Altruism, Insurance, And Costly Solidarity Commitments

Vesall Nourani* Christopher B. Barrett
Eleonora Patacchini Thomas Walker

May 22, 2020

Abstract

Whether understood as informal insurance, credit, or social taxation, the conceptual models of inter-household transfers rest on foundations of purely self-interested behavior. Using experimental data from rural Ghana we reject two core predictions of self-interested models. We then add altruism and social taxation to a standard model of limited commitment informal insurance contracting. The data support this model's more nuanced predictions, including that unobservable income shocks may facilitate altruistic giving targeted towards less-well-off individuals within one's network - especially friends, rather than family - and a large network can induce social taxation that overwhelms an altruistic agent, inducing her to cease giving.

*Nourani (corresponding author): Massachusetts Institute of Technology, vnourani@mit.edu. Barrett: Cornell University, cbb2@cornell.edu. Patacchini: Cornell University, ep454@cornell.edu. Walker: World Bank, twalker@worldbank.org. Harold Alderman, Taryn Dinkelman, Andrew Foster, Susie Godlonton, Joe Kaboski, Travis Lybbert, Annemie Maertens, Erin McDonnell, Daniel Putman, Ashish Shenoy, Liz Tennant, and seminar participants at Copenhagen, Cornell, Einaudi Institute for Economics and Finance, LSE, MIT,Notre Dame, UC-Davis, the U San Francisco, Williams, and the AAEA annual meetings provided excellent feedback on earlier versions of this paper. Special thanks to residents of the four villages used in this study for opening their homes to share their experiences, and to Robert Osei, Chris Udry and Jacqueline Vanderpuye-Orgle for crucial advice on the field data collection. Data used in this study were collected under USAID's AMA-CRSP Program (award number P686140). The corresponding author acknowledges support from the National Science Foundation under Grants No. DGE-1144153 and 1809427 as well as the African Development Bank Structural Transformation of African Agriculture and Rural Spaces project supported by the Korean government through the Korea-Africa Economic Cooperation Trust Fund.

Social solidarity networks have long been understood to play a central role in village economies. Although the possibility of altruism has been accommodated in some work within that literature, at least since Popkin (1979) and Posner (1980), the dominant framework for social scientists' understanding of transfers within social networks has rested on purely self-interested dynamic behavior, most commonly framed as self-enforcing informal insurance contracts among households whose risk averse preferences drive them to smooth consumption through transfers (Fafchamps, 1992; Coate and Ravallion, 1993; Townsend, 1994; Ambrus, Mobius and Szeidl, 2014).

A related but distinct literature emphasizes a darker side of self-interested behavior within social networks. Social pressures — often referred to as 'social taxation' — can place significant demands on those who enjoy observable income gains, thereby discouraging investment and potentially even trapping households in poverty (Platteau, 2000; Sen and Hoff, 2006; Jakiela and Ozier, 2016; Goldberg, 2017; Squires, 2017).

Informal insurance and social taxation both depend fundamentally on the public observability of income shocks.[1] The effectiveness of self-enforcing insurance among purely self-interested agents depends upon each party's ability to monitor others' income shocks so as to enforce the contract. Similarly, a social network can only tax the observable portion of one's income stream.

This implies two testable hypotheses. First, publicly observable (hereafter 'public') income shocks should lead to inter-household transfers, whether due to social taxation, informal insurance contracts, or some other self-interested mechanism. One should therefore reject the null hypothesis that public income shocks have no effect on inter-household transfers in favor of the one-sided alternate hypothesis of positive impacts. Second, income shocks that are not publicly observable (hereafter 'private') — in particular, positive private income shocks that a purely self-interested beneficiary would never divulge — should not prompt inter-household transfers. Of course, failure to reject the null that private income shocks have zero effect is a low-power test. We are

[1] More specifically, they rely on non-uniform shocks across households within the network so that exogenous income changes trigger redistribution through informal insurance, social taxation, or both. We use the term 'income shock' to imply non-uniform shocks.

unaware of any prior empirical tests of these two implications of the public observability of income shocks on inter-household transfers.

Neither prediction implied by the public observability hypothesis holds in a novel field experiment we conducted among households in southern Ghana. Over the course of a year we randomized private and public bimonthly cash payments to subjects whose informal gift networks we had previously mapped. In regressions of giving within subjects' social networks as a function of randomized private and public winnings, we fail to reject the public income shocks null but do reject the private income shocks null. We corroborate those findings with dyadic regressions on the flows between matched subjects.

These findings are inconsistent with framing inter-household transfers as solely a result of self-interested behaviors. Our results offer a more encouraging statement about human nature. People may act partly out of self-interest, but they also behave in ways that reflect more altruistic, other-regarding interests.

The experimental results motivate us to adapt the canonical stochastic dynamic model of self-enforcing insurance contracting. We introduce both costly altruistic giving, building on Foster and Rosenzweig (2001) and Ligon, Thomas and Worrall (2002), and social taxation pressures. Altruistic preferences directly explains why people might give from private income windfalls, and - as we show empirically - in a distributionally progressive manner, while still allowing for self-interested behavior. People can be both selfish and altruistic.

We incorporate both altruism and social taxation by introducing a costly, 'warm glow' component to altruistic preferences (Andreoni (1990)), in which the marginal gains from giving to others diminish faster than the marginal costs of giving. There then emerges a point at which even altruists cease giving. Everyone faces some outer limit to the pleasure they derive from beneficence or compliance with social taxation norms. We tease out the implications of this 'shutdown hypothesis' for inter-household transfer patterns. The key implication is that if social taxation demands increase with one's network size and with the public observability of income, then the larger one's network and the more observable one's income, the more likely one reaches the shutdown point. Over some domains, social taxation can counter altruistic and self-interested motives for interhousehold transfers, thereby explaining why people fail to give consistently from public income windfalls.

An important distributional implication is that when stochastic income realizations are publicly observable, the insurance or social taxation claims placed by less needy community members on another's windfall gains can crowd out the latter's altruistic giving to needier network members.In the presence of both altruism and social taxation, private rather than public income might better harness social networks so as to target the least well off in a population. This is precisely what we find in the experimental data.

These realistic model refinements obviate the sharp predictions of the public observability hypothesis implied by purely self-interested models. Interhousehold transfers now become non-monotone in response to public income shocks and may increase in private income shocks. These enhancements also suggest other testable hypotheses that match our experimental data remarkably well while still accommodating the core insights of the informal insurance and social taxation literatures. We find that the average inter-household transfer is larger following private than public windfall gains, and those transfers are more targeted to the neediest households within one's network, consistent with altruistic preferences. Over much of the range of network sizes, the number of gifts given is similar for public and private winnings. But the shutdown hypothesis holds; winners of publicly revealed cash prizes cease making transfers at all when they have too large a network. And after someone hits the shutdown point and stops giving to others after a public windfall, she becomes less likely to receive future transfers. Social taxation can thereby fray the informal insurance fabric networks create. We rule out a number of alternative explanations of these results in a series of robustness checks.

Our enhanced model and empirical findings reflect how our Ghanaian research subjects described to us the multifunctional nature of their social networks. Individuals value consumption smoothing and therefore leverage networks to accomplish that insurance goal. Limited risk pooling among family members holds following public income shocks when gift networks are of small-to-moderate size. For this special (but commonplace) case, the standard informal insurance model fits the data quite well. However, transfers to family members do not increase when windfall income is private, suggesting that private income is not easily observed by kin (De Weerdt, Genicot and Mesnard, 2019; Kinnan, 2019). Rather, private income increases transfers to the neediest within the village, mainly to friends, not family. This conforms with villagers' expression of (imperfectly) altruistic preferences. We thereby reconcile the informal insurance and social taxation literatures with each other, while allowing for the richer interhousehold transfer behaviors observed in our data. The social solidarity network is indeed multifunctional, (incompletely) pooling income risk across a network so as to (partially) smooth consumption as an insurance contract would, while also accommodating the social taxation pressures of kin and community members, and at the same time mediating altruistic transfers towards the least fortunate members of the network.

Our findings have practical policy implications, perhaps especially for cash transfer programs which have become the foundation for many social protection programs in the developing world. The purely self-interested informal insurance framework implies that social networks should (at least partially) correct targeting errors in publicly observable transfer programs, as non-recipients who suffer adverse shocks will enforce their claims on members of their network who enjoy windfall external transfers (Angelucci and De Giorgi, 2009). But if people have large networks, are at least partly altruistic and know better their network members' needs than external institutions do (Alderman, 2002; Bowles and Gintis, 2002), and are subject to social taxation pressures, then transparent transfer payments may diminish the progressivity of redistribution within networks by replacing altruistic motives with social taxation pressures.

1 Data and Descriptive Evidence

Our data combine a field experiment with repeated household surveys. The experiments were conducted bimonthly from March-October 2009 in conjunction with a year-long household survey in four communities in Akwapim South district of Ghana's Eastern Region, roughly 40 miles north of the nation's capital, Accra, but sufficiently far away that only a handful of respondents commute to the greater Accra metropolitan area for work. The sample consists of approximately 70 households from each of the four communities.[2] Individuals in the sample include the household head and his spouse.[3] There 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.

Experimental Data. Prior to survey rounds two through five we randomized cash and in-kind lotteries among the sample households so as to manufacture positive income shocks. The first (January) round of the survey was designed as a baseline, therefore no lottery took place in that round. One week before each subsequent survey round began we visited each village to remind villagers that we would soon return and to arrange visit dates with respondents. While in each village, we distributed 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 across the four lottery rounds and villages. Within each village and round, ten of the prizes were cash; the other ten were in the form of livestock. Approximately 58 percent of households won at least one prize of any type and 39 percent won cash prizes over the course of the year. For both cash and livestock winnings, five each were allocated publicly by lottery, and the other five (identical in type and value) were allocated in private, by lucky dip.[4] The values of the prizes varied from GH¢10 to GH¢70.[5] The prizes were of a substantial size - the largest prize is equivalent

[2]The first survey waves using this sample were conducted in 1997-98 (Conley and Udry (2010)) and 2004 (Vanderpuye-Orgle and Barrett, 2009). Most of the 70 households in each village were in the 1997-98 sample. 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 in 2009.

[3]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.

[4]Over the course of the year 23 percent of households won a private cash prize and 23 percent won a public cash prize, with little overlap. 4 households won a public prize twice (one of whom also won a private prize); 9 households won a private prize twice (2 of whom also won a public prize) and 1 household won private and public prizes 3 times each.

[5]During the course of our study, one GH¢ was roughly equivalent to 0.7 USD and prize

to roughly four times monthly per capita food consumption. In aggregate, each community's survey participants received GH¢370 of cash in each round to use however they would like. A detailed description of the events that took place during each lottery day are available in appendix section A.1.

Although private lottery winnings were not directly observable, villagers were aware who in the community was in the survey and might have made probabilistic inferences that a respondent may have won. But to the extent that private winnings were inferrable, if not observable, that would bias our empirical results towards finding no difference in impacts between private and public winnings. So this possibility only strengthens our findings.

The survey interviews in each round began one week after the lottery, deliberately delayed to allow winners to do something with their prize prior to the survey round. The interviews took place throughout the following three weeks, so that some winners were interviewed a week after receiving their prize, and others up to four weeks afterward.

Survey. Each respondent was interviewed five times during 2009, once every two months from February-November, with the two survey teams each alternating between two villages.[6] The survey covered a range of subjects: personal income, farming and non-farm business activities, inter-household gifts, transfers and loans, and household expenditures. In each round, both the husband and wife heading each household were interviewed separately on all of these topics.[7] The data were assembled mainly using information

values were equivalent to GH¢10, 20, 35, 50, and 70. In this paper, we are primarily interested in transfers of divisible windfall gains of constant known value among households within a round and where private winnings would be unobservable. 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 livestock differed by species and size: Chickens (GH¢10), two chickens (GH¢20), small goat (GH¢35), medium goat (GH¢50), and large goat (GH¢70). Different households may face different transaction costs, so the value of livestock, as opposed to cash, is heterogeneous across households, which would further complicate the use of livestock in the analysis, although we do control for livestock winnings in the empirical tests below. Most importantly for our purposes, it was impossible to 'privately' grant lottery winners a chicken or goat, while it was easy to award them the same amount in cash in ways others could not observe.

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

[7]There were some households with multiple spouses and others without a spouse. For

contained in the expenditure, gift and social network modules of the survey.

Inter-household Transfers. 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 identity, location and relationship to the respondent.[8] The value of the gift given and an estimated value for in-kind gifts were also recorded. In our preferred analysis we focus exclusively on cash gifts given since we are primarily interested in transfers of divisible windfall gains of constant known value among households within a round. Similarly, because livestock prizes won 'privately' were in practice almost always observable by other villagers, we cannot distinguish private from public livestock winnings in our analysis. With this in mind, we show that our results are not sensitive to the inclusion of either in-kind gifts given (as an outcome variable) or in-kind cash prizes (as an explanatory variable) into our main discuss analysis of our results in section 2.

We also focus on inter-household gifts to other households within the village and therefore drop gifts given or received to parties who reside outside of the village as well as incidents of within-household transfers — i.e., gifts transferred to one's spouse (which are studied in detail using the same data in Castilla and Walker (2013)) or co-residing children. In this context, the concept of gifts encompasses what one might think of as indemnity payments from an informal insurance contract, i.e., any inter-household transfer without an unconditional obligation to repay (i.e., not an explicit loan).[9]

Summary Statistics. Table D.2 reports the household aggregate measures

simplicity, throughout the paper we describe households as having a household head and spouse.

[8]When a gift is given to or received from another survey respondent, enumerators also provided the unique sample identifier for that individual. This enables dyadic analysis, which we explore in section 4.

[9]Table D.1 describes all reported gifts given and/or received across the five rounds of data. In total, 2,791 gifts were given to others from members of our sample (roughly 1.8 per two month period per household) and 3,006 gifts were received during the same period. Of these, we focus on gifts given to co-villagers and drop the 1,082 gifts given to others who reside outside of the village as well as any gifts to direct family members co-residing together. We remain with 1,542 gifts given to fellow villagers (a majority of all gifts given) of which 652 gifts (42%) are monetary gifts of commonly known value.

that form the basis of our analysis.[10] On average, each household has more than six members. Across the five rounds of data, households give and receive 0.41 and 0.25 cash gifts, respectively, to other households in the village over the course of two months.[11] Conditional on giving a gift, the average total value of the gifts given and received is 13.12 GH¢ and 11.77 GH¢, respectively. Note that the number and value of gifts given is larger than the number and value of gifts received. This would be the case if members of our sample increased participation in gift-giving, perhaps due to the influence of the experimental lottery, relative to those outside of the sample. The average value of winning either a publicly revealed or private cash prize is roughly 2.3 GH¢ in each of the four rounds in which we distributed cash prizes. We demonstrate that we are well-balanced along observable variables in appendix section A.2.

2 Testing the Public Observability Hypothesis

One typically cannot separate the private and public components of observed income streams without imposing rather herculean, untestable assumptions. Therefore, to date it has been infeasible to test the paired core predictions that inter-household giving increases in publicly observable income shocks and is invariant with respect to private income shocks unobservable to other households. Our experimental design, however, allows us to directly test this public observability hypothesis. Rejection of that hypothesis implies a need to enhance the core theory used to explain inter-household transfer behaviors.

Let y_{it} be the outcome of interest: the number of round t gifts distributed by household i in village v, the average amount per gift given, or the total amount given, i.e., the product of the first two outcomes. The two core hypotheses can be tested using the regression:

$$y_{it} = \alpha + \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} + \text{hh}_i + \text{r}_{tv} + \epsilon_{it}, \tag{1}$$

where β_v captures the extent to which round t gift-giving behavior is influenced by round t privately revealed lottery winnings and β_b captures the impact

[10] Disaggregated, individual-level, measures are available in Table D.3.

[11] Respondents may fail to report (perhaps especially small) gifts. But any such measurement error is almost surely uncorrelated with randomized lottery winnings, and thus would only lead to attenuation bias, further strengthening our findings.

TABLE 1: Prize Winnings and Gift Giving

		Gifts-Given		
Dependent Variable:		Value (Total) (1)	Value (Average) (2)	Number (3)
Randomized Explanatory Variables:				
Value of Private Cash Prize	β_v	0.222***	0.175***	0.238***
		(0.078)	(0.063)	(0.065)
Value of Public Cash Prize	β_b	0.109	0.0500	0.124*
		(0.087)	(0.070)	(0.072)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
Left-censored Obs.		1,182	1,182	1,182
Observations		1,586	1,586	1,586

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. The dependent variable equals log total value of cash gifts given in hh in column 1; log average value of cash gift given in column 2; number of gifts given in column 3. Value of Private/Public Cash prize is divided by $10 \in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns with a lower bound of zero. Table D.5 reports estimates of the number of gifts given using a Poisson estimator with qualitatively similar results as those in column 3.

of publicly revealed lottery winnings,[12] hh_i captures household fixed effects, r_{tv} captures village-specific round fixed effects that could affect giving by all households in a given village and period, and ϵ_{it} is the household-specific round t error term. For each specification we use the Tobit estimator where we integrate out censored observations equal to zero.

Table 1 reports the equation 1 estimation results. All three of the coefficient estimates for β_v are statistically significant and positive at the one percent level. Most notably, only one β_b coefficient estimate (number of gifts given) is statistically significant at the ten percent level. Moreover, all β_b point estimates are all smaller in magnitude than the β_v estimates. The p-values are only 0.09-0.17 on the one-sided test for differences between those two parameters. But recall that any partial observability or inferrability of private winnings will generate attenuation bias in this test.[13] Clearly, interhouse-

[12] We use 'b' to reflect public and 'v' to reflect private winnings throughout.

[13] Four robustness checks confirm Table 1's results. First, we use a Poisson count data

hold financial flows do not respond significantly to publicly observable income shocks but they do to private income shocks. This is the exact opposite of the public observability hypothesis that follows directly from purely self-interested models of inter-household transfers.

One candidate explanation for why respondents give from unobservable winnings is that households eventually, indirectly reveal unobservable income streams through observable consumption. Of course, that argument should hold equally for observable winnings, but does not. We nonetheless check for the possibility of incompletely hidden (i.e., semi-private) income by testing whether lagged income shocks - private or public - predict current consumption, following Kinnan (2019). That might help explain the otherwise-puzzling positive and significant β_v estimates in Table 1 . Toward that end, we estimate:

$$c_{it} = \alpha + \beta_v \text{Priv}_{it} + \beta_b \text{Pub}_{it} + \beta_{v,l} \text{Priv}_{it-1} + \beta_{b,l} \text{Pub}_{it-1} + \delta c_{it-1} + \text{r}_{tv} + \epsilon_{it} \quad (2)$$

which differs from equation 1 in that period t consumption, c_{it}, is the new dependent variable, we omit household fixed effects, and include lagged terms for the exogenous income shocks and consumption. Rejecting the null that $\beta_{v,l} = 0$ in favor of the one-sided alternate hypothesis that $\beta_{v,l} > 0$ would support the hypothesis that private winnings are not really hidden income and partly reconcile our results with the standard predictions of a purely self-interested insurance contracting model.

Appendix table D.8 reports results both using the "take one out" average of lagged consumption (column 1) and Round$\times$Village fixed effects (column 2). In all cases, the $\beta_{v,l}$ estimate is very small, slightly negative, and statistically insignificant. The data clearly do not support the hidden income hypothesis

estimator for the integer-valued dependent variable number of gifts given (Appendix Table D.5). Second, we include in-kind (livestock) lottery winnings in an alternate specification (Appendix Table D.6). In both cases, the results remain qualitatively unchanged. Third, when our measure of gifts given includes in-kind gifts, our results also remain qualitatively the same; the point estimates for β_v remain larger than β_b though they lose some precision (Appendix Table D.7). However, the value of large in-kind gifts are more difficult for households to estimate than small in-kind gifts, so it is likely that the value of gifts are less precisely measured. Indeed, household responses to number of gifts given remain statistically significantly positive following private (but not public) income shocks. When we remove outlying observations of gifts given (according to estimated value), our results are maintained with a higher level of confidence than prior to the removal of outliers.

as an explanation for rejection of the public observability hypothesis.

These findings motivate us to refine the canonical model of dynamic household choice over interhousehold transfers, incorporating a few small features informed by our discussions with and observations of our Ghanaian subjects. By incorporating these refinements into a fairly standard model of a dynamic endogenously enforceable contracting game between two agents facing stochastic income streams, we generate more nuanced predictions that reconcile fully with our data. The data fit the standard informal insurance model over an important range of the data; but over the full data set it becomes clear that social networks simultaneously mediate altruistic preferences and social taxation pressures in addition to informal insurance motives.

3 The Enhanced Model

We now show that a few empirically-grounded changes to a canonical model of risk-pooling can alter its predictions in important ways.[14] We build on Foster and Rosenzweig (2001) - and less directly on Ligon, Thomas and Worrall (2002) - who model transfers in a dynamic game in which two agents hold altruistic preferences over each other's consumption, in addition to standard preferences for their own consumption, and commitment to a transfer contract is imperfect due to lack of exogenous enforcement mechanisms. We 1) differentiate between private and public income, with the former not subject to informal contracting, 2) allow "warm glow" altruistic preferences that generate diminishing marginal utility in transfers made to others, 3) impose a cost associated with giving transfers, and 4) make the gift requests one receives an increasing function of both one's network size and the publicly revealed share of one's income, reflecting social taxation pressures that are endogenous to income observability.[15]

[14] Because we aim to reconcile informal insurance arrangements with social taxation and altruism, and thus with the data from our experiment, we need a stochastic dynamic model, so that insurance is relevant, and a game among at least two players, so as to allow for the agency inherent to informal contracting.

[15] Per the principle of parsimony, we rely on the relatively simple two household (as opposed to network) framework to illustrate the core empirical predictions, while keeping the state-contingent computations tractable.

Environment. We introduce two agents, $i = \{1, 2\}$ who receive stochastic incomes, $y_i(s_t) \geq 0$ that depend on the state, s_t, within the set of all states ($s \in \{1, 2, ..., S\}$), realized in period t. A sequence of the state history is characterized by $h_t = \{s_1, s_2, ..., s_t\}$.[16] We decompose income into its private (v) and public components (b) such that $y_i(s_t) = y_v(s_t) + y_b(s_t)$. We model the choice of history-dependent transfers from household 1 to household 2, $\tau(h_t)$, in period t; a negative $\tau(h_t)$ thus implies a flow from 2 to 1. In such an environment, household 1 knows both its private and public income but only its public income is revealed to household 2 (and vice versa). Both households have links with $g_1 - 1 = g_2 - 1 = g - 1$ other households where $g \geq 1$.[17]

We now introduce social taxation into the model. Depending on the realization of a particular state, households receive $gp_i(s_t)$ gift requests from their network in addition to the transfer arising from the informal contract.[18] We allow $0 \leq p_i(s_t) \leq 1$ to reflect the unconditional probability that a given household in one's network will request a transfer in period t. Consistent with evidence on social taxation, $p_i(s_t)$ is larger when the income realization is publicly revealed to i's network (Jakiela and Ozier (2016)).

Households will consider (self-interested) insurance obligations in an environment with limited commitment. To facilitate analysis, we impose two assumptions at this point:

Assumption 1 (Enforcement requires public observability of income)**.** *Insurance obligations can only be enforced when income is publicly observed. In other words, contracts only consider public income, $y_b(s_t)$.*

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

[17] We limit analysis to settings in which households 1 and 2 have the same network size for two reasons. First, we focus on the transfer decisions of household 1 as its network size increases. To understand the model's full network dynamics, one must model transfers across the entire network (a topic for a future paper). Second, setting $g_1 = g_2$ imposes symmetry in the transfer rule, simplifying the analysis. Otherwise (i.e., when $g_1 \neq g_2$), one must first identify conditions under which either household is a net sender or recipient of transfers and then adjust the denominator to the state-contingent equilibrium transfer level.

[18] These may include transfers one allocates to one's insurance network. We cannot differentiate insurance-relations from social taxation relations in our data, so we do not model the difference explicitly here.

Assumption 2 (Equal transfers to all)**.** *Households send transfers of equivalent size* $\tau(h_t)$ *to all households to which it transfers in a given period.*[19]

To simplify notation, let $r_1(s_t) = gp_1(s_t) + 1$ characterize the degree of social (taxation) pressure faced by household 1. As a result, net income for household 1 is $y_1(s_t) - \tau(h_t)r_1(s_t)$ and (observed) net income for household 2 is $y_{b2}(s_t) + \tau(h_t)$ since household 1 does not observe private income or any other transfers to household 2. Iff $\tau(h_t)$ >0, then household 1 is a net sender of transfers within the dyad.

Although we are interested in understanding how transfers change as a function of network size, we do not model network size as a choice variable. We acknowledge that there are implications for endogenous network choice that emerge from our enhanced model. We address the potential for endogenous networks below, empirically.

Preferences. Following Foster and Rosenzweig (2001), we assume households hold altruistic preferences towards others' single-period utilities. Individual i's altruistic preferences appear as household single-period utility that is separable in own and other household net income (assuming away savings and credit). Single-period utility for household 1 is:

$$u_1(y_1(s_t) - r_1(s_t)\tau(h_t)) + \gamma_1(r_1(s_t))u_2(y_2(s_t) + \tau(h_t)) \tag{3}$$

$u_1(\cdot)$ and $u_2(\cdot)$ are increasing, twice-differentiable, and concave functions and $\gamma_1(r_1(s_t))$ represents the altruism weight household 1 holds towards 2 in state s_t. Single-period utility for household 2 can be written in symmetric fashion.

Altruistic weights decrease in the social taxation pressure one faces. Intuitively, this indicates that the "warm glow" from giving is not as warm when one has to give to many others (Andreoni, 1990). It could also reflect the idea that altruism is not as effective an instrinsic motivator when the transfer feels forced. We do not distinguish between these motives. In practice, altruism weights diminish as a household's period-specific gift requests, $r_i(s_t)$, increase.

[19] This simplifying assumption is motivated by tractability concerns, observed inequality aversion among low-income rural households (e.g., Pitt, Rosenzweig and Hassan (1990), Berry, Dizon-Ross and Jagnani (2020)), and the clear belief in these communities that unequal treatment can inflame costly social rivalry (Cole, Mailath and Postlewaite (1992)).

Specifically (for household 1), let altruistic weights be a differentiable mapping of non-negative inputs $r_i(s_t) \in R^+$ onto $[0, 1)$ such that:

$$\gamma_i(0, s_t) = \overline{\gamma} < 1, \text{ and } \frac{\partial \gamma}{\partial r_i(s_t)} < 0 \text{ for all } r(s_t) > 0 \tag{4}$$

where $\overline{\gamma} < 1$ places an upper bound on household 1's altruism weight towards household 2, indicating that agents always value their own consumption more than others'.

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. We assume that if either party does not uphold the terms of the contract, they revert to an autarkic no transfer rule thereafter. Thus, we introduce the following assumption

Assumption 3 (Reneging on a contract leads to autarkic future consumption.). *A household i with positive public income, s' such that $y_{bi}(s') > 0$, drops out of future transfer arrangements if $\tau(s') = 0$. Furthermore, for every h_t such that $s' \in h_t$, $U_i(h_t) = 0$ and i engages in autarkic consumption.*[20]

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:

$$\begin{aligned} U_1(h_t) = \quad & u_1(y_1(s_t) - \tau(h_t) r_1(s_t)) - u_1(y_1(s_t)) + \\ & \gamma_1(r_1(s_t)) u_2(y_{b2}(s_t) + \tau(h_t)) - \gamma_1(r_1(s_t)) u_2(y_{b2}(s_t)) + \\ & \mathbb{E} \sum_{k=t+1}^{\infty} \delta^{k-t} \left\{ \begin{array}{l} u_1(y_1(s_k) - \tau(h_k) r_1(s_k)) - u_1(y_1(s_k)) + \\ \gamma_1(r_1(s_k)) u_2(y_{b2}(s_k) + \tau(h_k)) - \gamma_1(r_1(s_k)) u_2(y_{b2}(s_k)) \end{array} \right\} \end{aligned} \tag{5}$$

where δ represents the dynamic discount rate. The contract is enforced if the expected discounted utility surplus exceeds the (social) cost of maintaining a gift network of size g. We introduce this social cost as:

Assumption 4 (Cost of maintaining network increasing in network size.). *If a household does not pursue autarky, its lifetime discounted utility has to be at least as large as large as $\alpha(s', g_i) \geq 0$, a cost associated with maintaining participation in the network's inter-household transfer practices. This cost is increasing in network size such that $\frac{\partial \alpha}{\partial g} > 0$.*

[20] We abstract from states where both receive identical, positive public income shocks.

Here $\alpha(h_t, g_1)$ is the incremental cost to household 1 of maintaining a gift-giving link with household 2 given network size g_1. Together, these assumptions allow us to introduce the contract's implementability constraint, which states that gains from the contract be at least as high as the no-transfer rule for both parties: $U_1(h_t) \geq \alpha(h_t, g_1)$ and $U_2(h_t) \geq \alpha(h_t, g_2)$.

Limited Commitment Contract Solution. The above ingredients represent a simultaneous game in which agents find an implementable contract in the presence of limited commitment and no external enforcement mechanism. The solution will be a dynamic program conditional on the current state s, given targeted discounted utility gain for household 2, U_2^s, per 2's reaction function (Foster and Rosenzweig, 2001; Ligon, Thomas and Worrall, 2002).[21] Agents choose the net transfer by household 1 to 2 and the other recipients in its network, τ_s and the continuation utilities U_1^k and U_2^k for each possible state k. The value function for household 1 depends on current target utilities and resources observable to household 1: $U_2^s, \{y_1(s) + y_{2b}(s)\}$. Furthermore, the program includes five constraints: a promise-keeping constraint (λ_s) to ensure future transfers cohere with the utility promised in state s, two rationality constraints (μ_k and ϕ_k, one for each household), and two non-negativity constraints on consumption (ψ_{1s} and ψ_{2s}). The formal setup and description of constraints are described in appendix section B.1.

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

$$\frac{u_1'(y_1(s) - \tau_s g(s)) + \gamma_1(g,s)u_2'(y_{2b}(s) + \tau_s)}{\gamma_2(g,s)u_1'(y_1(s) - \tau_s g(s)) + u_2'(y_{2b}(s) + \tau_s)} = \lambda_s \tag{6}$$

$$-U_1^{k\prime}(U_2^k) = \frac{\lambda_s + \phi_k}{1 + \mu_k}, \quad \forall k \in S \tag{7}$$

$$\lambda_s = -U_1^{s\prime}(U_2^s). \tag{8}$$

[21] $U_2^s(U_1^s)$ is defined by equation 14 when one replaces all 1 subscripts with 2 and vice versa.

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.[22] For each state s, there is a history independent, non-empty interval $[\underline{\lambda}_s, \overline{\lambda}_s]$ that constitutes 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 $\lambda(h_t)$:

$$\lambda(h_{t+1}) = \begin{cases} \underline{\lambda}_s \text{ if } \lambda(h_t) < \underline{\lambda}_s \\ \lambda(h_t) \text{ if } \underline{\lambda}_s \leq \lambda(h_t) \leq \overline{\lambda}_s \\ \overline{\lambda}_s \text{ if } \lambda(h_t) > \overline{\lambda}_s. \end{cases} \tag{9}$$

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.[23]

From these conditions we can derive several useful propositions. First, there exists a finite network size $g(s)$ at and above which social taxation pressures induce household 1 to optimally resort to the no transfer rule. This is our 'shutdown hypothesis'.

Proposition 1 (Beyond some network size, social taxation induces "shut down".)**.** *For any $\epsilon > 0$, if household 1 is a net sender of transfers under contract rule $\overline{\lambda}_s - \epsilon$ and social taxation pressures $g(s)$, then there exists a $k > 0$ such that $g(s)' = g(s) + k$ will induce household 1 to revert to the no transfer policy.*

The proof is fairly straightforward. Because λ_s is bounded from above by $\overline{\lambda}_s$, and the left hand side of equation 6 is increasing in $g(s)$, there exists some $g(s)'$ that induces the left hand side to increase beyond the upper bound on λ_s. This is true whenever household 1 is a net sender of transfers at $\lambda = \overline{\lambda}_s$.

[22]For a formal proof, see Ligon, Thomas and Worrall (2002) or Thomas and Worrall (1988). Foster and Rosenzweig (2001) extend the result to altruistic preferences.

[23]Notice that at the top of each contract interval ($\lambda(h_{t+1}) = \overline{\lambda}_s$), U_1^s captures the entire surplus making $U_2^s = \alpha(r, g_2)$ when $\lambda(s_t) = \overline{\lambda}_s$. Thus, we can solve for all contract intervals by plugging the solution for τ_s along each endpoint and setting lifetime expected utilities equal to zero for the relevant value of λ_s.

However, when $g(s)$ is too large, household 1 will revert to the no transfer rule or "shut down."

Income shocks. We define two types of exogenous positive income shocks to allow the model to interact more directly with our experiment: 1) privately revealed income (denoted by y_v) and 2) publicly revealed income (y_b). Households that do not receive an income shock experience zero exogenous income shocks (y_z). While there are potentially nine different realized states s, we limit analysis to those states in which only up to one household receives a shock of any type. To simplify notation, we let the set of states be such that $s = \{zz, vz, bz, zv, zb\}$ denote the state combination of household 1 and 2, respectively.[24] We assume the following relationships across state-contingent incomes for household 1 (with an analogue for household 2):

Assumption 5 (Prize-winners have higher incomes)**.**

$$y_1(kz) > y_1(zz) \textit{ and } y_1(kz) > y_1(zk) \textit{ for all } k \in \{v, b\}$$

To highlight the role of social pressure, we make an assumption on the relationship between $r_i(s_t)$ and observability of income. Recall $r_i(s_t) = gp_i(s_t)+1$. Here, we 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 6 (Observable income induces social taxation)**.**

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

This assumption implies that households who enjoy observable windfall gains face social pressure to give a portion of those gains to others (e.g., Jakiela and Ozier (2016), Goldberg (2017) and Squires (2017)). It also reflects the infeasibility of hiding income in the public income state.

3.1 Model Implications

Given the complex state space, it is impossible to analytically 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

[24] zz, neither 1 nor 2 receive an income shock; vz household 1 receives a private shock; bz household 1 receives a public shock; zv, household 2 receives a private shock; zb, household 2 receives a public shock. Thus, $y_{2b}(s_t)$ in our model is equivalent to $y_2(zb)$.

Figure 1: Amount of Transfer by Network Size

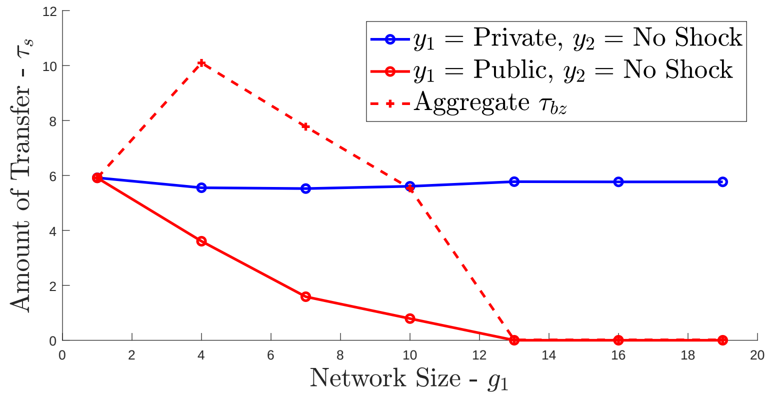

Note: The solid lines represent transfer amounts $\underline{\tau}^s$ from household 1 to household 2 when household 1 takes the entire share of the surplus ($\lambda_s = \underline{\lambda}_s$) and when household 1 wins a cash prize – i.e., the state is either vz or bz. The dashed red line represents the total value of transfers household 1 expects to give to all members of its network.

solutions using set values for model parameters while allowing network size to vary. These simulations are summarized in appendix section B.2.

This model leads to a set of formal empirical predictions.

Prediction 1 (The Shutdown Hypothesis)**.** *Households with large gift-giving networks that experience positive and publicly-revealed income shocks have an increased likelihood of shutting down, resulting in zero transfers (gross) to others. Similar households that experience positive and privately-revealed income shocks will continue to maintain positive net transfers to others over a larger range of network sizes.*

Figure 1 uses simulated gift transfers between households 1 and 2 to show the implications of the shut down hypothesis. Notice that when $g_1 = 1$, household 1 transfers the same amount to household 2 regardless of being in state vz or bz. This signals that the insurance motive is slack; altruism suffices to explain an observed dyadic transfer. Then, as household 1's network size initially increases, transfers to household 2 decrease; however, aggregate transfers to all households in its network increase. In this interval, household

1 is not overly burdened by social taxation pressures, so altruism and mutual insurance both motivate its transfers. However, as the network size increases beyond 4, aggregate transfers also begin to decrease. The increasing pressure to give decreases the marginal utility of household 1 transfers until it is no longer compelled to give anything out of its windfall gains at the shutdown threshold (network size greater than 13 in our simulation).[25] By contrast, household 1's total gifts in the wake of private income windfalls are largely invariant with respect to network size, although the identity of the recipient(s) could vary as network size grows. This leads to two additional 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 shut down threshold.*

Predictions 2 and 3 are related. Relative to a private income shock, a household with a public income shock faces more requests to share its windfall income and will therefore transfer a large number of gifts to other households if it does not shut down (Prediction 2). However, gifts given following private shocks are larger than those following public shocks (Prediction 2). This also implies that the total volume of transfers will be larger following public shocks, but only for households with small networks (Prediction 4).

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.*

Households' giving behavior should be reflected in the consumption be-

[25]Note that this prediction differs from that of small group advantage in collective action theory (Olson, 1971; Ostrom, 2015; Platteau, 2000). Here we assume away gains from collective action beyond those arising from the insurance contract between agents. Likewise, our two agent model differs from network models that predict that larger networks negatively affect outcomes because network size is negatively associated with network closure, and thus with trust that enhances cooperative behavior (Coleman, 1990; Allcott et al., 2007)

havior of connected households. Altruistic preferences and the demands on public vs. private income only imply that consumption will increase in average network *private* income shocks for households with relatively high marginal propensity to consume. This leads to prediction 5.

Prediction 5 (Consumption Increasing in Others' Private Winnings)**.** *A household's per capita consumption increases in its network's average private lottery winnings, especially in households with low per-capita consumption (and, hence, high marginal propensity to consume).*

A final prediction emerges from the dynamic nature of the limited commitment model and Assumption 3. Because households who default on insurance contracts revert to autarky in all future periods, households that "shut down" and renege on informal insurance contracts following a public income shock will be less likely to receive transfers in future periods. Relatedly, because households are not subject to social taxation following private income shocks, the contract intervals are more likely to overlap with the zero shock state (zz) regardless of network size. Thus, gift receipts are less likely to change in future periods for households who win the private income shock than for households facing public shocks that might induce them to cease giving in the face of excessive social taxation pressures.

Prediction 6 (The Dynamic Cost of Shutdown)**.** *Gift receipts are decreasing in network size for households who received a public income shock in a past period. This relationship arises because of the shutdown hypothesis that interacts network size and public income shocks. conversely, receipt of transfers is invariant with respect to network size for households who received a private income shock in a past period.*

4 Empirical Investigation

The model implications derived in Section 3 call for additional data. Specifically, Predictions 1 through 4 require measures of network size. Prediction 5 requires measures of consumption and network lottery winnings. We detail our methods of constructing each of these measures below. Then we describe how we test the predictions of the enhanced model.

4.1 Additional Data

Social Network Data. After selecting the sample but before collecting baseline data a detailed enumeration of respondents' social contacts was conducted. Each respondent was asked in turn (and in random order) about every other respondent in the survey sample from his or her community. We can exactly identify the directionality of giving, including each of the bi-directional, or reciprocal, gift links in our sample. We examine responses to the following two questions: 1) "Have you ever received a gift from [name$_j$]" and 2) "Have you ever given a gift to [name$_j$]"? When both i and j respond "yes" to these questions, a reciprocal gift link exists between these two individuals — i.e., $gi_{ij} = 1$ (and zero otherwise).

We consider two households to be linked in a reciprocal gift giving relationship if at least one household head or spouse engages in mutual (reciprocal) gift-giving with at least one head or spouse of the other household. Out of the 26,795 possible links observed between any two households in our data (across the four villages), we observe 3,866 instances of reciprocal gift-network links between households, 14.4% of possible links.[26] Network size was calculated using the total number of links created in this manner at the household level. Notice that the social network mapping and gift transfer data are sourced from separate modules — gift network responses precede information about actual gift transfers across all rounds. Further details regarding social network data construction are described in appendix section A.3.

Survey Data — Consumption. The survey expenditure module solicited detailed information on the quantities and values purchased of many items, including home produced and 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. In addition to reporting individual summary statistics, appendix table

[26] Of 3,866 observations, 42% of links are sustained by the household head only, 29% by the spouse only, and the remaining 29% by both the head and the spouse.

D.3 demonstrates 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.

This provides justification for a household-level analysis. Given that the household head and spouse seem to coordinate most closely around total household food consumption, and that the income shocks we generated experimentally are likely observable within households (even if unobservable to others outside the household), we aggregate variables at the household level.[27] We do this by taking the household sum of all expenditures reported by the individuals who incurred the expenditure.[28] We focus on food expenditures because the combination of the physiological need to eat frequently and the lack of any significant carryover of food over a period of two months between survey rounds ensures that food expenditures represent a period-specific flow measure of consumption, where ceremonial, durables, educational, health, or other expenditures are far more vulnerable to episodic or seasonal variability that can mask the consumption effects we seek to test.

Combining Experimental Data with Social Network Data — Lottery Winnings of the Gift Network. One approach to calculating gift network lottery winnings is to analyze the average cash winnings (private vs. public) of each household's gift network. The measurement of the network average lottery winnings, however, requires an additional consideration. We want to test the prediction that the consumption of poorer households increases when their network receives private income shocks. However, the size of the transfer received by the poorer household depends on the network size of the sending household. This is because the size of the transfer between, say, household 1, the one that receives the positive income shock, and house-

[27]For food expenditures, we sum 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 an analysis of how information asymmetry influences spending decisions within the household in these data.

[28]If one of either the head or the spouse was unable to report expenditure in a given round, we indicate that household expenditure is missing for that round.

hold 2, the household receiving the transfer, also depends on household 1's network size. Therefore, a more theoretically appropriate network average adjusts network winnings by household 1's network size.

To provide intuition, consider that household j has gift obligations to X other households. If household j receives a positive income shock and wants to allocate some portion of this shock, Y, to the X other households in its network, then, on average, $\frac{Y}{X}$ will be allocated to any given household in its network. Formally, the adjusted average amount received by household the adjusted network average is

$$\overline{\text{Private}}'_{it} = \sum_{j=1}^{N} \frac{\frac{\text{Private}_j}{\sum_{k=1}^{N} \mathbb{1}(g_{jk}==1)} \times \mathbb{1}(g_{ij} == 1)}{\sum_{j=1}^{N} \mathbb{1}(g_{ij} == 1)}. \tag{10}$$

The fraction in the numerator represents the weight placed on each household j's lottery winning in household i's network.

The top panel of Table D.9 presents our measure of network size. The average network size, defined by the number of inter-household reciprocal gift-giving links, is 11.3 but varies substantially with a standard deviation of 10.1.[29] Roughly 13% of the households do not have reciprocal gift giving links with any other household in the sample, consistent with observations in the 2004 survey round (Vanderpuye-Orgle and Barrett, 2009). Household per capita monthly food consumption, reported in the second panel, averages 24.20 GH¢, 75% of which is purchased food. So cash income clearly limits food consumption. Notice that the maximum size of the cash prize is close to four times the monthly per capita purchased food consumption. The bottom panel presents the average value of own and network cash winnings and shows that average prize winnings roughly correspond to the expected value of the cash prize of all households in the village sample.

4.2 Empirical Tests

The unique features of our experimental design allows us to test the model predictions in a straightforward manner. Let y_{it} again be the outcome of interest: either the (total or average) amount or number of round t gifts distributed

[29] Figure E.1 displays the distribution of network size. Our analysis is unlikely to be skewed by the presence of households with large, outlier network sizes.

by household i. The shut down hypothesis (Prediction 1 in Section 3) can be investigated using the following regression:

$$\begin{aligned} y_{it} = \alpha &+ \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} \\ &+ \beta_{vg} \text{Private}_{it} \times \text{Net-size}_i + \beta_{bg} \text{Public}_{it} \times \text{Net-size}_i \\ &+ \text{hh}_i + \text{r}_{tv} + \epsilon_{it}, \end{aligned} \tag{11}$$

where the estimation proceeds exactly as it did when testing the public observability hypothesis previously. The refinement here is to interact private and public winnings with the household's ex ante reciprocal gift network size (Net-size$_i$). Note that household fixed effects control for all time-invariant household factors, including the size of its gift network. Time-varying unobservable characteristics of household i are captured in the residual, ϵ_{it}.

Network size could proxy for an omitted variable or variables (e.g. personality traits, preferences, family background) that lead individuals to form smaller (larger) networks and also be more (less) generous when they earn windfall income. This could be a direct confound with the measure of baseline network size. This does not matter materially since we are interested in network size as a household attribute, which could of course proxy for other attributes.[30] Nevertheless, we show that our results are robust to alternative definitions of networks in section 5.

Table 2 contains the estimation results of equation 11 with three different outcome variables, with and without interaction terms. The significant negative coefficient in the fourth row (β_{bg}) of columns 1-3 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). The results combined suggest that when network size is small, the cash prizes substantially increase the number and value of gifts given whether or not the income shock is public or private. Furthermore, there is very little difference between gift-giving behavior in the public and private settings when network size is small — we cannot reject

[30]This is no different than how we interpret the gender or age or educational attainment of a household head as observable attributes that yield useful predictions despite being almost surely correlated with other, unobservable attributes.

Table 2: Testing the Shut Down Hypothesis

Dependent Variable		Gift Giving Value (Total) (1)	Gift Giving Value (Average) (2)	Gift Giving Number (3)
Randomized Explanatory Variables With Network Size Interaction				
Value of Private Cash Prize	$\beta_v > 0$	0.296**	0.261**	0.267**
		(0.130)	(0.104)	(0.108)
Value of Private Cash Prize × N	$\beta_{vg} \leq 0$	-0.006	-0.006	-0.003
		(0.008)	(0.006)	(0.006)
Value of Public Cash Prize	$\beta_b > 0$	0.408***	0.190*	0.471***
		(0.132)	(0.104)	(0.111)
Value of Public Cash Prize × N	$\beta_{bg} < 0$	-0.036***	-0.016*	-0.042***
		(0.012)	(0.009)	(0.011)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v = \beta_b$		.541	.628	.181
P-Value $H_0 : \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		.793	.284	.956
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		.139	.085	.070
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		.014	.067	.001
N at Shut Down		11.46	11.62	11.14
Left-censored Obs.		1,182	1,182	1,182
Observations		1,586	1,586	1,586

Note: $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals log total value of cash gifts given in hh in column 1; log average value of cash gift given in column 2; number of gifts given in column 3. Value of Private/Public Cash prize is divided by $10 = \in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns with a lower bound of zero. Model predictions associated with each coefficient (predicted sign) displayed next to each coefficient. Null hypotheses are tested using Wald tests of equivalence specified for network size (N) of 0, 5, 10 and 20. N at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$.

that gift-giving behavior is equivalent for a network size of zero to 5 across any of the specifications. However, by the time the network size is equivalent to, roughly, the average (11.3), we can reject similarity in gift-giving behavior across specifications 2 (average value of gift given) and 3 (number of gifts given). We calculate the shut down point predicted by the linear model as a network size of 11.46, 11.62, and 11.14 for columns 1-3 respectively. In other words, households give zero additional gifts following public income shocks when they have around 11 other households in their gift giving network.

Predictions 2 and 3 do not depend on heterogeneous network size. They can be tested by setting the interaction terms equal to zero. Prediction 2 states that $\beta_v > \beta_b$ with respect to the average value of gifts given, which is supported by our findings, while simultaneously allowing $\beta_b > \beta_v$ for total number of gifts given (Prediction 3) and total value of gifts given (Prediction 4) for small network sizes. Column 2 in table 2 shows that the point estimate for β_v is larger than β_b for all network sizes. Columns 1 and 3 show that β_b is larger than β_v only for small network sizes. All three results are consistent with predictions 2 through 4.

Evidence for predictions 3 and 4 are clearer when analyzed graphically. Figure 2 shows the linear combination of coefficients associated with a third-order polynomial specification of public and private winnings interacted with network size regressed against the total value of gifts given.[31] This semi-parametric specification shows that the confidence intervals of effect sizes have very little overlap with a network size of zero but the effect size of public winnings quickly decreases as network size increases while the opposite is true in private winnings.

Together, the results related to the first four predictions suggest a clear pattern of transfer behaviors. Households with small network sizes act similarly upon winning the privately revealed or publicly revealed cash prize: they increase the number of gifts given, the total value of gifts given and the average value of gifts given by roughly similar amounts. But as the network size increases, behaviors begin to diverge depending on the observability of the income windfall. Public income has no effect on giving with large network sizes, which suggests that the social demands on the lucky household induce default on informal sharing arrangements.

Testing Prediction 5. Empirical investigation of the model's implication for consumption (Prediction 5) relates household i's consumption expenditures to the average lottery winnings of i's gift network — i.e., the average network treatment effect on per capita food consumption, our preferred proxy for consumption in these data. We test this using the following equation:

$$y_{it} = \alpha + \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} + \beta_{vn} \overline{\text{Private}}'_{it} + \beta_{bn} \overline{\text{Public}}'_{it} + \text{r}_{tv} + \epsilon_{it}, \quad (12)$$

where y_{it} is log per capita household food consumption, $\overline{\text{Private}}'_{it}$ represents our theoretically preferred measure of network average private cash lottery

[31] Figure E.2 shows a similar plot regressed against the total number of gifts given.

FIGURE 2: Shut-down Hypothesis on Total Value of Gifts Given

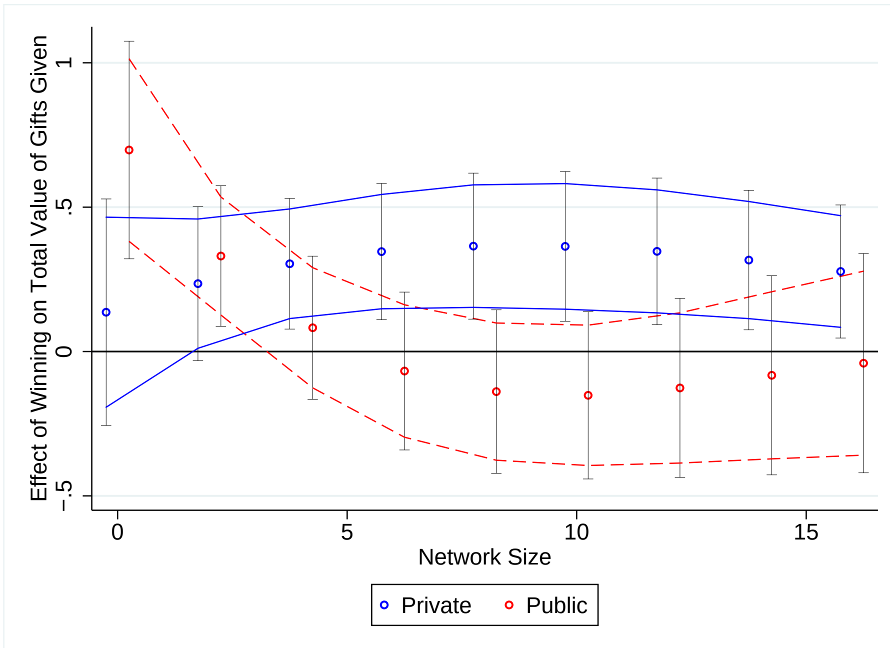

Note: Dependent variable equals log total value of gifts given. Estimation includes 3rd order polynomial interactions on network size. Dots represent point estimates. Blue (dotted red) line represents 90% confidence interval for private (public) winnings. Bars represent 95% confidence intervals. Plots of public coefficients offset by one for ease of viewing.

winnings in i's network at time t, and $\overline{\text{Public}}'_{it}$ is the analogous measure for the household's network's average public cash winnings that period.[32] We again include village-specific round fixed effects, r_{tv}.

Given the concavity of utility, households with lower period-specific food consumption should receive more support from their network. This has three implications for estimation. First, we no longer include household fixed effects because changes to consumption will be larger for households with lower levels of consumption. In other words, the average deviations implied with household fixed effects are not desirable. Second, we opt to use a quantile regression estimator to examine effects at different locations along the consumption distribution. We expect network effects of private income to be larger at the lower

[32] Appendix Figure E.3 shows similar results for $\overline{\text{Private}}_{it}$ and $\overline{\text{Public}}_{it}$.

end of the distribution. Third, we focus primarily on observations from rounds two and three of the data, the pre-harvest season when farming households are most food constrained as they await the next season's harvest.[33]

Finally, we note that, our measure of the network average is sensitive to outliers, which can negatively influence inference in the analysis. The distribution of $\overline{\text{Private}}'_{it}$ (or $\overline{\text{Public}}'_{it}$) approximates a normal distribution when network size is large. However, $\overline{\text{Private}}'_{it}$ can have very high values when network size is small. To allow for a more normal distribution of $\overline{\text{Private}}'_{it}$, we use log transformations of the adjusted network average.

We focus on the 1st, 12th, 25th, 50th and 75th quantiles to emphasize trends in the lower end of the food consumption distribution. We graphically depict the results of the simultaneous quantile estimation of Model 12 in Figure 3 (appendix Table D.10 shows estimation results for each quantile). The lower the per capita food consumption, the larger is the adjusted network average effect of private lottery winnings on food expenditures. In Figure 3, the coefficient estimates on private average network lottery winnings, represented by the blue dots and lines, are significantly positive and greater than zero for quantiles below the 50th percentile. Furthermore, the estimated increase in consumption following the network's private lottery winnings is statistically significantly larger than the estimated change in consumption following the network's public lottery winnings (the latter is insignificantly different from zero throughout). These results are consistent with both altruistic motives for giving and the shut down hypothesis, as reflected in Prediction 5 of our model.

Testing Prediction 6. The shut-down hypothesis implies that households choose to exit reciprocal transfer agreements when network size is too large. If they refuse to give in a state when others expect them to give, then they may become less likely to receive transfers in the future, a consequence of defecting from the informal contract (Prediction 6). In our case, we expect that households with large networks who also won the public cash prize will

[33] Appendix Figure E.4 shows how home-produced food consumption over the past month varies by survey date. Home-produced food is clearly least available from mid-March to early July, corresponding to survey rounds two and three.

FIGURE 3: Effect of Network Winnings on Food Consumption by Quantile

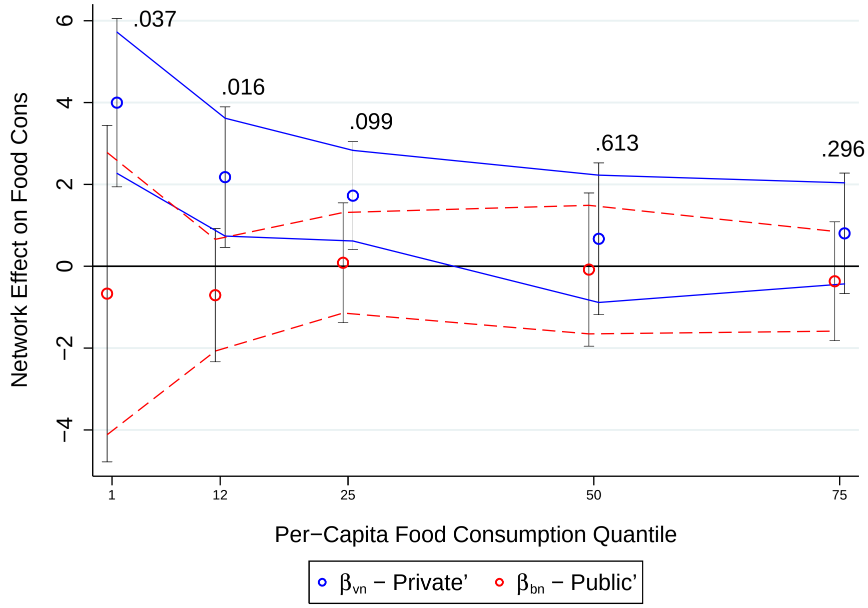

Note: Results of a simultaneous quantile regression at 1st, 12.5th, 25th, 50th, and 75th quantiles bootstrapped over 1,000 iterations. Dependent variable is log home-produced per capita food consumption over the last month. Blue dots (lines) show the coefficient estimates (90% confidence interval) on adjusted private network winnings, $\overline{\text{Private}}'_{it}$, at each quantile. Red represents public network winnings, $\overline{\text{Public}}'_{it}$. Blue dots offset for ease of viewing. The numbers above each point are the quantile-specific p-value of the Wald test $\text{H}_0: \beta_{vn} = \beta_{bn}$.

be less likely to receive subsequent transfers due to shutting down.

Table 3 tests this hypothesis by estimating a variant of equation 11 with the dependent variable the number of gifts received, which we regress on the income shock from the same round and lagged income shocks received in any round prior to round t. Column (1) only includes lagged regressors, column (2) only includes lagged regressors and omits households who received income shocks in round t and column (3) includes all observations as well as both contemporaneous and lagged regressors. The estimation results mimic those in Table 2. Households who win the public prize and have large networks are less likely to receive transfers in both the same round and a future round

Table 3: The Dynamic Cost of Shutting Down

Dependent Variable:	Number of Gifts Received		
	(1)	(2)	(3)
Lagged Randomized Explanatory Variables			
Won Private in Past	-0.123	-0.231	-0.148
	(0.273)	(0.297)	(0.283)
Won Private in Past × N	0.003	0.013	0.007
	(0.015)	(0.016)	(0.016)
Won Public in Past	0.718*	1.007***	0.689*
	(0.368)	(0.370)	(0.376)
Won Public in Past × N	-0.051**	-0.066***	-0.050**
	(0.022)	(0.024)	(0.023)
Contemporaneous Randomized Explanatory Variables			
Won Private			-0.168
			(0.335)
Won Private × N			0.011
			(0.022)
Won Public			0.604**
			(0.257)
Won Public × N			-0.041**
			(0.018)
Network Size			
N	0.007	0.009	0.009
	(0.007)	(0.007)	(0.007)
Round × Village FE	Yes	Yes	Yes
Observations	1,586	1,431	1,586

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals number of transfers received across all columns. Poisson estimator used in all columns. Standard errors in parantheses clustered at the household level. N denotes network size. Column (2) omits households who receive "contemporaneous" (round t) windfall gains. "Lagged" variables indicate receipt of windfall gain during any round prior to t.

(up to eight months after the initial public lottery winning). On the other hand, households with smaller networks who win the public lottery become more likely to receive gifts from the reciprocity of others, presumably because the early-round recipient demonstrated fidelity to the informal contract. We include both contemporaneous and lagged shocks to income to show that this effect takes place almost immediately.

The weak positive result on private winnings (β_v) suggests that households who give from private winnings may not see their gifts reciprocated in future rounds. This is expected with altruistic giving — one is not giving to others in expectation of a future reciprocated transfer. These results carry a powerful implication. If households have large networks, then public transfers may not only crowd out near-term altruistic transfers, they may also isolate individuals from extant gift networks, which could reinforce non-altruistic behaviors.

5 Robustness Checks and Extentions

Altruism → directional gifts to relatively needy. The extremely detailed micro-structure of our data offers an alternative estimation strategy to test the model's predictions and to look further into underlying mechanisms. We will first conduct an additional test of Prediction 5.

An additional test of prediction 5 is one in which we confirm that a "better off" household transfers resources to a relatively worse off household upon winning the private lottery, as opposed to the public lottery. In other words, the degree of giving out of private income depends on the difference between the giver's and recipient's food consumption. To examine this prediction in our data, we can estimate the following dyadic regression:

$$\begin{aligned} y_{ijtv} = {} & \alpha + \mathrm{r}_{tv} + \beta_v \mathrm{Priv}_{it} + \beta_b \mathrm{Pub}_{it} + \gamma(\mathrm{Food}_{it} - \mathrm{Food}_{jt}) \\ & + \beta_{vF} \mathrm{Priv}_{it} \times (\mathrm{Food}_{it} - \mathrm{Food}_{jt}) + \beta_{bF} \mathrm{Pub}_{it} \times (\mathrm{Food}_{it} - \mathrm{Food}_{jt}) + \epsilon_{ijt} \end{aligned} \tag{13}$$

where y_{ijt} represents giving from household i to household j either in terms of amount given or number of gifts given. Then, $(\mathrm{Food}_{it} - \mathrm{Food}_{jt})$ is the difference between household i and j's period t per capita food consumption. The larger the value, the more likely i is to give to j after winning the private lottery (under altruistic preferences), i.e., we predict $\beta_{vF} > 0$.

Of all the instances of within-village gift-giving reported in the survey's gift module, 10% of gifts given could be traced to gifts given to other sample households. Table 4 focuses on these instances of gift giving and columns 1 through 3 in limit the sample to those households who were linked to one another in

Table 4: Dyadic Regressions

Dependent Variable:		Gift Giving Within Dyad: From i to j Amount (1)	Number (2)	Amount (3)	Amount (4)
$(\text{Food}_{it} - \text{Food}_{jt})$	γ_F	0.073	0.029		
		(0.204)	(0.106)		
Network Size	γ_g			-0.036	-0.017
				(0.027)	(0.018)
Randomized Explanatory Variables With Interactions					
Value in Private	β_v	0.182	0.136*	0.318	0.239
		(0.153)	(0.078)	(0.235)	(0.157)
Value in Private × $(\text{Food}_{it} - \text{Food}_{jt})$	β_{vF}	0.305**	0.117**		
		(0.127)	(0.058)		
Value in Private × N	β_{vg}			-0.005	-0.007
				(0.009)	(0.009)
Value in Public	β_b	-0.286	-0.234	0.177	0.341**
		(0.265)	(0.166)	(0.399)	(0.164)
Value in Public × $(\text{Food}_{it} - \text{Food}_{jt})$	β_{bF}	-0.098	-0.055		
		(0.064)	(0.042)		
Value in Public × N	β_{bg}			-0.034	-0.044***
				(0.025)	(0.016)
Round × Village FE		Yes	Yes	Yes	Yes
All Dyads Included		No	No	No	Yes
P-value $H_0 : \beta_v = \beta_b$		0.12	0.05	0.76	0.64
P-value $H_0 : \beta_{vF} = \beta_{bF}$		0.00	0.01		
Left-censored Obs.		16,190		16,190	107,944
Observations		16,270	16,270	16,270	108,082

Note: $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable is log total value of (cash) gifts given from household i to household j in columns 1, 3 and 4 — Tobit estimates with observations censored at zero. Number of gifts in column 2, estimated using Poisson distribution. Value in Private/Public $\in \{0, 1, 2, 3.5, 5, 7\}$. $\text{Food}_{it} - \text{Food}_{jt}$ is Δ log per capita food consumption. Analysis only includes dyads in reciprocal gift-giving network at baseline in columns 1 through 3. All within-sample dyads represented in column 4. Standard errors clustered by dyad. N is network size.

the social network at baseline. We estimate equation 13 using Tobit and Poisson estimators when the amount given and number given are the respective dependent variables. Estimates in columns one and two reflect equation 13 estimated for the amount and the number of gifts given, respectively. The estimation results are consistent with Prediction 5. In both columns, gift giving increases after winning a private lottery but not after winning a public lottery. Furthermore, the effect is statistically significantly stronger when household i's food consumption is larger than household j's.

Selfish Network Formation? It is difficult to reconcile the strategic transfers following private shocks, i.e., selfish giving as a means of building network ties, with the observation that transfers flow towards relatively needy households. Furthermore, we do not find evidence that instances of transfers between two households with no prior reciprocal gift link increases following private lottery winnings. However, we do see significant increases in gift-giving to out-of-network households following public winnings.

Columns 3 and 4 in table 4 estimate the shut down hypothesis model in the dyadic setting. Column 4 includes out-of-network dyads while column 3 only includes households specified to maintain mutual gift-giving links at baseline. The results show that when gifts are given out of public lottery winnings, they are more likely to be given to individuals who are not in one's mutual gift giving network. However, this is not the case with respect to private winnings, which are more likely to be given to prior gift network members (columns 1 and 2). We see in column 3 that β_b is not significant while it is significant in column 4 with the expected shut down effect present in the negative β_{bg} coefficient. Thus, we do not find the behavior we would expect from households who are seeking to build network ties with their transfers.

Endogenous Network Size. As mentioned earlier, we acknowledge that network size could proxy for omitted variables, rendering it an endogenous regressor that biases our results. We explore alternative measures of networks in appendix Tables D.11 through D.13 and show consistency with the results obtained thus far. Specifically, Table D.11 shows that the total number of non-co-resident within-village family members (family network size) is the strongest predictor of gift network size. That one variable alone explains nearly 50% of the variation in gift network size in our sample. Assuming that family network size is exogenous, since fertility or marriage decisions in response to our lotteries experiment seems highly unlikely, we replace gift network size with family network size in Table D.12 and obtain similar results. We also generate a linear probability model to predict gift network size (column 1, Table D.11) and obtain qualitatively similar results when we use predicted network size (Table D.13). We conclude that endogenous network selection is

Table 5: Giving Private Lottery Winnings to Friends, not Family

Dependent Variable:		Value of Gifts Given (Average)		
Gifts directed to:		All Family (1)	Direct Family (2)	Village Friends (3)
Randomized Explanatory Variable With Network Size Interaction				
Won Private Cash Prize	β_v	-0.298	-1.065	0.875**
		(0.726)	(0.828)	(0.431)
Won Public Cash Prize	β_b	1.912***	2.029***	1.287***
		(0.686)	(0.652)	(0.491)
Won Private Cash Prize × N	β_{vg}	0.0237	0.0442	-0.0157
		(0.044)	(0.046)	(0.029)
Won Public Cash Prize × N	β_{bg}	-0.120**	-0.101**	-0.118**
		(0.051)	(0.049)	(0.048)
Round × Village FE		Yes	Yes	Yes
N at Shutdown		16	20	11
Left-censored Obs.		1,173	1,307	1,340
Observations		1,561	1,561	1,561

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable is log average value of (cash) gifts given. Column 1 shows gifts to all family, column 2 to direct family members (i.e., siblings, grandparents, parents) within the village, column 3 to village friends. Won in Private/Public $\in \{0, 1\}$. Tobit estimator used in all columns. N denotes network size.

not a major threat to our results.

Information Hypothesis. Households who win the private prize might not be able to conceal this fact from those who are close to them, such as non-co-resident family members within the village. This seems unlikely since within-family food consumption is likely to be correlated (and hence Prediction 5 would not have been confirmed).[34]

We explore this possibility in Table 5, differentiating gifts given according to links with varying likely quality of information about recipient households. We again estimate equation 11, with the dependent variable log value of gifts given to all kin (i.e., extended family) in column 1, to direct family (i.e., non-coresident siblings, parents, or grandparents) in column 2, and to village friends in column 3, assuming that information is more difficult to conceal from

[34]Furthermore, using the same experiment, Castilla and Walker (2013) show that even spouses did not necessarily know whether the other won a private prize.

non-co-resident family members. Contrary to the information hypothesis, gift giving to direct family members does not flow from private lottery winnings, while gift giving to village friends does. Gift giving to family and friends both shut down following public cash winnings. Thus there seems no information story to explain the patterns we observe in the data.

The Table 5 results, however, fit the prediction of the canonical limited commitment informal insurance model quite well for the special case of publicly observable income shocks within small-to-moderate-sized family networks — i.e., in the neighborhood of the sample mean or median. Specifically, in small networks, public income is shared among family members after winning the public prizes but not after winning the private prize. This suggests that an insurance motive is more likely when giving to family members. Our analysis suggests that there is a range of our data within which the familiar insurance model seems to work very well. For a large share of our sample households, however, their networks are too large to fit the canonical model without altruistic preferences and shutdown due to network overload.

Additional Extensions and Robustness Checks. Appendix section C provides comments on a series of other extensions and robustness checks we have explored. This discussion shows 1) households who win public prizes are not coordinating with one another to give less to others, 2) it is unlikely that mental accounting of stochastic unearned income explains the results, 3) households are unlikely to be endogenously opting out of informal insurance contracts when they receive large income shocks and 4) precautionary savings or investments in others are unlikely to motivate giving out of private shocks, especially since they are more progressively targeted than giving proceeding from public income shocks. In summary, our evidence points toward solidarity networks motivated only partly by insurance.

Test of Full Risk Pooling. We have demonstrated that altruistic preferences play an important role in motivating transfers between households. Our theoretical framework, however, allows for the possibility that multiple motives may be at play, including risk pooling. Thus, we also explore whether the familiar full-risk-pooling prediction (Townsend, 1994) holds in our data

— that the intertemporal 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. Details of this analysis are in appendix section C.5. We find limited risk pooling within the family network, but not within the village-friend network. These results, combined with those from Table 5, strongly suggest that gifts to village friends - rather than to family - are driven primarily by altruistic motives while transfers to family are more consistent with an insurance motive complicated by social taxation.

6 Conclusions

Inter-household networks within village economies serve multiple functions. They can mediate inter-household transfers that resemble credit, insurance, social taxes, altruistic gifts, or any combination of these. Given the observational equivalence of inter-household transfers - especially in cross-sectional or panel data of short duration - it is easy to misunderstand the function(s) of transfers and impose too restrictive a lens in interpreting inter-personal behaviors. We show empirically that more than pure self-interest seems at play in villages in southern Ghana. We adapt a standard economic model to allow for slightly richer behavior and show its predictions fit the data remarkably well. This more holistic framing of multifunctional social solidarity networks carries important policy and research implications.

Fully transparent - i.e., publicly observable - transfers often aim not only to reduce corruption but also to facilitate progressive redistribution within communities to rectify targeting errors. Our findings caution that if there exist non-trivial social taxation pressures and altruism within the community, then more discrete, private transfers might yield more progressive outcomes. Given the rapid expansion of transparent cash transfer programs globally, this issue seems to merit further research — especially if the degree of malfeasance that is avoided due to transparency is minimal.

Recent research has explored how accounting for moral motivations may enhance the efficacy of a given policy. Bursztyn et al. (2019), for example, show that making the concept of justice more salient in debt repayment messaging decreases default rates. Policy-makers might similarly consider how to make altruistic preferences salient in cash transfer programs, especially since such motives seem to generate progressive redistribution of resources.

Perhaps most importantly, our results caution against an overly simplistic approach to moral considerations in economic settings. In *The Moral Economy*, Bowles (2016) documents numerous instances in which reliance on policies to incentivize behavioral change, modeled around self-interested preferences, end up crowding out moral or ethical motives for actions. In reviewing the book, Kranton (2019) argues that economists need to study more closely social context and local norms so as to better understand the mechanisms through which a reliance on incentives might inadvertently lead to socially harmful outcomes. This paper takes that call to heart. Our results support a less jaundiced view of the social economic behaviors of rural villagers in low-income communities, allowing for greater richness associated with the coexistence of pro-social, altruistic preferences with self-interested behavior and costly social demands within multi-functional social networks (Barrett, 2005).

Bibliography

Alderman, Harold. 2002. "Do Local Officials Know Something We Don't? Decentralization of Targeted Transfers in Albania." *Journal of Public Economics*, 83(3): 375 – 404.

Allcott, Hunt, Dean Karlan, Markus M. Mobius, Tanya S. Rosenblat, and Adam Szeidl. 2007. "Community Size and Network Closure." *American Economic Review*, 97(2): 80–85.

Ambrus, Attila, Markus Mobius, and Adam Szeidl. 2014. "Consumption Risk-Sharing in Social Networks." *American Economic Review*, 104(1): 149–82.

Andreoni, James. 1990. "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving." *Economic Journal*, 100(401): 464–477.

Angelucci, Manuela, and Giacomo De Giorgi. 2009. "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?" *American Economic Review*, 99(1): 486–508.

Barrett, Christopher B., ed. 2005. *The Social Economics of Poverty.* Routledge.

Berry, James, Rebecca Dizon-Ross, and Maulik Jagnani. 2020. "Not Playing Favorites: An Experiment on Parental Fairness Preferences." *National Bureau of Economic Research working paper 26732.*

Bourlés, Renaud, Yann Bramoullé, and Eduardo Perez-Richet. 2017. "Altruism in Networks." *Econometrica*, 85(2): 675–689.

Bowles, S. 2016. *The Moral Economy: Why Good Incentives Are No Substitute for Good Citizens. Castle Lectures Series*, Yale University Press.

Bowles, Samuel, and Herbert Gintis. 2002. "Social Capital And Community Governance." *Economic Journal*, 112(483): F419–F436.

Bursztyn, Leonardo, Stefano Fiorin, Daniel Gottlieb, and Martin Kanz. 2019. "Moral Incentives in Credit Card Debt Repayment: Evidence From a Field Experiment." *Journal of Political Economy*, 127(4): 1641–1683.

Castilla, Carolina, and Thomas Walker. 2013. "Is Ignorance Bliss? The Effect of Asymmetric Information between Spouses on Intra-household Allocations." *American Economic Review*, 103(3): 263–68.

Coate, Stephen, and Martin Ravallion. 1993. "Reciprocity Without Commitment: Characterization and Performance of Informal Insurance Arrangements." *Journal of Development Economics*, 40(1): 1 – 24.

Cole, Harold L., and Narayana R. Kocherlakota. 2001. "Efficient Allocations with Hidden Income and Hidden Storage." *Review of Economic Studies*, 68(3): 523–542.

Cole, Harold L., George J. Mailath, and Andrew Postlewaite. 1992. "Social norms, savings behavior, and growth." *Journal of Political Economy*, 100: 1092–1125.

Coleman, James Samuel. 1990. *Foundations of Social Theory. Belknap Press Series*, Belknap Press of Harvard University Press.

Conley, Timothy G., and Christopher R. Udry. 2010. "Learning About a New Technology: Pineapple in Ghana." *American Economic Review*, 100(1): 35–69.

De Weerdt, Joachim, Garance Genicot, and Alice Mesnard. 2019.

"Asymmetry of Information within Family Networks." *Journal of Human Resources*, 54(1): 225–254.

Dubois, Pierre, Bruno Jullien, and Thierry Magnac. 2008. "Formal and informal risk sharing in LDCs: Theory and empirical evidence." *Econometrica*, 78: 679–725.

Fafchamps, Marcel. 1992. "Solidarity Networks in Preindustrial Societies: Rational Peasants with a Moral Economy." *Economic Development and Cultural Change*, 41(1): 147–174.

Foster, Andrew D., and Mark R. Rosenzweig. 2001. "Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas." *Review of Economics and Statistics*, 83(3): 389–407.

Goldberg, Jessica. 2017. "The Effect of Social Pressure on Expenditures in Malawi." *Journal of Economic Behavior & Organization*, 143: 173 – 185.

Jackson, Matthew O., Tomas Rodriguez-Barraquer, and Xu Tan. 2012. "Social Capital and Social Quilts: Network Patterns of Favor Exchange." *American Economic Review*, 102(5): 1857–97.

Jakiela, Pamela, and Owen Ozier. 2016. "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies." *Review of Economic Studies*, 83(1): 231–268.

Kinnan, Cynthia. 2019. "Distinguishing Barriers to Insurance in Thai Villages." *Journal of Human Resources*.

Kranton, Rachel. 2019. "The Devil Is in the Details: Implications of Samuel Bowles's The Moral Economy for Economics and Policy Research." *Journal of Economic Literature*, 57(1): 147–60.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies*, 69(1): 209–244.

Okeke, Edward N., and Susan Godlonton. 2014. "Doing Wrong to Do Right? Social Preferences and Dishonest Behavior." *Journal of Economic Behavior & Organization*, 106: 124 – 139.

Olson, Mancur. 1971. *The Logic of Collective Action: Public Goods and the*

Theory of Groups,. American Studies Collection. Second Printing with a New Preface and Appendix ed., Harvard University Press.

Ostrom, Elinor. 2015. *Governing the Commons: The Evolution of Institutions for Collective Action. Canto Classics*, Cambridge University Press.

Pitt, Mark M., Mark R. Rosenzweig, and Md Nazmul Hassan. 1990. "Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries." *American Economic Review*, 80: 1139–56.

Platteau, J.P. 2000. *Institutions, Social Norms, and Economic Development. Fundamentals of Development Economics*, Harwood Academic.

Popkin, Samuel L. 1979. *The Rational Peasant : The Political Economy of Rural Society in Vietnam.* University of California Press Berkeley.

Posner, Richard A. 1980. "A Theory of Primitive Society, with Special Reference to Law." *Journal of Law and Economics*, 23(1): 1–53.

Santos, Paulo, and Christopher B. Barrett. 2011. "Persistent Poverty and Informal Credit." *Journal of Development Economics*, 96(2): 337 – 347.

Sen, Arijit, and Karla Hoff. 2006. "The Kin System As a Poverty Trap?" In *Poverty traps.* , ed. Samuel Bowles, Steven Durlauf and Karla Hoff, Chapter 4, 95–115. Princeton:Princeton University Press.

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

Thaler, Richard H. 1999. "Mental Accounting Matters." *Journal of Behavioral Decision Making*, 12(3): 183–206.

Thomas, Jonathan, and Tim Worrall. 1988. "Self-Enforcing Wage Contracts." *Review of Economic Studies*, 55(4): 541–554.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica*, 539–591.

Vanderpuye-Orgle, Jacqueline, and Christopher B. Barrett. 2009. "Risk Management and Social Visibility in Ghana." *African Development Review*, 21(1): 5–35.

Walker, Thomas. 2011. "Risk Coping, Social Networks and Investment in Rural Ghana." *PhD Dissertation.*

APPENDICES

For Online Publication

Altruism, Insurance, And Costly Solidarity Commitments

Vesall Nourani, Christopher B. Barrett, Eleonora Patacchini, Thomas Walker

May 22, 2020

Table of Contents

Appendix Figures

Appendix Tables

A Additional Data Descriptions

A.1 Description of Lottery

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 (i.e., draws were identical and independently distributed, without replacement). A village meeting was held in a central area of the community, commonly outside the village school, and all respondents were explicitly 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 public meeting.[35] There were usually a number of non-respondents at these meetings as well. At each gathering we thanked the participants for their continued support and participation. We explained that survey respondents, and only respondents, had a chance to win one of 20 prizes that day, framing the prizes as a gratuity for their participation in the survey.[36] 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.[37] 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, usually at the home survey visit.

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

[36] Following a protocol approved by Cornell's IRB, respondents signed an informed consent form at the start of the survey, explaining how they would be remunerated for their survey participation. Entry in the lottery and lucky dip was part of this remuneration. In addition to the chance to win a prize, each respondent was given a little cash for their participation, which varied across rounds and was used as an endowment in a private provision of public goods experiment as part of a separate study (Walker, 2011).

[37] If the winner was not present, the prize would be put to the side and delivered to the winner at a later date. But everybody present at the draw heard the name of the individual who won the prize, so the windfall was clearly public knowledge, even if the physical transfer took place privately, later.

After the public lottery prizes were distributed, we conducted a second lottery round in private. Respondents were asked to identify themselves to a member of the survey team, who took their thumbprint or signature and issued them with a ticket displaying their name and survey 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 cash or livestock prizes.[38] Those who drew winning tokens were informed immediately which prize they had won, 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 commonly hid their prizes in their clothes before leaving the room.

A.2 Discussion of Balance Test

Appendix Table D.4 presents balance tests of baseline variables according to whether any member of the household won any public or private lottery over the course of the year. 136 of the households in the study are thus in our "treatment" group while the remaining 181 did not win a cash prize. We analyze baseline values of the core variables used in our analysis and we cannot reject differences in mean values for any of these measures at the 10% confidence level. Furthermore, the P-value associated with the F statistic of joint significance is 0.53. In total, these statistics suggest that randomization was successful.

[38] 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. There were usually 25-30 non-attendees and less than three prizes remaining. There were many checks (and staff) in place to ensure that cash prizes were distributed to their intended households — we do not suspect any problems similar to those reported in Okeke and Godlonton (2014) where vouchers were misallocated by field staff.

A.3 Additional Network Data Details

In addition to questions regarding gift-giving links, the social network module of the survey asked whether they knew the person, by name or personally, how often they saw him/her, 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. Questions regarding gift links are only asked if individual i identified that he or she knew individual j personally. Across definitions of network links we compare individual i's response regarding j's gift-giving behavior with individual j's response of i's gift-giving behavior. Since we are primarily interested in mutual gift-giving links, we only include links in which both sides confirm the existence of a historic gift-giving relationship.

To better understand how we construct gift-giving links at the household-level, consider the following example. Consider households A and B, each with one male (M) and one female (F) head/spouse, we consider A and B linked if any one of the four possible reciprocal networks exists between paired individuals: AM-BM, AM-BF, AF-BM, AF-BF. Otherwise, no reciprocal link exists between the two households. Replacing the individual index with a household index, we define g_{ij} as the link between households i and j and impose that $g_{ij} = \max\{g_{i1,j1}, g_{i1,j2}, g_{i2,j1}, g_{i2,j2}\}$ when both household i and j have one head (indexed 1) and one spouse (indexed 2).

It may be helpful to further clarify the difference between the social network mapping and the gift transfer data here. We collected data on actual transfers between households in the gifts module in each of rounds one through five as described in section 1. The gift networks, however, were solicited prior to any of our survey modules and reflect gift-giving patterns prior to the collection of any other data. Recall, the gift modules solicited the identity of the gift recipient/giver depending on whether a gift was listed as a gift received or transferred to another household. We matched the identity of the giver/receiver to our sample IDs when the recipient of the gift was a member of our sample. However, in the gift module, respondents also listed gifts given to other villagers outside of our sample. The data from the gift module were not used to calculate network size. We analyze gifts given within networks and gifts given outside of networks in section 5.

B Theoretical Model

B.1 Formal Model Setup

Formally, we set up the model in section 3 in the following manner:

$$\begin{aligned} U_1^s(U_2^s) = \quad \max_{\tau_s,(U_1^k,U_2^k)_{k=1}^S} \quad & u_1(y_1(s) - \tau_s r_s) - u_1(y_1(s)) \\ & + \gamma_1(r_s)u_2(y_{2b}(s) + \tau_s) - \gamma_1(r_s)u_2(y_{2b}(s)) \\ & + \delta \sum \pi_{sk} U_1^k(U_2^k) \end{aligned} \tag{14}$$

subject to

$$\begin{aligned} \lambda_s\text{:} \quad & u_2(y_{2b}(s) + \tau_s) - u_2(y_{2b}(s)) \\ & + \gamma_2(g, s)u_1(y_1(s) - \tau_s r_s) - \gamma(r_s)u_1(y_1(s)) \\ & + \delta \sum_{k=1}^{S} \pi_{sk} U_2^k \geq U_2^s \end{aligned} \tag{15}$$

$$\delta\pi_{sk}\mu_k\text{:} \quad U_1^k(U_2^k) \geq \underline{U}_1^k = \alpha(k, g_1) \quad \forall k \in S \tag{16}$$

$$\delta\pi_{sk}\phi_k\text{:} \quad U_2^k \geq \underline{U}_2^k = \alpha(k, g_2) \qquad \forall k \in S \tag{17}$$

$$\psi_{1s}\text{:} \quad y_1(s) - \tau_s r_s \geq 0 \tag{18}$$

$$\psi_{2s}\text{:} \quad y_{2b}(s) + \tau_s \geq 0, \tag{19}$$

where π_{sk} represents the transition probability from state s to k. Equation 15 says that transfer and future utility allocations will satisfy the promise-keeping constraint — i.e., future transfers to household 2 in state k will reflect the promises made to maintain future discounted utility contracted in state s. Equations 16 and 17 reflect an individual rationality constraint; allocated utility in any state k must 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 18 and 19 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.

B.2 Model Simulations

For the purposes of the simulation, we use the following utility function for our analysis: $u_i(c) = c^{\frac{1}{2}}$ 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 windfall income is equal to 25 — for the sake of simplicity, all income is either private or public in a given state, Otherwise, income is equal to 5. For example, $y_1(vz) = y_p(vz) + y_b(vz) = 25 + 0 = 25$; $y_1(bz) = y_p(bz) + y_b(bz) = 0 + 25 = 25$; and $y_1(zz) = y_p(zz) + y_b(zz) = 5 + 0 = 5$. We set the altruism weight to $\gamma_i = \frac{1}{2 \cdot (g_i p(s)+1)}$. Transition probabilities reflect probabilities of a household winning the experimental lottery and are $\pi_{zz} = 0.7$, $\pi_{zv} = \pi_{bz} = \pi_{zv} = \pi_{zb} = 0.075$. When a household receives a publicly revealed prize, it will receive gift requests from a majority of network members and $p_1(bz) = p_2(zb) = 0.7$. Otherwise, the household will not receive any gift requests and will focus on its transfer agreement with only the other household, whose decisions are reflected in the simultaneous solution. The cost of maintaining participation in transfers in state s is represented by $\alpha(s) = 2.5 \times g_1 p(s)$. Finally, the discount rate is set to $\delta = 0.85$ for both households.

Without loss of generality, we focus our analysis on household 1's behavior when it wins either the public or private prize and household 2 does not receive a shock to income, states bz and vz respectively. We assume that household 1 does not consider 2's network size in contracts and further does not calculate any social pressure household 2 might incur — this creates asymmetry in the bounds of the contract intervals when we solve the model from household 1's perspective.

Figure B.1 shows the evolution of the optimal (log) contract intervals as network size increases for household 1.As network size increases, household 1 initially demands a larger share of the surplus since both altruistic weight is decreasing and social pressures are increasing — the latter decreases household 1's income and generates deadweight loss. The decrease in both the altruistic weight and effective transferable income shrinks the contract interval $[\underline{\lambda}, \overline{\lambda}]$

FIGURE B.1: Contract Intervals

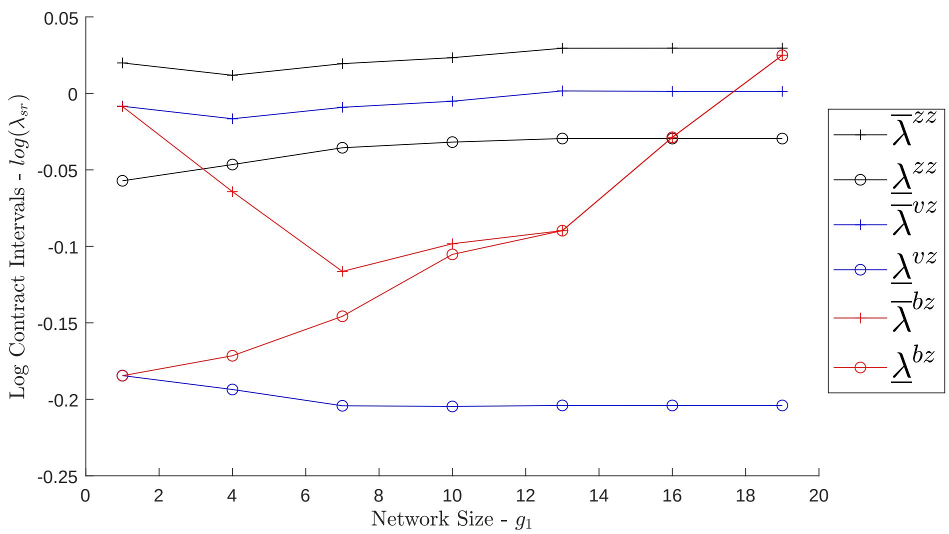

Note: Contract interval solutions as a function of network size. Logged values of λ on the y-axis and network size on x-axis. Contract intervals, $[\underline{\lambda}, \overline{\lambda}]$ in state bz contract as g_1 increases, reflecting the difficulty of negotiating a contract when income is observable and social pressures are large. The "first-best" insurance contract (consumption equal to a stationary share of aggregate output) is only available when network size is less than five.

that household 1 is willing to commit to in terms of transfers to household 2. The upper end of this contract interval is initially decreasing, indicating that household 1 demands a larger share of the surplus than before as network size increases. This is likely due to the fact that household 1 anticipates that the higher degree of social pressure will decrease its consumption in the bz state, which is problematic when marginal returns to own consumption are also increasing (due to the decreasing altruistic weight).

The increasing social pressure to give reflected in larger network size (larger than 7 in our simulations) leads to household 1 giving more of its income to other households in the village collecting on social tax revenues. This leaves household 1 with less income to transfer to household 2. The ratio of the rate of change between smaller effective transferable income and lower altruistic

FIGURE B.2: Discounted Lifetime Expected Utility

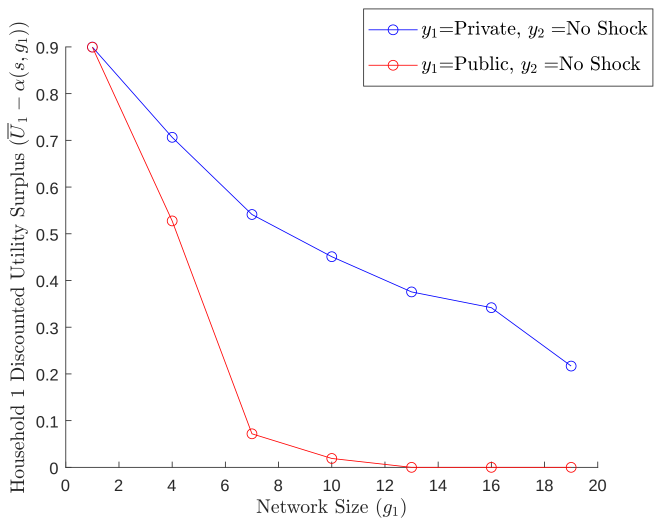

Note: Discounted lifetime expected utility minus the cost of maintaining participation in social transfer and social tax arrangements. Utility modeled for household 1 when the initial state is vz vs. bz and when household 1 takes all available surplus from the transfer arrangement. Utility values in the bz state decrease at faster rates than state vz throughout. The cost of maintaining each network tie, arbitrarily set to $\alpha(g_1) = 2.5 \times g_1 p(s)$ is increasing in network size and intersects with $\overline{U}_1^{bz}$ at a threshold of $g_1 = 12$. Beyond this point, household 1 shuts down all gift transactions when it reaches the bz state.

weight generates an inflection point in the upper bound of the contract interval as network size increases. The important point, however, is that the bounds of the contract interval are always increasing in network size in the bz state due to both the added pressures to give to others as well as the decreasing altruistic weight.

The shut-down point is reached when the contract interval in the bz state is represented by a singleton. At this point, transfers from household 1 to household 2 will equal zero and the shut-down condition takes shape because discounted lifetime utility in the upper bound of the contract interval cannot increase beyond the cost of maintaining participation in informal transfers and norms of social taxation. This is reflected in figure B.2.

C Additional Empirical Extensions and Robustness Checks

C.1 Coordinated Giving

It is possible that the nature of a publicly observable prize allows households to coordinate with one another to distribute prizes across overlapping subnetworks. This would reduce the amount that any one household gives out of publicly observed prizes, which would be observationally equivalent to the shutdown hypothesis if the likelihood of overlapping subnetworks is increasing in one's network size, or in networks with a higher degree of support as in Jackson, Rodriguez-Barraquer and Tan (2012). To provide some intuition for this possibility, suppose households A, B and C are connected to one another through a gift network and households A and C receive publicly observable income shocks. Households A and C should optimally coordinate with one another to each give less to household B than they would if they could not coordinate. If this were true, then the inclusion of the average network winnings of the second degree network in our specification should nullify the negative coefficient on the network size interaction in Table 2.

To test this alternative hypothesis, we construct the average network winnings of the second degree network by averaging across the adjusted *network* winnings of household i's gift network (removing household i's contribution to this average), labeled $\overline{\text{Public}}^{2'}_{it}$, and include it and its $\overline{\text{Private}}^{2'}_{it}$ counterpart as control variables in our specification containing the test of the shutdown hypothesis. If households coordinate giving out of public winnings, the negative coefficient on β_{bg} should no longer be significant and would be replaced by a negative coefficient in front of $\overline{\text{Public}}^{2'}_{it}$.

Table D.14 presents results that confirm the shutdown hypothesis with the number of gifts given as the dependent variable.[39] First, the effect outlined in the above paragraph does not take place. Instead, the negative sign of

[39]The results are also consistent with our prior results for the other dependent variables in our paper: total value of gifts given, and average value of gifts given. These estimates are available upon request.

β_{bg} is maintained and the coefficient in front of $\overline{\text{Public}}_{it}^{2'}$ is generally positive, which runs counter to the idea that households coordinate with one another to reduce giving. Column 2 examines whether public lottery winners change giving patterns when the second-degree network also wins more in public than average. It could be that households with large gift networks are the only ones who coordinate around giving, so the negative coefficient we would expect out of this hypothesis would only be present if we interact $\overline{\text{Public}}_{it}^{2'}$ with network size. Column 3 shows this is not the case. Instead, there is a positive and significant effect again. The result is consistent with the first and second column. Households who win the public lottery are more likely to give when the second degree network wins the public lottery. This is opposite to the result we would expect under the coordination hypothesis.

These results counter the notion that households coordinate giving with one another when an income shock is publicly observable to the village and reinforce our findings of a shutdown hypothesis. This result is explainable through a tweak of our analytical model. Consider a network setting in which the probability that household j requests a gift from household i decreases when another household in j's network, k, receives a publicly observable windfall. This decrease in household i's probability of receiving a gift request decreases the probability of receiving gift requests beyond the shut-down threshold, reducing the negative effect of network size on gift giving.[40]

C.2 Mental Accounting of Stochastic Unearned Income

One might believe that mental accounting over stochastic unearned income generates different decision making patterns relative to earned income (Thaler, 1999). This could be true, but would not explain our primary results: that people behave differently when stochastic unearned income is publicly versus privately distributed, and that network size matters to gift giving differentially depending upon the public observability of the income shock. Furthermore,

[40]It is unlikely that the shutdown effect arises due to a "bystander effect" wherein the lack of coordination among households results in a type of free-riding behavior in which no one gives to anyone. If this were the case, we would not observe the positive coefficient estimates on the interaction terms with $\overline{\text{Public}}_{it}^{2'}$ in columns 2 and 3 of Table D.14.

cash transfers as a policy device are by definition unearned, so the analysis of behavior following unearned income shocks is important in policy terms.

C.3 Opting Out of Informal Insurance

Models of limited commitment, endogenously enforceable informal insurance contracts suggest that households will choose to opt out of the contract if and when they receive a large enough income shock to make exit and default on one's contractual obligation preferable to payment and remaining in the arrangement. Thus, households receiving large income shocks may "shut down" due to this mechanism as opposed to the mechanism in our model. To test among these two explanations, we replicate Table 2 and divide windfall lottery winnings into "small" (less than or equal to 35 GH¢) and "large" (greater than 35 GH¢). In table D.15 we show that the shut-down effect (negative β_{bg}) remains for both small and large public windfall gains. Default as a result of unusually large windfall gains is unlikely to be a driver of the shut down effect.

C.4 Precautionary Savings and Investments in Others

Another prospective motive for giving out of private winnings is to increase one's savings by transferring cash to sympathetic friends in the form of interest-free loans — this could either be viewed as a callable deposit that can be withdrawn in future periods or as an investment in relatively productive households. In either case, it seems irrational to target gifts out of private winnings to those with the highest marginal propensity to consume. Such households are unlikely to have sufficient supply of liquid assets to give to their friends when called upon. Similarly, they are either unlikely to be among the relatively more productive households in the village or they are unlikely to use such transfers to invest in productive activities. Households looking to invest in others for their own future gain will target households of moderate or better existing wealth (Santos and Barrett, 2011).

C.5 Test of Full Risk Pooling

We test whether social networks also serve the informal insurance purpose of smoothing members' consumption by distributing income shocks across the network. The familiar full-risk-pooling prediction, following Townsend (1994), is that the intertemporal 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 full risk-pooling hypothesis null is 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. Given that within our model inter-household transfers serve multiple purposes beyond merely informal insurance, we expect to reject the full-risk-pooling null in favor of the one-sided alternate hypothesis that the coefficient is less than one. We likewise expect to reject the no-risk-pooling null that change in consumption is uncorrelated, in favor of the one-sided alternate hypothesis that they are positively correlated, reflecting that transfers serve in part as (incomplete) insurance. The incompleteness of the informal insurance occurs because of the shutdown hypothesis and because altruistic households will not share private winnings with networks members who do not exhibit great material need. The social solidarity network fulfills some insurance function, but incompletely, in part because it also serves members' altruistic objectives and because excessive social taxation pressures can induce optimal defection. Notice that incomplete insurance in this setting need not arise due to more familiar mechanisms of moral hazard, limited commitment or hidden income (Cole and Kocherlakota (2001),Ligon, Thomas and Worrall (2002),Dubois, Jullien and Magnac (2008), Kinnan (2019)).

Table D.16 reports results of those hypothesis tests. We show that limited risk pooling occurs within the full gift network and the family-only network in columns 1 and 2, respectively. The respective point estimates of 0.31 and 0.33 are statistically significantly greater than 0 but also statistically significantly less than 1. However, when we exclude family members from the gift network (column 3) we cannot reject the zero risk pooling null (and strongly reject the full risk pooling null). These results combined with those from Table 5 strongly suggest that gifts to village friends - rather than to family - appear driven primarily by altruistic motives while transfers to family are more consistent with a pure insurance motive. Columns 4 through 6 look at three more combinations of gift vs. family networks and conclude that the network with the highest degree of insurance-related sharing corresponds to those networks that include family members with whom one has a prior gift exchange relationship.

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, Bramoullé and Perez-Richet (2017). From this result, we conclude that inter-household gift networks are multi-functional. They may include limited risk pooling, especially among family, but likely also involve altruistic solidarity among network ties, especially non-family members within the village.

D Appendix Tables

TABLE D.1: Detailed Gift Data

	Gifts Given			Gifts Received		
	HH Average		Total	HH Average		Total
	Value	Number	Number	Value	Number	Number
All gift variables						
All gifts	17.27	1.76	2,791	35.39	1.90	3,006
Gifts to/from spouse	1.38	0.08	119	1.73	0.08	119
Gifts to/from children	1.75	0.11	168	9.01	0.35	563
Gifts outside of village	10.07	0.68	1,082	28.16	1.16	1,844
In kind gifts	5.27	0.93	1,470	9.96	0.82	1,305
Used in Analysis						
Cash gifts given to fellow villagers	3.34	0.41	652	1.97	0.25	404
All gifts to fellow villagers	5.60	0.97	1,542	4.07	0.59	935

Note: This table provides a description of the data from the gift module from the household survey. The survey asked respondents to indicate both the number of "monetary" and "non-monetary" (or in-kind) gifts "you have given anyone" (first three columns) or "anyone has given you" (latter three columns) "which they/you do not need to repay" in the last two months. For each gift, the respondent indicated the relationship to the recipient/giver and the gift's value in GH¢(estimated for in-kind gifts). HH Average indicates the household-level average response for the total value of gifts and the number of gifts in each round. The column indicating "Total Number" reflects the total number of gifts given/received across all rounds in the row-specified category of gifts.

Table D.2: Household Summary Statistics

				Percentile	
	N	Mean	Sd	5th	95th
HH size	318	6.64	2.64	3	11
Cash Gifts Given (last 2 months):					
Number	1,586	0.41	0.86	0	2
Value GH¢ (Total Given)	1,586	3.34	18.16	0	15
Value GH¢ (Conditional on Giving)	404	13.12	34.19	1	50
Cash Gifts Received (last 2 months):					
Number	1,586	0.25	0.71	0	2
Value GH¢ (Total Received)	1,586	1.97	12.08	0	10
Value GH¢ (Conditional on Receiving)	265	11.77	27.57	1	31
Own Lottery Winnings (GH¢):					
Value of Private Cash Prize	1,251	2.35	10.52	0	20
Value of Public Cash Prize	1,251	2.29	10.45	0	10

Note: HH size is fixed over the year in which data is collected, other values vary over the five rounds of data collection. Total value of all gifts given/received are reported conditional on giving or receiving a gift. Cash prizes are distributed prior to each of rounds two through five, so round one observations are not included here. In the analysis, we impose a value of zero on these variables in round one.

Table D.3: Individual Summary Statistics

Panel A: Individual Summary Statistics

				Percentile	
	N	Mean	Sd	5 p-tile	95 p-tile
HH Size	610	5.81	2.65	2.00	11.00
Gift Network Size	605	9.87	10.13	0.00	31.00
Cash Gifts Given (last 2 months):					
N gifts given	3,006	0.22	0.60	0.00	1.00
Total value of all gifts given	3,006	1.76	13.02	0.00	8.00
Food Consumption (last month):					
PC Food Consumption	3,005	15.54	18.53	0.00	43.20
PC Purchased Food	3,005	11.71	17.48	0.00	39.65
PC Home-produced Food	3,005	3.83	6.12	0.00	15.00

Panel B: Intrahousehold Differences

	HH Head		Spouse		
	Mean	N	Mean	N	P-Value
Gift Network Size	11.47	288	8.42	317	0.00
Cash Gifts Given (last 2 months):					
N gifts given	0.30	1,439	0.14	1,567	0.00
Total value of all gifts given	2.87	1,439	0.75	1,567	0.00
Food Consumption (last month):					
PC Food Consumption	11.93	1,439	18.86	1,566	0.00
PC Purchased Food	5.37	1,439	17.53	1,566	0.00
PC Home-produced Food	6.55	1,439	1.32	1,566	0.00

Note: HH size is fixed over the year in which data is collected, other values vary over the five rounds of data collection. P-value reflects the difference in the household head and spouse responses to variables in panel B. "Head" and "Spouse" averages do not sum to the household average since a subset of households have heads with multiple spouses or are single-headed.

TABLE D.4: Balance Tests Along Baseline Household Statistics

	Did Not Win		**Won Lottery**		
Baseline (R1) Statistics	N	Mean	N	Mean	P-Value
HH size	181	6.52	137	6.81	0.33
Gift Network Size	181	11.69	137	10.80	0.44
Cash Gifts Given (last 2 months):					
N gifts given	181	0.54	136	0.38	0.14
Total value of all gifts given	181	4.47	136	5.63	0.67
Food Consumption (Last Month):					
PC Food	174	31.11	124	30.09	0.76
PC Purchased Food	174	24.16	124	23.28	0.79
Test of Joint Significance (P-Value of F Statistic)					0.53

Note: Balance test of round one ("Baseline") observations. "Won Lottery" reflects statistics for households who won the cash lottery at any point during rounds two through five. "Did Not Win" reflects statistics for households who never won a cash lottery. P-Value (t-test) associated with the null hypothesis of equal mean at baseline across the two categories. When "Won Any" is replaced with households who only won the private (or public) lottery, balance tests maintain insignificant differences between winners and non-winners.

Table D.5: Prize Winnings Influence Gift-Giving - Count Data

		(1) Number
Randomized Explanatory Variables		
Value of Private Cash Prize	β_v	0.154***
		(0.045)
Value of Public Cash Prize	β_b	0.0919**
		(0.046)
Household FE		Yes
Round × Village FE		Yes
Observations		1,586

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals number of cash gifts given in hh. Value of Private/Public Cash prize is divided by $10 \in \{0, 1, 2, 3.5, 5, 7\}$. Estimated using Poisson estimator with hh and Round × Village FE.

Table D.6: Prize Winnings and Gift Giving - Including Livestock Winnings

		Gifts-Given		
Dependent Variable:		Value (Total) (1)	Value (Average) (2)	Number (3)
Randomized Explanatory Variables:				
Value of Private Cash Prize	β_v	0.224***	0.175***	0.241***
		(0.078)	(0.063)	(0.065)
Value of Public Cash Prize	β_b	0.111	0.0504	0.127*
		(0.087)	(0.070)	(0.072)
Value of Private Livestock Prize		0.0399	0.0138	0.0661
		(0.085)	(0.068)	(0.070)
Value of Public Livestock Prize		-0.0278	-0.0123	-0.0501
		(0.099)	(0.079)	(0.082)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
Left-censored Obs.		1,182	1,182	1,182
Observations		1,586	1,586	1,586

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. The dependent variable equals log total value of cash gifts given in hh in column 1; log average value of cash gift given in column 2; number of gifts given in column 3. Value of Private/Public Cash or Livestock prize is divided by 10 $\in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns with a lower bound of zero.

TABLE D.7: Prize Winnings and Gift Giving - Including In-kind Gifts in Dependent Variable

		Gifts-Given					
		Outliers Included			Outliers Excluded		
Dependent Variable:		Value (Total) (1)	Value (Average) (2)	Number (3)	Value (Total) (4)	Value (Average) (5)	Number (6)
Value of Private Cash Prize	β_v	0.085	0.052	0.181**	0.103*	0.062	0.202***
		(0.061)	(0.041)	(0.074)	(0.056)	(0.039)	(0.076)
Value of Public Cash Prize	β_b	0.046	-0.001	0.082	-0.005	-0.007	0.014
		(0.060)	(0.041)	(0.073)	(0.058)	(0.040)	(0.078)
Value of Private Livestock Prize		0.000	-0.001	0.051	-0.055	-0.038	-0.037
		(0.064)	(0.043)	(0.078)	(0.061)	(0.042)	(0.082)
Value of Public Livestock Prize		-0.100	-0.058	-0.123	-0.100*	-0.047	-0.157*
		(0.066)	(0.045)	(0.079)	(0.060)	(0.041)	(0.080)
Left-censored Obs.		914	914	914	894	894	894
Observations		1,586	1,586	1,586	1,431	1,431	1,431

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. In this table we combine all cash and in-kind gifts in the dependent variable. As in-kind gifts increase in value, households find their value harder to estimate, often by over-valuing in-kind gifts and making the measurement imprecise. In columns 1 through 3 we combine all cash and in-kind gifts regardless of estimated value. In columns 4 through 6 we exclude total values of gifts given that exceed the 95th percentile of gift-value (both cash and in-kind). The dependent variable equals log total value (estimated for in-kind gifts) of gifts given in hh in columns 1 and 4; log average value of gift given in columns 2 and 5; number of gifts given in columns 3 and 6. Value of Private/Public Cash prize is divided by $10 \in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns with a lower bound of zero. "Number of Gifts" is most precisely measured and coefficient is consistently positive in private cash prize.

Table D.8: Testing the Hidden Income Model

Dependent Variable:	Per Capita Consumption	
	(1)	(2)
Lagged Total Consumption		
Per Capita Total$_{t-1}$	0.686***	0.427**
	(0.030)	(0.174)
Average(Per Capita Total$_{t-1}$)$_{-i}$	0.056	
	(0.077)	
Randomized Explanatory Variables		
Value of Private Cash Prize	0.013	0.010
	(0.015)	(0.013)
Value of Public Cash Prize	0.034***	0.034***
	(0.011)	(0.011)
Lagged Randomized Explanatory Variables		
Value of Private Cash Prize (1 Period Lag)	-0.006	-0.010
	(0.010)	(0.010)
Value of Public Cash Prize (1 Period Lag)	0.015	0.007
	(0.013)	(0.012)
Village × Round FE	No	Yes
Village × Round FE × Food($t-1$)	No	No
Village × Round FE × Cons($t-1$)	No	Yes
R^2	0.47	0.54
Observations	1,205	1,205

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. The dependent variable equals log per capita household total consumption in columns 1 and 2. OLS estimator used throughout and standard errors (in parentheses) are clustered by household. In column 1 we use the mean of lagged average consumption in favor of Village× Round fixed effects — these are used in column 2. The insignificant and small (negative) effect on lagged private winnings indicates households are unlikely to strategically hide income in to finance future consumption in our setting.

Table D.9: Additional Household Summary Statistics

				Percentile	
	N	Mean	Sd	5th	95th
Network Size:					
Gift Network Size	318	11.30	10.08	0	32
Food Consumption (last month, GH¢):					
PC Food	1,462	24.20	17.54	7.43	52.88
PC Purchased Food	1,462	18.14	16.59	3.75	45.20
Network Average Lottery Winnings (GH¢):					
Adjusted Average Value (Private)	1,272	0.20	1.20	0	0.63
Adjusted Average Value (Public)	1,272	0.20	1.10	0	0.74

Note: Networks were collected prior to baseline making network size fixed over the year in which data is collected, other values vary over the five rounds of data collection. Per capita (PC) food consumption per household sums all food purchases by the head of household or the spouse and divides by household size. If either was not present for a particular round of the survey, then we report the variable as missing for the household during that round.

TABLE D.10: Quantile Regression Estimates

Dependent Variable:	Log PC Food Consumption				
Quantile:	1st	12th	25th	50th	75th
Adjusted Network Private ($\overline{\text{Private}}'_{it}$)	3.998***	2.178**	1.725**	0.671	0.804
	(1.047)	(0.874)	(0.672)	(0.944)	(0.750)
Adjusted Network Public ($\overline{\text{Public}}'_{it}$)	-0.669	-0.707	0.085	-0.081	-0.367
	(2.094)	(0.829)	(0.746)	(0.953)	(0.740)
Value of Private Cash Prize	0.111*	-0.009	-0.034	-0.026	-0.019
	(0.063)	(0.043)	(0.033)	(0.043)	(0.025)
Value of Public Cash Prize	0.102**	0.053*	0.032	0.034	-0.008
	(0.045)	(0.031)	(0.026)	(0.026)	(0.022)
Round × Village FE	Yes				
Observations	594				

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Simultaneous quantile regression bootstrapped 1,000 times at 1st, 12th, 25th, 50th and 75th quantiles. Dependent variable is log total per capita food consumption in household over the last month. Log transformations of network averages. We limit analysis to observations of households surveyed during the "hungry" season (see Figure E.4). Coefficient estimates represented graphically in Figure 3.

Table D.11: Predictors of Gift Network Size

Dependent Variable:	N in Gift Network (1)	N in Gift Network (2)
HH Size (Present in Village)	-0.116	
	(0.313)	
HH Members Living Away from Village	-0.121	
	(0.170)	
Adult HH Members (Present)	0.768*	
	(0.391)	
N of Non-Coresident Family Members in Sample	0.419***	0.427***
	(0.030)	(0.031)
Share of HH Members aged 5 to 18 who attend school	2.477***	
	(0.928)	
Log Per-Capita HH Food Consumption	0.853	
	(1.984)	
Observations	318	318
R-squared	0.48	0.46

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals number of households in gift network (as defined in section 4) in round 1. Coefficients estimated using OLS. Robust standard errors in parentheses.

Table D.12: Testing the Shut Down Hypothesis — Network Defined by Number of Family Members

Dependent Variable:		Gift Giving Value(Total) (1)	Gift Giving Value (Average) (2)	Gift Giving Number (3)
Randomized Explanatory Variables				
Value of Private Cash Prize	$\beta_v > 0$	0.226*	0.217**	0.226**
		(0.117)	(0.093)	(0.097)
Value of Private Cash Prize × N_{FAM}	$\beta_{vg} \leq 0$	-0.000	-0.002	0.000
		(0.005)	(0.004)	(0.004)
Value of Public Cash Prize	$\beta_b > 0$	0.372***	0.183*	0.431***
		(0.133)	(0.105)	(0.112)
Value of Public Cash Prize × N_{FAM}	$\beta_{bg} < 0$	-0.020**	-0.010	-0.024***
		(0.008)	(0.006)	(0.007)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v = \beta_b$		0.4	0.81	0.16
P-Value $H_0 : \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		0.74	0.54	0.48
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.70	0.28	0.74
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.07	0.07	0.02
$\hat{N}$ at Shut Down		18.85	18.58	18.33
Left-censored Observations		1,182	1,182	1,182
Observations		1,586	1,586	1,586

Note: $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals log total value of cash gifts given in household in column 1; log average value of cash gift given in column 2; number of gifts given in column 3. Value of Private/Public Cash prize is divided by $10 = \in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns. Null hypotheses are tested using Wald tests of equivalence specified for network size (N_{FAM}) of 0, 5, 10 and 20. P-values reported under each column for each of the hypotheses. N_{FAM} denotes network size — network definition equals number of non co-resident adult family members residing in the village. N_{FAM} at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$.

Table D.13: Testing the Shut Down Hypothesis — Network Defined by Predicted Gift-Network Size

		Gift Giving		
Dependent Variable:		Value(Total) (1)	Value (Average) (2)	Number (3)
Randomized Explanatory Variables				
Value of Private Cash Prize	$\beta_v > 0$	0.258*	0.254**	0.240**
		(0.145)	(0.116)	(0.120)
Value of Private Cash Prize $\times$ $\hat{N}$	$\beta_{vg} \leq 0$	-0.003	-0.007	-0.01
		(0.010)	(0.008)	(0.008)
Value of Public Cash Prize	$\beta_b > 0$	0.418**	0.170	0.548***
		(0.173)	(0.135)	(0.150)
Value of Public Cash Prize $\times$ $\hat{N}$	$\beta_{bg} < 0$	-0.034**	-0.013	-0.048***
		(0.017)	(0.013)	(0.016)
Household FE		Yes	Yes	Yes
Round $\times$ Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v = \beta_b$		0.47	0.64	0.11
P-Value $H_0 : \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		0.97	0.33	0.57
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.21	0.12	0.11
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.05	0.24	0.00
$\hat{N}$ at Shut Down		12.14	12.73	11.43
Left-censored Observations		1,182	1,182	1,182
Observations		1,586	1,586	1,586

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals log total value of cash gifts given in household in column 1; log average value of cash gift given in column 2; number of gifts given in column 3. Value of Private/Public Cash prize is divided by $10 = \in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns. Null hypotheses are tested using Wald tests of equivalence specified for network size ($\hat{N}$) of 0, 5, 10 and 20. P-values reported under each column for each of the hypotheses. $\hat{N}$ denotes network size — network definition equals predicted network size (first stage reflected in table D.11). $\hat{N}$ at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$.

Table D.14: Testing the Coordination Hypothesis

		Number of Gifts Given		
		(1)	(2)	(3)
Own Winnings and Network Size Interaction Variables				
Private (Own)	$\beta_v > 0$	0.247**	0.289**	0.257**
		(0.108)	(0.112)	(0.108)
Private (Own) × N	$\beta_{vg} \leq 0$	-0.002	-0.001	-0.003
		(0.006)	(0.006)	(0.006)
Public (Own)	$\beta_b > 0$	0.458***	0.512***	0.450***
		(0.110)	(0.116)	(0.110)
Public (Own) × N	$\beta_{bg} < 0$	-0.041***	-0.038***	-0.039***
		(0.011)	(0.013)	(0.011)
Second Degree Adjusted Network Average Winnings ($\overline{\textbf{Private/Public}}_{it}^{2'}$)				
$\overline{\text{Private}}_{it}^{2'}$		5.349	6.263*	8.166*
		(3.320)	(3.394)	(4.916)
$\overline{\text{Public}}_{it}^{2'}$		6.765**	7.886***	-0.509
		(2.821)	(2.888)	(4.264)
$\overline{\text{Private}}_{it}^{2'}$ × Private (Own)			-2.977	
			(2.571)	
$\overline{\text{Public}}_{it}^{2'}$ × Public (Own)			-5.942	
			(9.608)	
$\overline{\text{Private}}_{it}^{2'}$ × N				-0.318
				(0.435)
$\overline{\text{Public}}_{it}^{2'}$ × N				0.794**
				(0.355)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.09	0.31	0.11
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.00	0.03	0.00
Left-censored Obs.		1,172	1,172	1,172
Observations		1,573	1,573	1,573

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals number of gifts given. Private/Public (Own) equals the value of the private/public cash prize won by the HH divided by 10. $\overline{Private}_{it}^{2'}$ reflects the private cash winnings of the second degree network. Tobit estimator used in all columns with a lower bound of zero. Wald tests of equivalence specified for network size (N) of 10 and 20.

Table D.15: Default Following Large Windfall Gain Only?

		Gift Giving		
Dependent Variable:		Value(Total) (1)	Value (Average) (2)	Number (3)
Private Windfall				
Private-Small	$\beta_{v,small} > 0$	0.660**	0.466*	0.504**
		(0.304)	(0.244)	(0.251)
Private-Small × N	$\beta_{vg,small} \leq 0$	-0.043*	-0.026	-0.031
		(0.023)	(0.018)	(0.019)
Private-Large	$\beta_{v,large} > 0$	0.245*	0.230**	0.241**
		(0.144)	(0.116)	(0.120)
Private-Large × N	$\beta_{vg,large} \leq 0$	-0.001	-0.004	0.000
		(0.008)	(0.007)	(0.007)
Public Windfall				
Public-Small	$\beta_{b,small} > 0$	0.451*	0.256	0.448**
		(0.251)	(0.202)	(0.206)
Public-Small × N	$\beta_{bg,small} \leq 0$	-0.025	-0.013	-0.030**
		(0.018)	(0.015)	(0.015)
Public-Large	$\beta_{b,large} > 0$	0.454***	0.190	0.531***
		(0.157)	(0.122)	(0.131)
Public-Large × N	$\beta_{bg,large} \leq 0$	-0.053***	-0.022*	-0.059***
		(0.019)	(0.013)	(0.016)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
Left-censored Obs.		1,182	1,182	1,182
Observations		1,586	1,586	1,586

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals log total value of cash gifts given in hh in column 1; log average value of cash gift given in column 2; number of gifts given in column 3. Value of Private/Public Cash prize is divided by $10 = \in \{0, 1, 2, 3.5, 5, 7\}$. Small values less than or equal to 3.5; large values greater than 3.5. Tobit estimator used in all columns with a lower bound of zero.

TABLE D.16: Tests of Risk-Sharing

Dependent Variable:	Δlog (PC Food)					
	G	F	G $\notin$ F	F $\notin$ G	G $\cap$ F	$\notin$ (G $\cup$ F)
	(1)	(2)	(3)	(4)	(5)	(6)
First Difference of Network Average Per Capita Food Consumption						
Δlog(Network PC Food)$_{it}$	0.306***	0.328***	0.102	0.034	0.257***	0.022
	(0.087)	(0.098)	(0.077)	(0.063)	(0.078)	(0.224)
Randomized Explanatory Variables						
Value of Private Cash Prize	-0.001	0.011	0.002	0.013	0.002	0.007
	(0.010)	(0.015)	(0.011)	(0.014)	(0.010)	(0.013)
Value of Public Cash Prize	0.006	0.007	0.014	0.004	0.008	0.004
	(0.012)	(0.011)	(0.013)	(0.011)	(0.013)	(0.011)
$\overline{\text{Private}\|\text{Network}}_{it}$	0.005	0.057	-0.012	0.025	0.014	-0.320**
	(0.027)	(0.043)	(0.030)	(0.021)	(0.023)	(0.156)
$\overline{\text{Public}\|\text{Network}}_{it}$	-0.006	-0.001	0.016	0.006	-0.038	-0.077
	(0.032)	(0.021)	(0.022)	(0.019)	(0.031)	(0.175)
Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Network Definition						
Gift Network	Yes	—	Yes	No	Yes	No
Family Network	—	Yes	No	Yes	Yes	No
Left-censored Obs.	265	268	233	263	245	303
Observations	969	979	844	961	897	1,107

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable equals change in log per capita food consumption in household from round t to $t-1$. Estimated using OLS. Standard errors clustered by household. Each column analyzes a different network: 1) Reciprocal gift network, 2) Family (including extended) network, 3) Reciprocal gift links that are not family members 4) Family members that are not reciprocal gift links 5) Reciprocal gift links that are family members and 6) Neither in family nor gift network. We drop observations when the specified network contains zero links. We reject full insurance across all specifications and observe the highest degree of insurance motives in family networks. This suggests that gift-giving among friends follows mainly from altruistic motives and gift-giving among family mixes altruistic and insurance motives.

E Appendix Figures

Figure E.1: Distribution of Gift Network Size

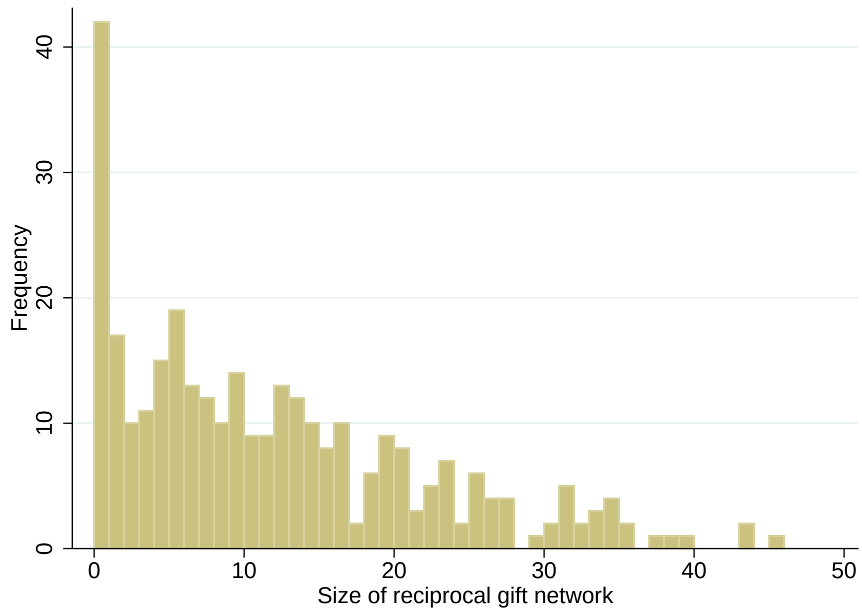

Note: Each bin is one unit wide.

Figure E.2: Shut-down Hypothesis on Number of Gifts Given

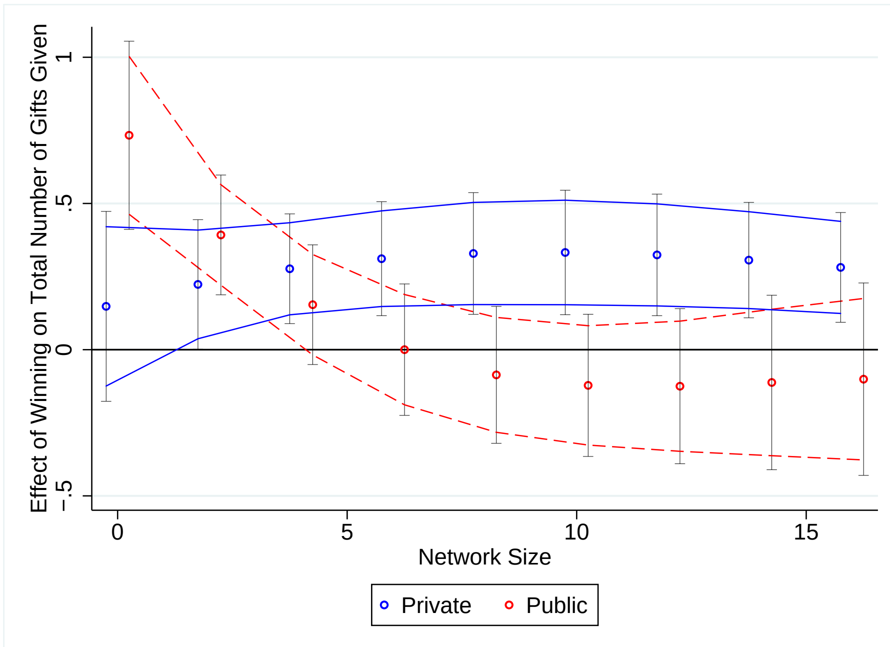

Note: Dependent variable equals number of gifts given. Estimation of Model 11 with the inclusion of 2nd and 3rd order polynomial interactions on network-size variable (with respective coefficients $\beta_{bg^2}, \beta_{vg^2}, \beta_{vg^3}$ and β_{bg^3}). Dots represent point estimates of $\beta_b + \beta_{bg} \times N + \beta_{bg^2} \times N^2 + \beta_{bg^3} \times N^3$ (repeat for private, β_v). Blue line represents 90% confidence interval for linear combination of private coefficients; dotted red line represents the 90% confidence interval for linear combintation of public coefficients. Bars represent 95% confidence intervals. Plots of public coefficients offset by one for ease of viewing.

Figure E.3: Effect of Unadjusted Network Winnings on Food Consumption by Quantile

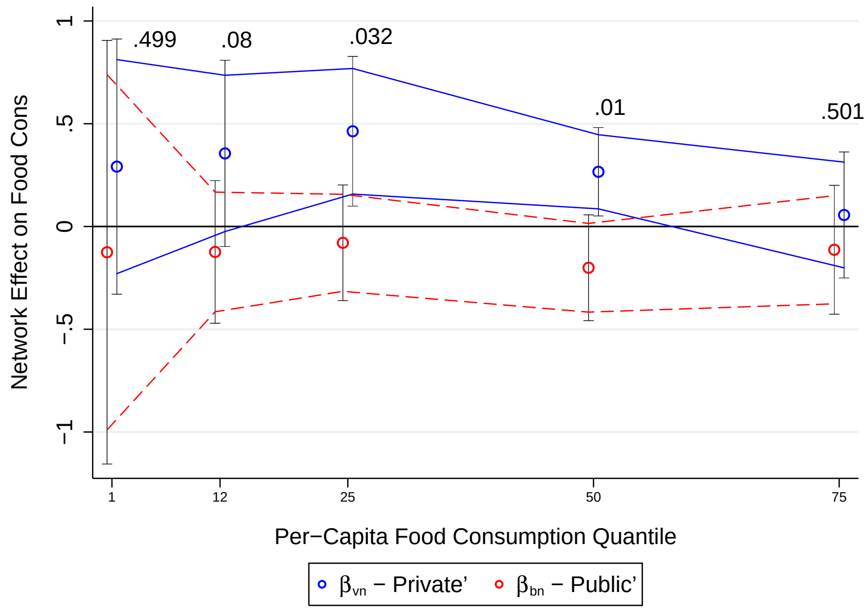

Note: Results of a simultaneous quantile regression at 1st, 12.5th, 25th, 50th, and 75th quantiles bootstrapped over 1,000 iterations. Dependent variable is log home-produced per capita food consumption over the last month. Quantiles represented on the x axis. Blue dots (lines) show the coefficient estimates (90% confidence interval) on private network winnings, $\overline{\text{Private}}_{it}$, at each quantile. Red represents public network winnings, $\overline{\text{Public}}_{it}$. The numbers above each point represent the quantile specific Wald test of $\text{H}_0: \ \beta_{vn} = \beta_{bn}$.

Figure E.4: Home Produced Food Over The Course of the Year

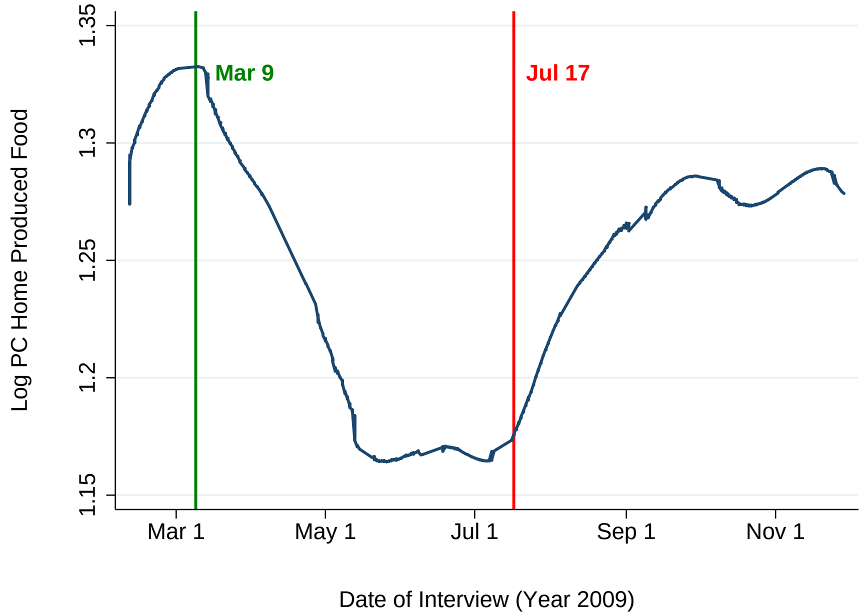

Note: Log home-produced per capita food consumption over the last month on the y axis. Date of interview on the x axis. Blue line shows the lowess smoothed curve by date with a bandwidth of 0.4. The peak of the average home produced food consumption is around March 9. After this point, average home produced food consumption begins to decrease until its nadir on around July 17. We include all observations between the vertical green line and vertical red line in our quantile regression analysis in Section 4. Households with negligible per capita home food production (N=46) of between GH¢0 and 1.5 are excluded from the calculations in this graph in order to gain a clearer understanding of home-produced food availability over the course of the year.