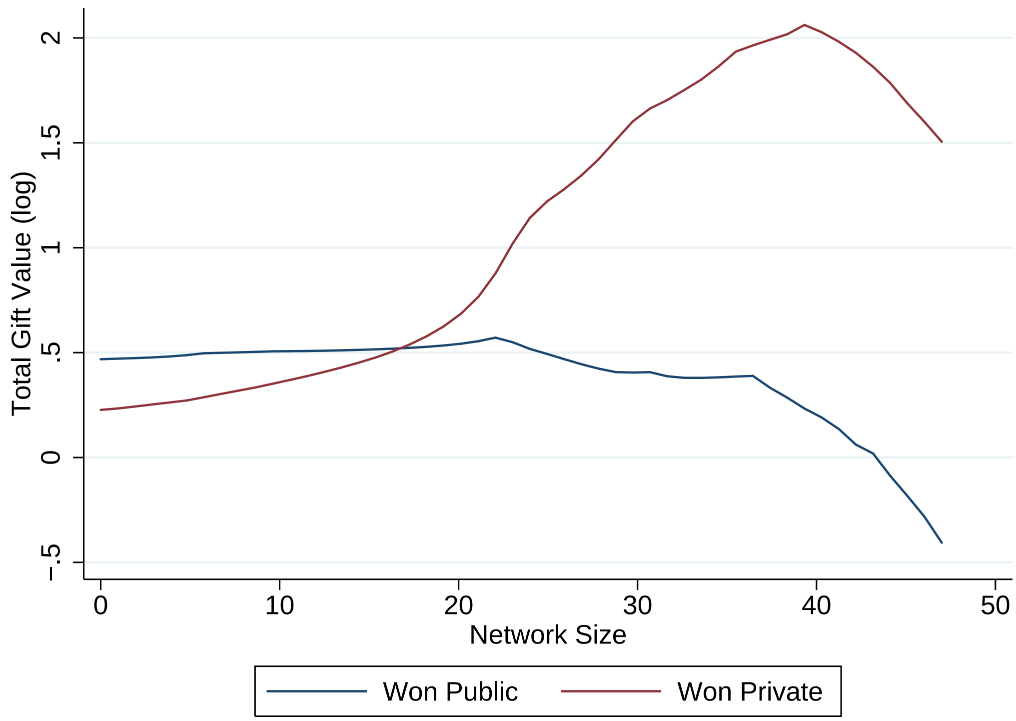

Note: Values on the y-axis adjusted for household fixed effects. See description in figure 7

Figure C.3: Results of Quantile Regression by Quantile - All Rounds

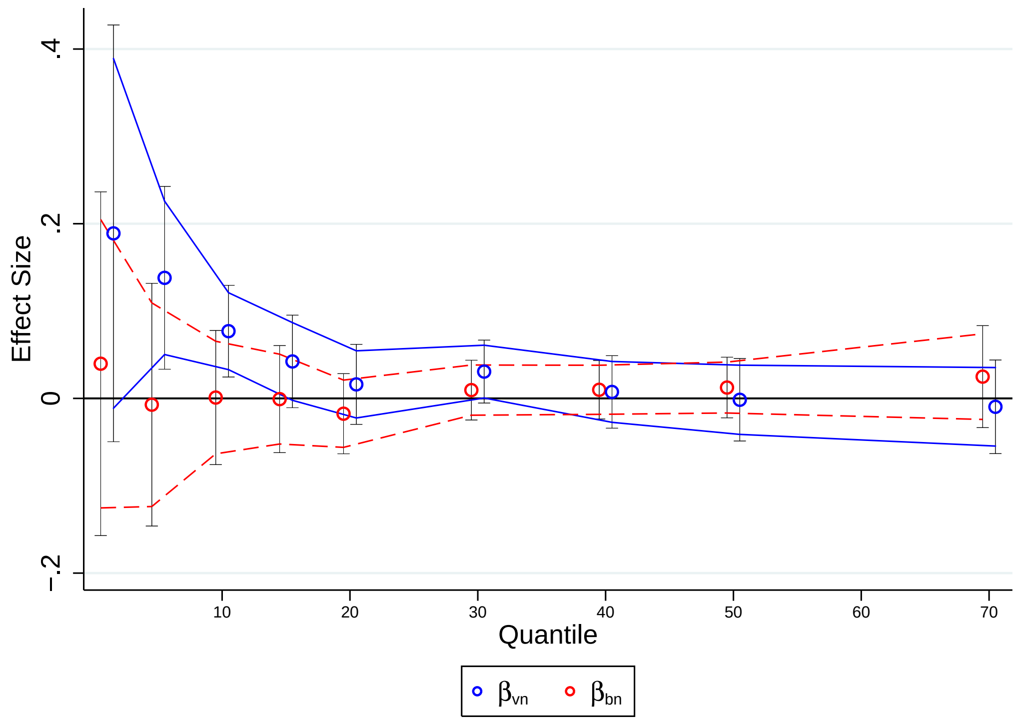

Note: Estimation of equation 15 using simultaneous quantile regression estimator (100 bootstrap repetitions). Dependent variable equals log per-capita food consumption. All rounds of data used. Blue represents average network treatment effect of privately revealed lottery winnings and red represents publicly revealed lottery winnings. Evaluated at the 1%, 5%, 10%, 15%, 20%, 30%, 40%, 50%, and 70%-tiles. Coefficient estimates offset relative to x-axis for ease of viewing.