This version, November, 2019

Altruism, Insurance, And Costly Solidarity Commitments[1]

Abstract. Inter-household transfers play a central role in village economies. Whether understood as informal insurance, credit, or social taxation, the dominant conceptual models used to explain transfers rest on a foundation of self-interested dynamic behavior. Using experimental data from households in rural Ghana, where we randomized private and publicly observable cash payouts repeated every other month for a year, we reject two core predictions of the dominant models. We then add impure altruism and social taxation to a model of limited commitment informal insurance networks. The data support this new model's predictions, including that unobservable income shocks may facilitate altruistic giving that better targets less-well-off individuals within one's network, and that too large a network can overwhelm even an altruistic agent, inducing her to cease giving.

Vesall Nourani
Department of Economics
Massachusetts Institute of Technology

Christopher B. Barrett
Charles H. Dyson School
Cornell University

Eleonora Patacchini
Department of Economics
Cornell University

Thomas Walker
World Bank

[1]Corresponding author: vnourani@mit.edu. Harold Alderman, Andrew Foster, Susie Godlonton, Annemie Maertens, Liz Tennant, and seminar participants at Cornell University, Einaudi Institute for Economics and Finance, London School of Economics, the University of Copenhagen, Massachusetts Institute of Technology, Williams College, 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 Graduate Research Fellowship under Grant No. DGE-1144153 and the African Development Bank Structural Transformation of African Agriculture and Rural Spaces project supported by the Korean government through the Korea-Africa Economic Cooperation Trust Fund.

1. INTRODUCTION

Social solidarity networks have long been understood to play a central role in village economies. There can be both altruistic and self-interested drivers behind such networks' functioning (Ligon and Schechter (2012)). Although the possibility of altruism has been accommodated in some work within that literature (notably Foster and Rosenzweig (2001)), at least since Popkin (1979) and Posner (1980), the dominant framework for social scientists' understanding of transfers within social networks has rested on self-interested dynamic behavior, most commonly framed as self-enforcing informal insurance contracts (Ambrus, Mobius, and Szeidl, 2014; Coate and Ravallion, 1993; Fafchamps, 1992; Townsend, 1994). This contractual framing helps explain risk pooling among households whose risk averse preferences drive them to seek to smooth consumption (Coate and Ravallion, 1993; Townsend, 1994). An important policy implication of this framework is 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 recipients within their network to share windfall gains arising from public (or charitable) transfers (Angelucci and De Giorgi, 2009).

A related but distinct literature emphasizes a dark side of self-interested sharing within social networks. Social pressures — often referred to as 'social taxation' — can place significant demands on those who enjoy income growth, discouraging investment and potentially even trapping households in poverty. A range of studies find strong empirical evidence supporting the existence of social taxation (Goldberg, 2017; Jakiela and Ozier, 2016; Platteau, 2000; Sen and Hoff, 2006; Squires, 2017). This contrary perspective raises important questions about prospective limits to the value of extensive social networks.

The informal insurance and social taxation literatures both depend fundamentally on the observability of income, or at least of shocks to income.[2] 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, social taxation can only apply to the observable portion of others' income streams that is observable, and thus taxable by one's social network.

[2] More specifically, they rely on non-uniform shocks across households within the network so that exogenous change in incomes triggers the redistributive mechanism implied by informal insurance, social taxation, or both. We use the term 'income shock' to imply non-uniform shocks.

In this paper we take a step towards reconciling and transcending these literatures. We start by reconsidering whether dynamic self-interest suffices to explain observed patterns of inter-household transfers. Both of the above frameworks' dependence on public observability of income shocks implies two testable hypotheses. First, publicly observable income shocks should lead to inter-household transfers, whether due to social taxation, informal insurance contracts, or both. One should be able to reject the null that public income shocks have no effect on inter-household transfers in favor of the one-sided alternate hypothesis of positive impacts. Second, unobservable income shocks — in particular, positive private income shocks that a purely self-interested beneficiary would never divulge — should not prompt inter-household transfers. Failure to reject the null that private income shocks have zero effect is a low-power test of the public observability hypothesis. To date, however, we are unaware of any research that empirically tests the implications of publicly observable (hereafter 'public') versus unobservable (i.e., 'private') income shocks on inter-household transfers.

We show that neither prediction 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. Regressions of giving within subjects' social networks as a function of exogenous (randomized) private and public winnings clearly fail to reject the public income shocks null but do reject the private income shocks null. We corroborate those findings with regressions of how subjects' consumption varies with winnings within one's network and with dyadic regressions reflecting the flows between any two subjects. These findings imply rejection of the framing of inter-household transfers as solely a result of self-interested informal insurance contracting or of social taxation. At a broader level, these results offer a more encouraging statement about human nature; people may act out of self-interest, but they likewise behave in ways that clearly reflect more noble, other-regarding interests.

Those empirical results imply a need to refine our theoretical understanding of inter-household transfers. We adapt the canonical dynamic model of self-enforcing insurance contracting to introduce an altruistic motive for households to give to others, following Foster and Rosenzweig (2001). Adding an altruistic component to preferences directly addresses the second hypothesis above, explaining why people might give from private income windfalls, while still allowing for self-interested behavior. People can be both selfish and altruistic.

Our model includes two key refinements, however, reflecting how our research subjects in rural Ghana describe to us the operation of sharing arrangements within their social networks. First, we include a costly, impure, 'warm glow' component to altruistic preferences (following Andreoni (1990)), the gains from which diminish as one gives more gifts within one's network. Following the logic of social taxation, network members make demands on individuals who enjoy positive, observable income shocks. But while individuals might vary in the extent of their altruism, everyone faces some outer limit to the pleasure they derive from beneficence or compliance with social taxation norms. If giving has constant marginal cost and the marginal returns to giving diminish,[3] there then emerges some point at which even altruistic individuals cease giving because of the excessive social taxation pressures they face. High rates of social taxation, which might arise due to large networks, can thereby induce low giving from public income shocks. We term this the 'shutdown hypothesis'.

Second, when stochastic income realizations are publicly observable, the insurance claims of less needy members of the solidarity network to share in a windfall can crowd out altruistic giving to those with greater need.[4] This reinforces the shutdown hypothesis. And it implies that in the presence of altruism, private rather than public income might better harness social networks so as to target the least well off in a population.

These two key, realistic refinements eliminate the sharp predictions of standard informal insurance or social taxation models as regards the effect of private and public income shocks on inter-household transfers. Inter-household transfers now become non-monotone in response to public income shocks and potentially increasing in private income shocks. These enhancements also allow us to draw out several other, more subtle, testable hypotheses that match our data remarkably well.

Our model thereby fits the experimental data while still accommodating the core, sensible insights of the informal insurance and social taxation literatures. Individuals value consumption smoothing and seek to leverage networks to accomplish that goal. They also face pressures from within their network to surrender scarce resources and would therefore like to shield their gains from others. By re-introducing the possibility of (imperfectly) altru-

[3]This just requires that the marginal returns to giving diminish faster than the marginal costs of giving.

[4]Concave utility implies that altruistic individuals would like to target their giving toward the neediest members of their social network.

istic preferences, we show that one can reconcile the informal insurance and social taxation literatures with each other and with the data, while also allowing for a richer set of observed behaviors.

We then successfully test these more refined hypotheses in the field experimental data. First, we confirm the prediction that the average size of gifts one gives within one's network is larger for private than for public windfall gains, consistent with altruistic preferences. Second, and relatedly, those with unobservable, private income gains target their giving to the neediest households within their networks. Private, altruistic giving is more sensitive to correcting maldistribution than is sharing of public gains that necessarily addresses the insurance and social taxation motives within networks as well.

Third, over a significant range of network sizes, the number of gifts given is similar for public and private winnings, consistent with greater network demand for transfers when windfalls are observable. But, fourth, the shutdown hypothesis holds. Winners of publicly revealed cash prizes cease making transfers at all when they have too large a network. We rule out a number of alternative explanations of these results in section 6.

We also demonstrate the multifunctionality of these networks by showing that limited risk pooling among family members holds when income is publicly observed. Public income shocks increase transfers among family ties 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. This suggests that private income is not easily observed among family ties, as likewise found by De Weerdt, Genicot, and Mesnard (2019) and Kinnan (accepted 2019). By directly estimating giving as a function of private income shocks, we show that altruism drives some transfers. In particular, private income increases transfers to the neediest within the village, often individuals who are not members of one's family. Cumulatively, these results suggest that attempts to test for the dominance of one inter-household transfer mechanism over another may mislead, as these behaviors reflect a blend of insurance and altruism motives, mediated by social taxation pressures that likely arise primarily from kinship ties.

Our findings have practical policy implications, especially for cash transfer programs which have, over the past decade or two, become the foundation for many social protection programs throughout the developing world. For example, if networks are sufficiently well-connected and populations are motivated by the well-being of others in the network, then transparency may limit the efficiency of redistributive behaviors within networks. These results suggest that communities in many parts of the world have intimate knowledge of their members' needs and can potentially allocate resources more efficiently than state institutions (Alderman, 2002; Bowles and Gintis, 2002). Although transparency of the transfer is essential in informal insurance arrangements among purely self-interested agents, transparent cash transfer policies may impede altruistic agents' ability to focus their giving on the most needy as they are compelled to respond to social taxation or informal insurance demands from the less needy within their network, especially extended family.

2. DATA AND DESCRIPTIVE EVIDENCE

We combine a field experiment with household surveys to construct the data used in the analysis. The field experiments were conducted between March and October 2009 in conjunction with a year-long household survey in four communities in Akwapim South district of Ghana's Eastern Region. This district lies some 40 miles north of the nation's capital, Accra, but is sufficiently far away that only a handful of respondents commute to the greater Accra metropolitan area for work. The sample consists of approximately 70 households from each of the four communities.[5] Individuals in the sample include the household head and his spouse.[6] 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 round of the survey was designed as a baseline, therefore no lottery took place in that round. One week before each subsequent round we visited each village to distribute

[5] The survey was part of a three-wave panel, the first two waves having been conducted in 1997-98 (e.g., in Conley and Udry (2010)) and 2004 (Vanderpuye-Orgle and Barrett, 2009). More than half of the 70 households 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.

[6] 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.

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 38 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.[7] The values of the prizes varied from GH¢10 to GH¢70 as described in Figure 1.[8] The prizes were of a substantial size - the largest prize is equivalent to a month's worth of food consumption for an average household with six or so members. In aggregate, each community's survey participants received GH¢370 of cash in each round to use however they would like.

The lotteries took place one week before the commencement of the survey interviews. We took great care to make clear to participants that the allocation of prizes was random, and that each individual had an equal chance of winning in each round (i.e., draws were identical and independently distributed, without replacement). A village meeting was held in a central area of the community, 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.[9] 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

[7] In this manner, 23 percent of households won a private cash prize and 23 percent won a public cash prize over the course of the year with very little overlap. Of the households who won, 4 won a public prize on two occasions (one of these households also won a private prize); 9 households won a private prize on two occasions (2 of these also won a public prize) and one very lucky household won a private prize on 3 occasions and a public prize on 3 occasions.

[8] During the course of our study, one GH¢ was roughly equivalent to 0.7 USD. In this paper, we are primarily interested in transfers of divisible windfall gains of constant known value among households within a round, 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 further complicates the use of livestock in the analysis. Additionally, in this study context, it is more difficult to 'privately' grant lottery winners a large goat than it is to privately grant them the same amount in cash.

[9] 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.

10 Cash prizes per village
- 5 Private (GH¢10, 20, 35, 50, 70)
- 5 Public (GH¢10, 20, 35, 50, 70)

Figure 1: Experimental Data: Lottery Payouts

prizes that day, framing the prizes as a gratuity for their participation in the survey.[10] 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.[11] 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.

After the public lottery prizes were distributed, we conducted a second 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 identification number. They then waited to enter a closed school room, one at a time, where an enumerator invited them to draw a bottle cap without replacement from a bag. There was one bottle cap for each of the N respondents in the community. Of these, N - 10 were non winning tokens (red colored) and ten were winning tokens, marked distinctively to indicate one of the ten prizes listed in Figure 1.[12] Those who drew winning tokens were informed

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

[11] 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.

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

immediately that they had won a prize, which was identified to them, and were told that they did not have to tell anyone else that they had won. We emphasized that the survey team would not divulge the identities of winners who won in private. Cash prizes were given to the winners immediately and winners commonly hid their prizes in their clothes before leaving the room. The survey interviews in each round commenced one week after the lottery, deliberately delayed to allow winners to receive their prize and do something with it. The interviews took place in no specified order throughout the following three weeks, so that some winners were interviewed a week after receiving their prize, and others up to four weeks afterward.

Survey. Each respondent was interviewed five times during 2009, once every two months between February and November.[13] Each survey round took approximately three weeks to complete, with the two survey teams each alternating between two villages. The survey covered a wide range of subjects including personal income, farming and non-farm business activities, inter-household gifts, transfers and loans, and household consumption expenditures. In each round, both the husband and wife heading each household were interviewed separately on all of these topics.[14] Our data set is assembled mainly using information 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.[15] The value of the gift given and an estimated value for in-kind gifts were also recorded. In our 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. We also focus on gifts to other households within the village and we, therefore, drop gifts given to parties who

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 anticipate problems similar to those reported in Okeke and Godlonton (2014) where vouchers were misallocated by field staff.

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

[14] There were some households with multiple spouses and others without a spouse. For simplicity, throughout the paper we describe households as having a household head and spouse.

[15] When a gift is given to or received from another respondent, our enumerators also provided the unique sample identifier for that individual in our data. This enables dyadic analysis, which we explore in section 5.

TABLE 1

HOUSEHOLD SUMMARY STATISTICS

				Percentile	
	N	Mean	Sd	5th	95th
HH size	315	6.66	2.64	3	11
Cash Gifts Given (last 2 months):					
Number	1,561	0.74	1.22	0	3
Value GH¢ (Total Given)	1,561	9.77	62.73	0	35
Value GH¢ (Conditional on Giving)	615	24.79	98.11	1	80
Cash Gifts Received (last 2 months):					
Number	1,561	0.26	0.71	0	2
Value GH¢ (Total Received)	1,561	2.00	12.17	0	10
Value GH¢ (Conditional on Receiving)	264	11.81	27.61	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.

reside outside of the village and we drop 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). With respect to gifts received, we are interested in gifts from others who win lottery prizes. Thus, we drop observations of gifts received from others who do not reside within the village. 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).

Summary Statistics. Household aggregate measures that form the basis of our analysis are represented in Table 1. On average, each household has more than six members. Across the five rounds of data, households give and receive 0.74 and 0.26 cash gifts, respectively, to any other household in the village over the course of two months. Conditional on giving a gift, the average total value of the gifts given and received is 24.79 GH¢ and 11.81 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 2.4 GH¢ in each of the four rounds in which we distributed cash prizes.

Appendix Table C.2 presents balance tests of variables collected at baseline according to whether one member of the household won any of the public or private lottery at any point over the course of the year. 119 of the households in the study are thus in our "treatment" group while the remaining 190 did not win a cash prize. We also separate the test according to the households that won the privately revealed vs. publicly revealed lottery. The table suggests that randomization was successful — of the 21 tests along which we seek to reject balance, one is significant at the five percent level and another is significant at the ten percent level. For the others, balance cannot be rejected at the ten percent significance level.

3. 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 of canonical models of purely self-interested informal insurance and social taxation: 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 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, 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:

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

TABLE 2

Prize Winnings and Gift Giving

		Gift Giving		
Dependent Variable:		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Randomized Explanatory Variables				
Value of Private Cash Prize	β_v	0.149^{**}	0.129^{**}	0.166^{***}
		(0.069)	(0.055)	(0.057)
Value of Public Cash Prize	β_b	0.00789	-0.0265	0.0639
		(0.071)	(0.057)	(0.058)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-value: $\beta_v = \beta_b$		0.15	0.05	0.21
P-value: $\beta_v <= \beta_b$		0.08	0.02	0.10
Left-censored Obs.		946	946	946
Observations		1,561	1,561	1,561

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. The 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 with a lower bound of zero. Table C.3 reports estimates of the number of gifts given using a Poisson estimator with qualitatively similar results as those in column 3.

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 of publicly revealed lottery winnings, 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 2 reports the estimation results of model 1 with three different outcome variables: log total value of gifts given, log average value of gifts given, and the total number of gifts given per household. None of the (β_b) coefficient estimates is statistically significant at the ten percent level. Moreover, the point estimates are all smaller in magnitude than the (β_v) estimates, each of which is statistically significantly positive at the five percent level. [16]

[16]Three robustness checks confirm the results in Table 2. First, the number of gifts given is integer-

Clearly, interhousehold financial flows do not significantly respond to publicly observable income shocks but they do respond positively to private income shocks.

We can therefore overwhelmingly reject the paired core predictions of purely self-interested models of inter-household transfers. This motivates us to turn in the next section to refining the canonical model of dynamic household choice, incorporating a few small features informed by our discussions with and observations of our Ghanaian subjects. We show that by building impure altruism and social taxation into a fairly standard model of a dynamic game among agents facing stochastic income streams, we generate more nuanced predictions that reconcile fully with our data.

4. THE ENHANCED MODEL

In the model that follows, we show that a few reasonable, empirically-grounded changes to a canonical model of risk-pooling can alter its predictions in important ways. We build on Foster and Rosenzweig (2001), who model transfers in the context of a two agent game in which agents can hold altruistic preferences over each other's consumption, in addition to standard preferences for their own consumption, and the commitment to a transfer contract is imperfect due to lack of exogenous enforcement mechanisms. We add to this model by 1) allowing "warm glow" altruistic preferences that generate diminishing marginal utility in the number of gifts given, 2) imposing a cost associated with gift giving, and 3) making the number of gift requests one receives an increasing function of the share of one's income that is publicly versus privately revealed, reflecting social taxation pressures. These seemingly innocuous adjustments, grounded in our observation of solidarity network activity and our field research subjects' descriptions of behaviors within their villages, generate more nuanced predictions that obviate the public observability hypothesis that we just rejected in our experimental data. Rather, we show that when inter-household transfer motives are not limited to myopically self-interested dynamic behavior, risk pooling may be incomplete, and larger networks and publicly observable income may not be desirable. While one could model this type of giving on a full network, our core predictions do not depend on the strategic

valued, so we also use a Poisson count data estimator for that dependent variable (Appendix Table C.3). Second, in-kind (livestock) lottery winnings could contaminate our results, so we include them in an alternate specification (Appendix Table C.4). In both cases, the results remain qualitatively unchanged. Finally, there are 24 missing household-round observations out of the 1,585 possible across the five rounds of data (i.e., 1.5%). We find no evidence of attrition bias in our sample (Table C.5).

interplay of gift giving along the network. We therefore rely on the simpler, established two household 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 $h_t = \{s_1, s_2, ..., s_t\}$.[17] We model the choice of history-dependent transfers from household 1 to household 2, $\tau(h_t)$, in period t. Both households have gift links with $g_1 = g_2 \geq 1$ other households. Depending on the realization of a particular state, households will receive $g_i p_i(s_t)$ different gift requests from their network, where $0 \leq p_i(s_t) \leq 1$ reflects the unconditional probability that a given household in one's network will request a transfer in period t — $p_i(s_t)$ is larger when the income realization is publicly revealed to i's network.

To focus attention on transfers between households 1 and 2, we assume that net transfers with all other households in one's network equal zero. Thus, net income for household 1 is $y_1(s_t) - \tau(h_t)$ and net income for household 2 is $y_2(s_t) + \tau(h_t)$. If $\tau(h_t) > 0$, then household 1 (2) is a net sender (receiver) of transfers. Otherwise, if $\tau(h_t) < 0$ household 1 (2) is a net receiver (sender) of transfers within the dyad.

We note that while 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, however, that there are implications for endogenous network choice that emerge from the principles reflected in our enhanced model. We address the potential for endogenous networks below, empirically, and preserve this phenomenon for future analysis and discuss potential next steps in the conclusion.

Preferences. Following Foster and Rosenzweig (2001), we assume households hold altruistic preferences towards others' single-period utilities. We introduce individual i's altruistic preferences by assuming that household single-period utility is separable in own and other

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

household consumption. Single-period utility for household 1 is:

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

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

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

$$
\gamma_1(g_1, s_t) = \min\{\bar{\gamma}_1^F + \frac{\bar{\gamma}_1^W}{g_1 \cdot p_1(s_t)} \mathbb{1}(\tau(h_t) \neq 0), \overline{\gamma}_1\} \tag{3}
$$

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

Explicitly stated, we assume that the amount of warm glow gains household 1 derives from transfers to household 2 decreases in the total number of household 1's period t gift obligations, $g_1 \cdot p_1(s_t)$. This reflects the idea that warm glow increases at a diminishing rate in the number of discrete transfers each household participates in — intuitively, the warm glow of giving dims as transfers become more commonplace. And so long as utility is concave in consumption, the marginal warm glow from giving will be higher when transfers are directed to otherwise-poorer households, because the marginal utility of added consumption is higher among the poor. Without loss of generality, we will set $\bar{\gamma}_1^F = 0$ and focus our analysis around warm glow altruism — thus, when we speak of altruism moving forward, we are no longer referring to "pure" altruism. Intuitively, and taken together, each household is altruistic towards others, but not without limit. Households may vary in the "stock" of altruism (or altruistic capital as in Ashraf and Bandiera (2017)) they possess, but will be limited in the degree of altruism they exercise towards other households.

Dynamic Payoffs and Transfer Choices. At period t, households seek to maximize their expected lifetime utility, which requires agreeing upon a history-contingent transfer contract that is preferable to zero transfers across all states. Thus, we assume that households compare payoffs from the dynamic contract to payoffs from a no-transfer rule.[18] To set up the household's problem, we define $U_1(h_t)$ as 1's expected discounted utility gain from the risk-sharing contract with 2 relative to a no-transfer rule after history h_t:

$$
\begin{aligned}
U_1(h_t) = \quad & u_1(y_1(s_t) - \tau(h_t)) - u_1(y_1(s_t)) \\
& + \gamma_1(g_1, s_t) u_2(y_2(s_t) + \tau(h_t)) - \gamma_1(g_1, s_t) u_2(y_2(s_t)) \\
& + \mathbb{E} \sum_{k=t+1}^{\infty} \delta^{k-t} \left\{ \begin{array}{l} u_1(y_1(s_k) - \tau(h_k)) - u_1(y_1(s_k)) \\ +\gamma_1(g_1, h_t) u_2(y_2(s_k) - \tau(h_k)) - \gamma_1(g_1, h_t) u_2(y_2(s_k)) \end{array} \right\} \\
& - \alpha_1(g_1)
\end{aligned}
\tag{4}
$$

where δ represents the dynamic discounting factor. The cost term $\alpha_1(g_1)$ represents a second way in which our model diverges from others' — it is the incremental cost to household 1 of maintaining a gift-giving link with household 2 given network size g_1. We assume that the cost of maintaining such a link is weakly convex in network size and can be thought of as the effort required to maintain a social bond and, for example, awareness of household 2's realized income. The contract is enforced if the expected discounted utility surplus is nonnegative. The contract requires an implementability constraint that states that gains from the contract be at least as high as the no-transfer rule for both parties: $U_1(h_t) \geq 0$ and $U_2(h_t) \geq 0$. Together, the economic environment, payoffs and transfer decision represent a simultaneous game in which agents seek to find a contract that can be implemented in the presence of limited commitment and no external enforcement mechanism.

[18] Households in Foster and Rosenzweig (2001) revert to a sequence of history-dependent Nash equilibria (SHDNE) in which transfers are maintained even when a household defaults from the contract. Such an environment is not crucial for the type of analysis we conduct in our study. Nevertheless, appendix section A.1 shows how one can adapt our model to reflect such SHDNE default transfers.

Limited Commitment Contract Solution. Following Foster and Rosenzweig (2001) and Ligon, Thomas, and Worrall (2002), the solution to the utility maximization problem will be a dynamic program in which the current state is given by state s, given targeted discounted utility gain for household 2, U_2^s.[19] Choice variables in the programming problem will be consumption levels c_1, c_2 and the continuation utilities U_1^r and U_2^r for each possible state r, reflecting the next period's optimization problem. This enables us to write the value function for household 1 as dependent on current target utilities and collective resources: $U_2^s, \{y_1(s) + y_2(s)\}$. Formally, we write the dynamic programming problem as

$$
\begin{aligned}
U_1^s(U_2^s) = \quad \max_{\tau_s,(U_1^r,U_2^r)_{r=1}^S} \quad & u_1(y_1(s) - \tau_s) - u_1(y_1(s)) \\
& + \gamma_1(g_1(s))u_2(y_2(s) + \tau_s) - \gamma_1(g_1(s))u_2(y_2(s)) \\
& - \alpha_1(g_1) + \delta \sum \pi_{sr} U_1^r(U_2^r)
\end{aligned} \tag{5}
$$

subject to

$$
\begin{aligned}
\lambda: \quad & u_2(y_2(s) + \tau_s) - u_2(y_2(s)) \\
& + \gamma_2(g_2(s))u_1(y_1(s) - \tau_s) - \gamma_2(g_2(s))u_1(y_1(s)) \\
& - \alpha_2(g_2) + \delta \sum_{r=1}^{S} \pi_{sr} U_2^r \geq U_2^s
\end{aligned} \tag{6}
$$

$$
\delta\pi_{sr}\mu_r: \quad U_1^r(U_2^r) \geq \underline{U}_1^r = 0 \quad \forall r \in S \tag{7}
$$

$$
\delta\pi_r\phi_r: \quad U_2^r \geq \underline{U}_2^r = 0 \quad \forall r \in S \tag{8}
$$

$$
\psi_1: \quad y_1(s) - \tau_s \geq 0 \tag{9}
$$

$$
\psi_2: \quad y_2(s) + \tau_s \geq 0, \tag{10}
$$

where π_{sr} represents the probability of state r occuring. Equation 6 says that transfer and future utility allocations will satisfy the promise-keeping constraint. Equations 7 and 8 state

[19] $U_2^s(U_1^s)$ is defined by equation 5 (and its constraints) when all subscripts with 1 are replaced with a 2 and vice versa.

that allocated utility in any state r will be at least as high as the lower bound utility household 1 and, respectively, 2 can receive via defaulting to the no-transfer arrangement. Equations 9 and 10 place non-negativity constraints on consumption allocations in period s. The actual contract can be computed recursively, starting with an initial value for U_2^s.

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

$$\frac{u_1'(y_1(s_t) - \tau(h_t)) + \gamma_1(g_1(h_t))u_2'(y_2(s_t) + \tau(h_t))}{u_2'(y_2(s_t) + \tau(h_t)) + \gamma_2(g_2(h_t))u_1'(y_1(s_t) - \tau(h_t))} = \lambda + \frac{\psi_2 - \psi_1}{u_2'(y_2(s_t) - \tau(h_t)} \tag{11}$$

$$-U_1^{r\prime}(U_2^r) = \frac{\lambda + \phi_r}{1 + \mu_r}, \quad \forall r \in S \tag{12}$$

$$\lambda = -U_1^{s\prime}(U_2^s). \tag{13}$$

Taken together, these three conditions imply that a constrained-efficient contract can be characterized in terms of the evolution over time of λ, where $-\lambda$ is the slope of the Pareto frontier.[20] For each state s, there is a history independent 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 $\lambda(h_t)$ in the following manner

$$\lambda(h_{t+1}) = \begin{cases} \underline{\lambda}_s \text{ if } \lambda(h_t) < \underline{\lambda}_s \\ \lambda(h_t) \text{ if } \underline{\lambda}_s \leq \lambda(h_t) \leq \overline{\lambda}_s \\ \overline{\lambda}_s \text{ if } \lambda(h_t) > \overline{\lambda}_s. \end{cases} \tag{14}$$

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

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

non-linear system of equations.

Income shocks. We now add more structure to the model to study the importance of the transparency of cash transfers. Let us define two types of exogenous income shocks: 1) privately revealed cash prizes (denoted by v) and 2) publicly revealed cash prizes (b). Households that do not receive cash prizes experience zero exogenous income shocks (z). Thus, there are potentially nine different states that can be realized, though we limit our analysis to states in which only up to one household receives a prize of any type: neither 1 nor 2 receive a prize (zz), 2 receives a private prize (vz), 2 receives a public prize (bz), 1 receives a private prize (zv), and 1 receives a public prize (zb).[21] Explicitly, here we are assuming that the prize-winning household receives a higher income than the non-prize winning household and the private and public prizes are equal in value:

Assumption 1 (Prize-winners Have Higher Incomes)

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

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

Assumption 2 (Observability of Income)

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

This assumption reflects the notion that households who enjoy observable windfall gains will experience some 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. Rather than introducing an endogenous cost to hide one's income, the above assumption is analytically equivalent to a modeling framework in which it is infeasible to hide income in the public income state and costless to hide income in the private income state (so that income will never be revealed in this state). This yields the sharp binary distinction between the private and public income states in the probability of

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

receiving a gift request.

This assumption also implies that the warm glow altruism weight household 1 holds towards household 2, for example, decreases when household 1 wins a publicly revealed lottery and is likely to face additional request for transfers, fulfillment of which also entails transactions costs beyond the amount transferred.

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

We find that as network size increases, the marginal utility of participating in a risk-sharing contract decreases in network size, but decreases at a faster rate in the state when a household wins a public prize — this is because the degree to which altruism motivates the transfer is smaller in the public than private state. When gift-giving links are also costly to maintain, a household will "shut down" all giving — beyond a certain network size threshold, if requests for gifts are too large (public income state), then the household will not give any gifts. Additionally transfers will in most cases be **larger** when a household wins a privately-revealed prize and can concentrate its giving to those who generate the greatest warm glow.

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 2 uses simulated gift transfers between households 1 and 2 to show the empirical implications of the shut down hypothesis. Notice that at small gift network sizes, household 1 transfers the same amount to household 2 regardless of being in state zv or zb. In this interval, household 1 is not overly burdened by transfer requests, so each of its transfers continue to

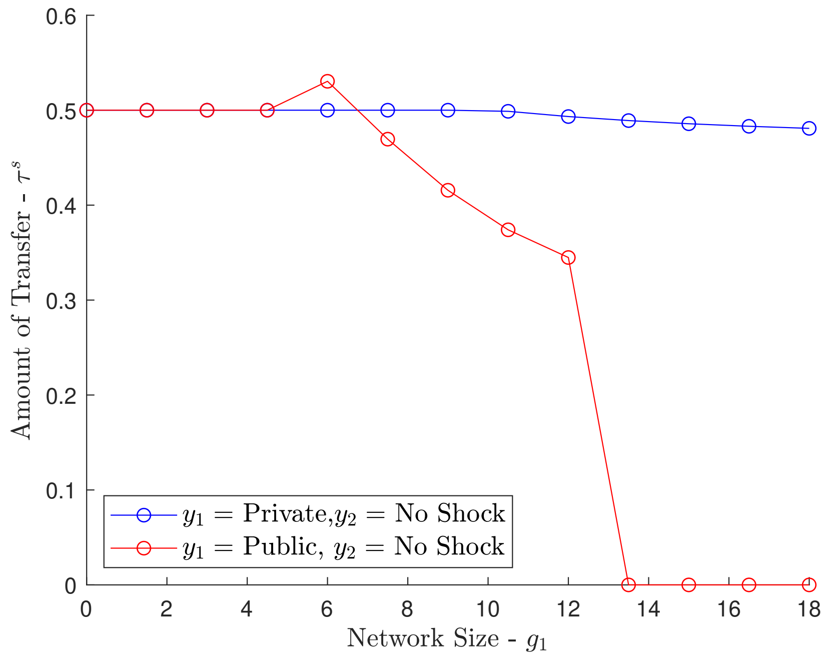

Note: This figure represents transfer amounts $\underline{\tau}^s$ from household 1 to household 2 when household 2 takes the entire share of the surplus (U_1^s is set to zero) and when household 1 wins a cash prize. Thus, it also represents the average transfer amount from household 1 to any other household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

FIGURE 2: Amount of Transfer by Network Size

maintain a high degree of altruistic motives. However, as the network size increases, gift requests in the public income state will also increase. The decreasing marginal utility causes households to decrease the amount of their transfer until they fall to zero at the shutdown threshold and beyond.[22] This relationship leads to two additional empirical implications:

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

Prediction 3 (Publicly Revealed Prize → Higher Number of Gifts Given) *The av-*

[22]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 (Allcott, Karlan, Mobius, Rosenblat, and Szeidl, 2007; Coleman, 1990)

erage 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. Due to the requests to share their publicly revealed income, the household that receives the public shock will naturally give a larger number of gifts in total to other households if they do not shut down (Prediction 2). However, each gift given in the public income state will be of smaller value on average than gifts given in the private income state (Prediction 2).

Depending on the average value of the gifts given in each state, the above two predictions potentially imply that the total value of gifts given by households who receive public income shocks are greater than the total value of gifts given by households with private income shocks, but only prior to reaching the shut down threshold. This prediction relies on the rate at which transfers decrease after the zb state ceases to overlap with the zv state.

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

So far we have discussed how the model generates predictions regarding the gift transfer behavior of household 1. Naturally, if household 2 receives gifts from household 1, we should be able to symmetrically identify changes in household 2's consumption as a function of household 1's lottery winnings. This implies that household 2's consumption levels will be higher on average when their gift-giving network wins a prize. However, since transfers are predicted to be larger when the peer household wins a privately revealed prize, it is likely that the effect will only be observed in the private income state. Furthermore, since household 1's marginal utility is decreasing in household 2's consumption, we should see stronger and more progressive patterns of gift giving through the private lottery when the consumption gap between households 1 and 2 is large. It is straightforward to show via simulation that average transfer sizes increase as the gap between 1 and 2's per-period income increases.[23] This leads to the final prediction:

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

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

5. EMPIRICAL INVESTIGATION

The model implications derived in Section 4 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.

5.1. *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. More specifically, 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.

Due to the nature of the data, we can exactly identify the directionality of giving, including each of the bi-directional, or reciprocal, gift links in our sample. We do this by comparing individual i's response regarding j's gift-giving behavior with individual j's response of i's gift-giving behavior. 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, we establish that a reciprocal gift link exists between these two individuals. We define g_{ij} as the reciprocal link between individuals i and j in the sample and $g_{ij} = 1$ if both individuals confirm the existence of a reciprocal gift-giving link 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.[24] 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.[25] Network size was calculated using the total number of links created in this manner at the household level.

A point of clarification regarding the difference between the social network mapping and the gift transfer data is warranted at this point. We collected data on actual transfers between households in the gifts module in each of rounds one through five as described in section 2. 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.[26]

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. Appendix Table C.1 reports individual summary statistics. This table 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.

[24] 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. 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).

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

[26] We analyze gifts given within networks and gifts given outside of networks in section 6 below.

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

$$\overline{\text{Private}}_{it} = \sum_{j=1}^{N} \frac{\text{Private}_j \times \mathbb{1}(g_{ij} == 1)}{\sum_{j=1}^{N} \mathbb{1}(g_{ij} == 1)}, \tag{15}$$

where $\text{Private}_j \in \{0, 10, 20, 35, 50, 70\}$ are the values of cash prizes household j can win and $\mathbb{1}$ represents the indicator function.

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

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

the positive income shock, and household 2, the household receiving the transfer, also depends on household 1's network size. Recall that as network size increases, the average value of the gift given by household 1 decreases as well. Therefore, a more theoretically appropriate network average adjusts network winnings by household 1's network size.

We therefore construct an "adjusted average network value" by weighting 2's network winnings by the inverse of 2's network size. To provide intuition, consider that household 2 has gift obligations to X other households. If household 2 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

$$\overline{\text{Private}}'_{it} = \sum_{j=1}^{N} \frac{\frac{\text{Private}_j}{\sum_{k=1}^{N} \mathbb{1}(g_{jk}==1)} \times \mathbb{1}(g_{ij} == 1)}{\sum_{j=1}^{N} \mathbb{1}(g_{ij} == 1)}. \tag{16}$$

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 3 presents our measure of network size. The average network size, defined by the number of inter-household reciprocal gift-giving links, is 11.4 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.

[29] Figure D.1 displays a histogram of the distribution of network size, showing that our analysis is unlikely to be skewed by the presence of households with large, outlying, values on network size.

TABLE 3

Household Summary Statistics for the Enhanced Model

	N	Mean	Sd	Percentile 5th	Percentile 95th
Network Size:					
N of HH in Network	315	11.40	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¢):					
Average Value of Private Network Prize	1,257	2.30	5.24	0	9.23
Average Value of Public Network Prize	1,257	2.08	3.93	0	8.75
Adjusted Average Value (Private)	1,257	0.20	1.20	0	0.63
Adjusted Average Value (Public)	1,257	0.20	1.10	0	0.74

Note: Gift 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. Network average lottery winnings calculate the average lottery winnings of a household's network. The adjusted average calculates an average of a household's network lottery winnings divided by the networked household's network size.

5.2. *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 by household i. The shut down hypothesis (Prediction 1 in Section 4) can be investigated using the following regression:

$$\begin{aligned} y_{it} = \alpha &+ \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} \\ &+ \beta_{vg} \text{Private}_{it} \times \text{Net-size}_i + \beta_{bg} \text{Public}_{it} \times \text{Net-size}_i \\ &+ \text{hh}_i + \text{r}_{tv} + \epsilon_{it}, \end{aligned} \tag{17}$$

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. 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. Nevertheless, we show that our results are robust to alternative definitions of networks in section 6.

Predictions 2 and 3, that do not depend on heterogeneity in network size, can simply be tested by setting the interaction terms equal to zero.

Table 4 contains the estimation results of equation 17 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 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 the average (10.4), we can reject similarity in gift-giving behavior across all specifications. We calculate the shut down point predicted by the linear model as a network size of 9.15, 7.27, and 10.25 for columns 1-3 respectively. In other words, households give zero additional gifts following public income shocks when they have around 10 other households in their gift giving network.

TABLE 4

Testing the Shut Down Hypothesis

		Gift Giving		
Dependent Variable:		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Randomized Explanatory Variables With Network Size Interaction				
Value of Private Cash Prize	$\beta_v > 0$	0.296***	0.199**	0.226**
		(0.114)	(0.092)	(0.094)
Value of Private Cash Prize × N	$\beta_{vg} \leq 0$	-0.012*	-0.005	-0.005
		(0.007)	(0.006)	(0.006)
Value of Public Cash Prize	$\beta_b > 0$	0.264**	0.115	0.420***
		(0.111)	(0.088)	(0.091)
Value of Public Cash Prize × N	$\beta_{bg} < 0$	-0.029***	-0.016**	-0.041***
		(0.010)	(0.008)	(0.008)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0: \beta_v = \beta_b$		0.84	0.50	0.13
P-Value $H_0: \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		0.32	0.15	0.88
P-Value $H_0: \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.05	0.02	0.05
P-Value $H_0: \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.02	0.02	0.00
N at Shut Down		9.15	7.27	10.25
Left-censored Obs.		946	946	946
Observations		1,561	1,561	1,561

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. 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. P-values reported under each column for each of the hypotheses. N denotes network size. N at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$.

Our model predicts that $\beta_v > \beta_b$ with respect to the average value of gifts given, which is supported by our findings. Furthermore, Table 4 shows that the relationship between these point estimates is maintained at a network size of zero in column 2, though we cannot reject equivalence in small networks (up to a network size of around 7).

We also predict that the number of gifts given following public cash prizes will be larger than private cash prizes in small networks (Prediction 3). When we interact network size in Table 4, we see that the relationship between the point estimates flips relative to Table 2, precisely as suggested by our model. Furthermore, we can reject the equality null in favor of the one-tailed alternate hypothesis that $\beta_b < \beta_v$ in column 3.[30]

Finally, Prediction 4 provides a parallel hypothesis with respect to the total value of gifts given. While the point estimate for β_b in column 1 table 4 increases relative to its analogous coefficient in Table 2, we cannot reject that the total volume of gifts given following private and public cash prizes are equivalent.[31]

Together, the results associated with the first four predictions suggest a clear pattern of behaviors that emerge following private versus public cash transfers. 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. First to decrease is the total value of gifts given: households with network size of around five households give slightly more, but smaller, gifts upon winning a publicly revealed cash transfer than households with similar network size who win privately revealed transfers. Once network size reaches around 10, one unit less than the mean network size, publicly revealed prizes no longer have any effect whatsoever on giving. Figures D.2 and D.3 further suggest that households with very large networks, give significantly fewer gifts upon winning a public prize than they would without having won a public prize.[32] This suggests, that the transparent cash transfer causes households with large networks to shut down their giving, even to households to whom they otherwise would have transferred gifts. This 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, Section 4) 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

[30]This relationship is demonstrated graphically with the aid of a third-order polynomial estimation of equation 18 in Figure D.2.

[31]This relationship is also demonstrated graphically in Figure D.3.

[32]The number approximates 15, which translates to the 70th percentile of gift-network size across all villages.

capita food consumption, our preferred proxy for consumption in these data. We test this using the following equation:

$$y_{it} = \alpha + \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} + \beta_{vn} \overline{\text{Private}}'_{it} + \beta_{bn} \overline{\text{Public}}'_{it} + \text{r}_{tv} + \epsilon_{it}, \tag{18}$$

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 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.[33] We again include village-specific round fixed effects, r_{tv}.

Given the assumed concavity of utility in consumption and in the presence of altruism, we expect that households with lower levels of period-specific food consumption will receive more support from their network. This feature of the model 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 to be larger at the lower 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.[34]

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.

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 18 in Figure 3 (appendix Table C.6 shows es-

[33] We repeat the analysis for $\overline{\text{Private}}_{it}$ and $\overline{\text{Public}}_{it}$ in appendix Figure D.4 with qualitatively similar, but noisier, results.

[34] Appendix Figure D.5 uses a simple lowess estimator to demonstrate how home-produced food consumption over the past month varies with survey date. Food availability is clearly most constrained from around the middle of March to early July, corresponding to survey rounds two and three.

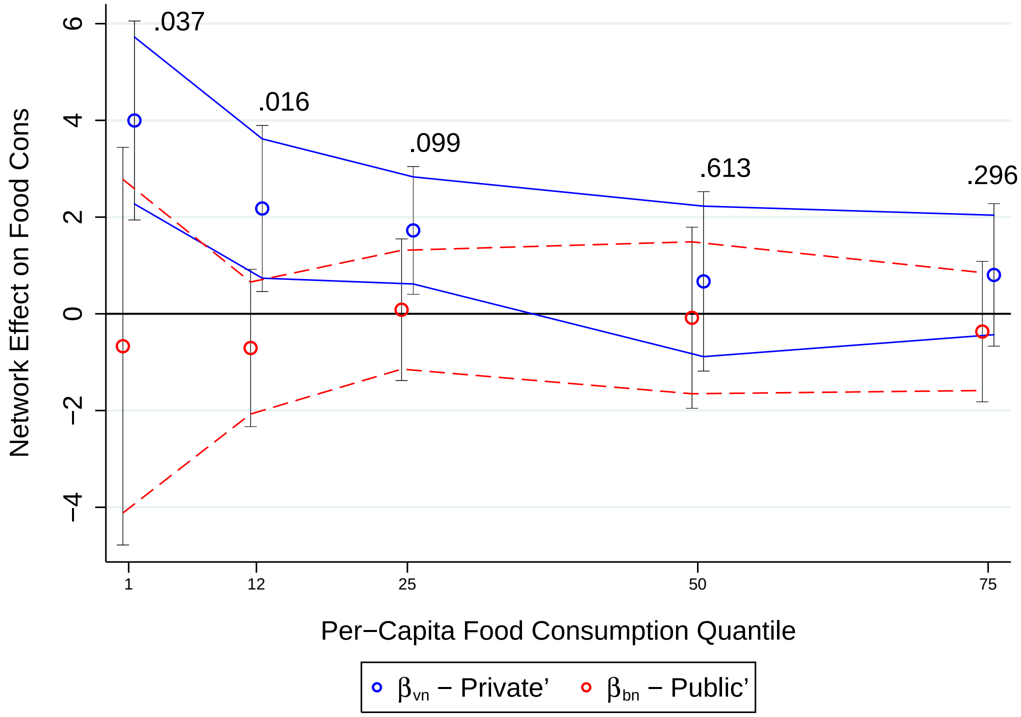

Note: Results of a simultaneous quantile regression at 1st, 12.5th, 25th, 50th, and 75th quantiles bootstrapped over 1,000 iterations. Dependent variable is log home-produced per capita food consumption over the last month. Quantiles represented on the x axis. Blue dots (lines) show the coefficient estimates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

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

timation 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. By contrast, the coefficient estimates on a household's network's public lottery winnings, depicted with red dots and lines, are insignificantly different from zero throughout the distribution. 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. These results are consistent with both altruistic motives for giving and the shut down hypothesis, as reflected in Prediction 5 of our model.

6. 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. The quantile regression analysis above is powerful because it tests whether the consumption of a gift-recipient household increases in network gift-giving. We found that this is true for households at the lower end of the food consumption spectrum.

A more direct test, however, should 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:

$$
\begin{aligned}
y_{ijtv} = \alpha &+ \beta_v \text{Private}_{it} + \beta_b \text{Public}_{it} \\
&+ \beta_{vF} \text{Private}_{it} \times (\text{Food}_{it} - \text{Food}_{jt}) + \beta_{bF} \text{Public}_{it} \times (\text{Food}_{it} - \text{Food}_{jt}) \\
&+ \gamma(\text{Food}_{it} - \text{Food}_{jt}) + \text{r}_{tv} + \epsilon_{ijt}
\end{aligned} \tag{19}
$$

where y_{ijt} represents giving from household i to household j either in terms of amount given or number of gifts given. Then, $(\text{Food}_{it} - \text{Food}_{jt})$ 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 β_{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 5 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 the social network at baseline. We estimate Model 19 using Tobit and Poisson estimators when the amount given and number given are the respective dependent variables. Estimates in columns one and two reflect Model 19 estimated for the amount and the number of gifts given, respectively. The estimation results are consistent

TABLE 5

Dyadic Regressions

		Gift Giving Within Dyad: From i to j			
Dependent Variable:		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 equals log total value of (cash) gifts given from household i to household j in columns 1, 3 and 4 — estimated using Tobit with observations censored to the left by 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 difference in 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 denotes network size.

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 could be argued that transfers are strategic, following selfish motives, as a means of building network ties. If this is the case, it is difficult to reconcile this 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 innings. However, we do see significant increases in gift-giving to out-of-network households following public winnings.

Columns 3 and 4 in table 5 estimate the shut down hypothesis model in the dyadic setting. Column 4 includes all out-of-network households while column 3 only includes households specified to maintain 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 — β_v and β_{vg} are not significant in either specification. Thus, we do not find the behavior we expect to see 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 an omitted variable or variables, rendering it an endogenous regressor that biases our results. We explore alternative measures of networks in appendix Tables C.7 through C.9 and show consistence with the results obtained thusfar. Specifically, Table C.7 shows that the total number of non-co-resident within village family members (family network size) is the strongest predictor of gift network size — it explains nearly 50% of the variation in gift network size in our sample. Assuming that family network size is exogenous, we replace gift network size with family network size in Table C.8 and obtain similar results. We also generate a linear probability model to predict gift network size (column 1, Table C.7) and obtain qualitatively similar results when we use predicted network size. We conclude that endogenous network selection is not a major threat to our results.

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).

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).[35]

Nevertheless, we explore this possibility in Table 6, differentiating gifts given according to links with varying likely quality of information about recipient households. We again estimate equation 17 where the dependent variable is the log value of gifts given to all kin (i.e., extended family) in column 1, to direct family in column 2 and to village friends in column 3, assuming that information is more difficult to conceal from 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 experience the shut-down condition following public cash winnings. Thus there seems no information story to explain the patterns we observe in the data.

The striking thing about the Table 6 results, however, is that the prediction of the canonical limited commitment informal insurance model seem to fit the data 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 model nests the familiar insurance model. Our analysis

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

TABLE 6

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 equals log average value of (cash) gifts given in household. Column 1 consists of gifts to all family, column 2 to direct family members (i.e., siblings, grandparents, parents) who have their own households 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.

suggests that there is a range of our data within which that 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 B 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) households who "shut-down" are less likely to be transfer recipients in future rounds, 3) it is unlikely that mental accounting of stochastic unearned income explains the results and 4) households are unlikely to be endogenously opting out of informal insurance contracts when they receive large income shocks.

Test of Full Risk Pooling. We conclude this section by testing 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.

Table 7 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 6 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 con-

TABLE 7

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)
$\overline{\text{Private}\vert\text{Network}}_{it}$	0.005	0.057	-0.012	0.025	0.014	-0.320**
	(0.027)	(0.043)	(0.030)	(0.021)	(0.023)	(0.156)
$\overline{\text{Public}\vert\text{Network}}_{it}$	-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.

sumption 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. In summary, our evidence points toward solidarity networks motivated only partly by insurance. Combined with the significant giving we see from private winnings, altruism and social taxation appear to play far more prominent roles driving inter-household transfers than is implied by the self-interested informal insurance motives that underpin the dominant models employed by economists for the past generation.

7. CONCLUSIONS

Inter-household networks within village economies are multi-functional. They can mediate inter-household transfers that resemble credit, insurance, social taxes, altruistic gifts, or some combination of these. Moreover, pure gifts, informal insurance indemnity payments, informal loans, precautionary savings through networks, and social taxes are observationally equivalent in data on flows among households.

Under the dominant models of purely self-interested behavior, transfers should be increasing in a household's publicly observable income shocks, but not with respect to its private, unobservable windfall gains. Our novel experimental data enable us to test, and overwhelmingly reject, this hypothesis. Indeed, we obtain the opposite result — that transfers increase (and are more efficiently redistributed) after a household wins a cash windfall in private, as opposed to households whose cash windfalls are revealed to the whole village. In this setting at least, the transfers that lubricate the village economy reflect more than merely self-interested informal insurance or seemingly compulsory social taxation.

In our alternative model, altruistic motives coalesce with insurance motives and social pressures. Our data fully support the predictions of this refined model of multi-functional village solidarity networks. First, on average, more gifts are given out of private cash winnings than public cash winnings, signaling that altruistic preferences - not just self-interested behavior within an endogenously enforceable insurance scheme - must be a significant driver of inter-household transfers. Second, winners of privately revealed prizes target giving to the neediest households within their networks, indicating greater social welfare gains from altruistic transfers than from insurance transfers. Third, winners of publicly revealed cash prizes do not make transfers when they have very large networks; they break the informal contract due to network size. Fourth, we can reject the null hypotheses of both full and no risk pooling, signalling incomplete risk pooling.

The results give rise to a series of new questions that our framework engenders. For example, what are the longer run, evolutionary implications of the model? How do norms of social taxation interact with altruistic preferences? If large networks are costly to maintain, why and how are such large networks built and how might they be preserved? To what extent does the observability of income determine the size of a gift-giving network? How much (or how little) information do these results imply regarding villagers' knowledge of each other's incomes, and what does that mean for the enforceability of a risk-sharing contract?

These questions are beyond the scope of this paper to answer, however we can consider them in a new light given the framework we develop. For example, in considering optimal network size our paper challenges the reductionist vision of the role social networks play in mediating inter-household cash transfers. Observed patterns appear at least as well explained by a combination of altruism and social taxation, i.e., duty-motivated (deontological) rather than outcome-driven (consequential) behavior. An optimal network size must therefore take into consideration the multiple functions networks play.

Given these theoretical insights and empirical corroboration of the model's predictions, we highlight one last point. 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 co-existence of pro-social, altruistic preferences with self-interested behavior and costly social demands within multi-functional social networks.

REFERENCES

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

ALLCOTT, H., D. KARLAN, M. M. MOBIUS, T. S. ROSENBLAT, AND A. SZEIDL (2007): "Community Size and Network Closure," *American Economic Review*, 97(2), 80–85.

AMBRUS, A., M. MOBIUS, AND A. SZEIDL (2014): "Consumption Risk-Sharing in Social Networks," *American Economic Review*, 104(1), 149–82.

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

ANGELUCCI, M., AND G. DE GIORGI (2009): "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?," *American Economic Review*, 99(1), 486–508.

ASHRAF, N., AND O. BANDIERA (2017): "Altruistic Capital," *American Economic Review*, 107(5), 70–75.

BOURLÉS, R., Y. BRAMOULLÉ, AND E. 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, S., AND H. GINTIS (2002): "Social Capital And Community Governance," *Economic Journal*, 112(483), F419–F436.

CASTILLA, C., AND T. WALKER (2013): "Is Ignorance Bliss? The Effect of Asymmetric Information between Spouses on Intra-household Allocations," *American Economic Review*, 103(3), 263–68.

COATE, S., AND M. RAVALLION (1993): "Reciprocity Without Commitment: Characterization and Performance of Informal Insurance Arrangements," *Journal of Development Economics*, 40(1), 1 – 24.

COLEMAN, J. S. (1990): *Foundations of Social Theory*, Belknap Press Series. Belknap Press of Harvard University Press.

CONLEY, T. G., AND C. R. UDRY (2010): "Learning About a New Technology: Pineapple in Ghana," *American Economic Review*, 100(1), 35–69.

DE WEERDT, J., G. GENICOT, AND A. MESNARD (2019): "Asymmetry of Information within Family Networks," *Journal of Human Resources*, 54(1), 225–254.

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

FOSTER, A. D., AND M. 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, J. (2017): "The Effect of Social Pressure on Expenditures in Malawi," *Journal of Economic Behavior & Organization*, 143, 173 – 185.

JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (2012): "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *American Economic Review*, 102(5), 1857–97.

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

Kinnan, C. (accepted 2019): "Distinguishing Barriers to Insurance in Thai Villages," *Journal of Human Resources*.

Kranton, R. (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, E., and L. Schechter (2012): "Motives for Sharing in Social Networks," *Journal of Development Economics*, 99(1), 13 – 26.

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

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

Olson, M. (1971): *The Logic of Collective Action: Public Goods and the Theory of Groups,*, American Studies Collection. Harvard University Press, Second Printing with a New Preface and Appendix edn.

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

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

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

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

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

Sen, A., and K. Hoff (2006): "The Kin System As a Poverty Trap?," in *Poverty traps*, ed. by S. Bowles, S. Durlauf, and K. Hoff, chap. 4, pp. 95–115. Princeton University Press, Princeton.

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

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

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

Townsend, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, pp. 539–591.

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

Walker, T. (2011): "Risk Coping, Social Networks and Investment in Rural Ghana," *PhD Dissertation*.

APPENDIX A: MATHEMATICAL APPENDIX

A.1. *Adding a Sequence of History-dependent Nash Equilibria (SHDNE) Transfers*

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

$$\tau^D(h_t) = \begin{cases} r \text{ s.t. } u_1'(y_1(s_t) - r)/u_2'(y_2(s_t) + r) = \gamma_1(g_1(h_t)) \\ \qquad \text{if } u_1'(y_1(s_t))/u_2'(y_2(s_t)) < \gamma_1(g_1(h_t)) \\ r \text{ s.t. } u_1'(y_1(s_t) - r)/u_2'(y_2(s_t) + r) = 1/\gamma_2(g_1(h_t)) \\ \qquad \text{if } u_1'(y_1(s_t))/u_2'(y_2(s_t)) > 1/\gamma_2(g_1(h_t)) \\ 0 \text{ otherwise.} \end{cases} \tag{20}$$

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

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

(21)

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

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

A.2. *Model Simulations*

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

Without loss of generality, we focus our analysis on household 1's behavior when it wins either the public, state zb, or private, state zv, prize and household 2 does not receive a shock to income. Figure A.1 shows the evolution of the optimal (log) contract intervals as network size increases. At low network size values, g_1 less than 5, the contract intervals overlap and are unchanging — they are unchanging as an artefact of the model assumptions because we limit warm glow altruism towards household 2 to a maximum of 0.5. Once network size increases beyond 5, the influence of warm glow altruism decreases in the state in which household 1 wins a publicly revealed lottery — zb. The lower- and upper-bound intervals start to increase until they no longer overlap with state zz and then with state zv. In our example, the contract intervals in state zz and zv overlap over the entire domain in Figure A.1. Intuitively, as household 1's marginal (altruistic) utility in household 2's consumption decreases, it will request a larger share of the available surplus in its dynamic transfer arrangement with household 2.

Figure A.2 shows the resulting discounted lifetime expected utility of such a contract when the initial state is either zv or zb and when household 1 extracts all the possible

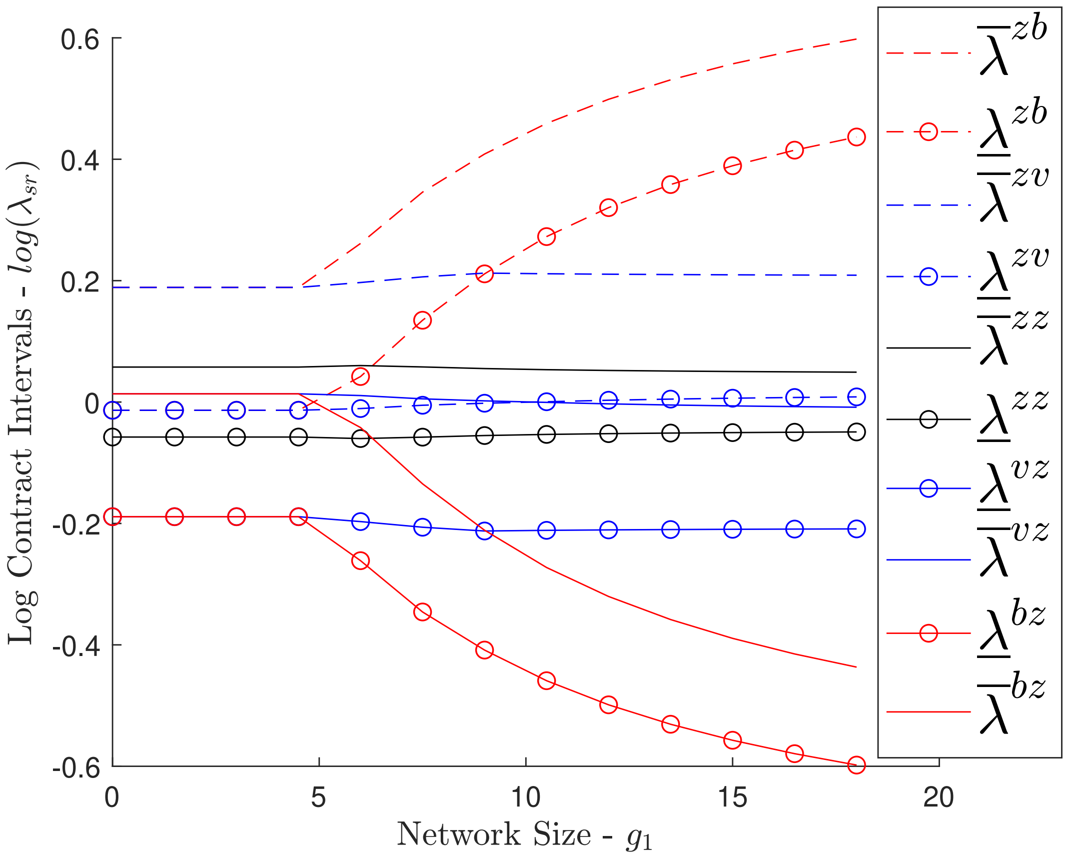

Note: Contract interval solutions as a function of network size with log utility (i.e., $u_1() = u_2() = ln()$. Logged values of λ on the y-axis and network size on x-axis. Contract intervals in state zb increase when $g_1 > 5$ and no longer overlaps with zz when $g_1 > 5$. Furthermore, it is non-overlapping with zv when $g_1 > 7$. The "first-best" insurance contract (consumption equal to a stationary share of aggregate output) is only available when network size is less than five.

FIGURE A.1: Contract Intervals

surplus — in other words, in the initial state, the contract is specified by $\lambda(h_1) = \underline{\lambda}(s_1)$ since household 1 extracts the highest surplus when household 2's surplus is set to zero. When $g_1 < 5$, the warm glow altruism in state zb and zv is the same ($\gamma_1 = 0.5$) resulting in the same levels of discounted lifetime expected utility. After this point, discounted utility decreases at a faster rate as network size increases in the zb state relative to the zv state. This necesarily means that households receiving public income shocks will reachh the shutdown point prior to households receiving private income shocks.

The intuition is as follows: as household 1's network size increases from zero, it will receive more requests for transfers when its income is publicly observed and will derive less

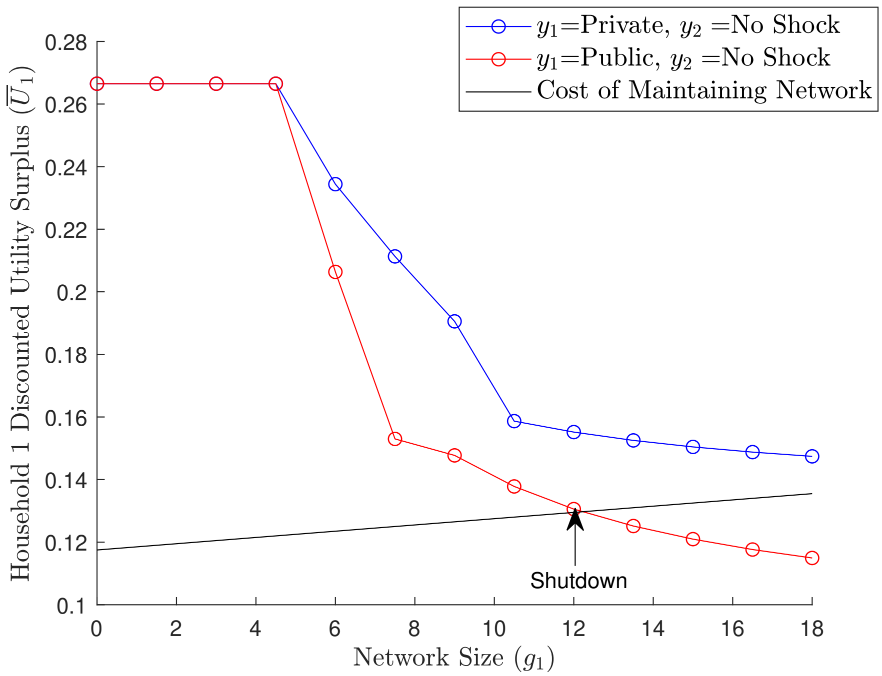

Note: Discounted lifetime expected utility for household 1 when the initial state is zv vs. zb and when household 1 takes all available surplus from the transfer arrangement. Utility values in the zb state decrease at faster rates than state zv throughout. The cost of maintaining each network tie, arbitrarily set to $\alpha(g_1) = .1175 + .001g_1^{1.2}$ is increasing in network size and intersects with $\overline{U}_1^{zb}$ at a threshold of $g_1 = 12$. Beyond this point, household 1 shuts down all gift transactions when it reaches the zb state. We plot $\overline{U}_1^{zb}$ without the possibility of shutdown; however, surplus utility is $\overline{U}_1^{zb} = 0$ whenever $g_1 > 12$.

FIGURE A.2: Discounted Lifetime Expected Utility

(altruistic) utility from each transfer it allocates to other households as a result. As such, when the household reaches a particular network size (greater than five in our example), it will want to hold onto an increasing share of the available surplus. Household 2 will only consent to this arrangement if it receives a higher share of the surplus in future states, which reduces the share of the discounted expected utility surplus received by household 1. In other words, when household 1's surplus demands are too large, household 2 will not consent to giving household 1 the entire surplus in zb if it will not receive the entire surplus in zv. Figure A.2 also includes a plot of the cost of maintaining one's gift-giving ties, $\alpha(g_1)$. Once discounted utility falls beneath this line, household 1 will choose to no longer participate in inter-household transfers and will shut down all giving to other households when state zb is realized.

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 4.

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 C.10 presents results that confirm the shutdown hypothesis with the number of gifts given as the dependent variable.[36] First, the effect outlined in the above paragraph does not take place. Instead, the negative sign of β_{bg} is maintained and the coefficient in front of $\overline{\text{Public}}^{2}_{it}{}'$ is weakly positive, which runs counter to the idea that households coordinate with one another to reduce giving. 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 $\overline{\text{Public}}^{2}_{it}{}'$ with network size. Column 2 shows this is not the case. Instead, there is a positive and insignificant effect. Column

[36]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.

3 examines whether public lottery winners change giving patterns when the second-degree network also wins more in public than average. 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.[37]

The Social Cost of Shutting Down. The shut-down hypothesis implies that households choose to exit reciprocal transfer agreements when network size is too large. Of course, 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. In our case, we expect that households with large networks who also won the public cash prize in the past will be less likely to receive transfers from their network subsequent to their public cash winnings. Table C.11 tests this hypothesis by estimating a variation of equation 17 in which the dependent variable is gifts received, which we regress on a binary variable equal to one if the household ever won a public or private prize *in any round prior to round t.*

The estimation results mimic those in Table 4. Households who win the public prize and have large networks are less likely to receive future transfers from their network ($\beta_{bg} < 0$). On the other hand, households with smaller networks who win the public lottery become more likely to benefit from reciprocity in future rounds ($\beta_b > 0$), presumably because the early-round recipient demonstrated fidelity to the informal contract, thereby earning reciprocal

[37]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 $\overline{\text{Public}}^2_{it}{}'$ in columns 2 and 3 of Table C.10.

treatment subsequently. Strikingly, the shut down point ($-\frac{\beta_b}{\beta_{bg}}$) is between 14 and 16 across the three columns. In figure D.2, this approximates the point at which the public prize *decreases* gift-giving relative to the status quo.

The weak positive result on private winnings ($\beta_v = 0$ not rejected) suggests that households who give gifts from private winnings do not necessarily see their gifts reciprocated in future rounds β_v. This is expected in a setting 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.

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

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 4 and remove instances of the largest public winnings (households winning 70 GH¢in the public lottery) and show that the shut-down effect (negative β_{bg}) remains, especially in the number of gifts given (Table C.12). Default as a result of unusually large windfall gains is unlikely to be a driver of the shut down effect.

APPENDIX C: APPENDIX TABLES

TABLE C.1
Individual Summary Statistics

	N	Mean	Sd	Percentile 5	Percentile 95
Fixed Over Time:					
HH size	606	5.09	2.23	2	9
Gift Network Size	597	9.94	10.10	0	31
Gifts and Loans (last 2 months):					
N Gifts Given	2,983	0.82	1.37	0	4
N Gifts Received	2,983	0.30	0.80	0	2
Total Value of all Gifts Given	1,175	20.02	75.25	1	66
Total Value of all Gifts Received	542	12.58	35.75	1	35

Intrahousehold Differences in Food Expenditure: Head vs. Spouse

Food Consumption (last month):	**Head**	**Spouse**	**HH Total**	**SD**
PC Food Consumption	10.43	16.71	26.45	20.77
PC Purchased Food	3.11	15.86	19.42	18.83
PC Home-produced Food	7.78	1.60	8.63	7.98

Note: Gift Network data missing for a subset of observations. N of gifts given/received equal zero if none given/received. Value of gifts contingent on having received at least one. Gift data excludes within-household transfers and exclude all gifts whose destination or origin is outside of the study village. In the bottom panel we report the amount the household head (usually male) reported on monthly per capita (PC) food consumption, the amount that the spouse reported, the total household food consumption, and the standard deviation (SD) of household food consumption. T-tests of equivalent spending between household head and spouse are strongly rejected (P-Value = 0.00 across all categories).

TABLE C.2

Balance Tests Along Baseline Household Statistics

	N Winners		Win-at-all		Win-Private		Win-Public	
	N-No	N-Win	Diff	P-Value	Diff	P-Value	Diff	P-Value
Fixed Over Time:								
HH size	181	134	0.32	0.28	0.51	0.13	-0.00	1.00
N of HH in Solidarity Network	181	134	-0.68	0.56	-1.10	0.40	-0.21	0.87
Gifts Given:								
Number Gifts Given	181	129	-0.06	0.72	-0.06	0.76	0.06	0.75
Number Gifts Received	181	129	-0.06	0.48	-0.05	0.61	-0.06	0.51
Food Consumption (Last Month):								
PC Food	174	124	-1.03	0.76	-3.85	0.30	0.52	0.89
PC Purchased Food	174	124	-0.87	0.79	-2.33	0.52	0.44	0.91

Note: Balance test of round one observations. N Winners separates the sample according to those households that won any type of cash lottery over rounds two through five and those who did not win a lottery. Win-at-all subtracts the average round one responses of lottery winners from the round one responses of those who did not win the lottery (diff = winners minus non-winners)these two categories of households. We test whether observable characteristics are different across groups — P-values represent outcomes of t-tests of equal means across the two groups. Win-private (public) are t-tests of round one differences across households that won the private (public) lottery as compared to households that never won either lottery.

TABLE C.3

Prize Winnings Influence Gift-Giving - Count Data

		Gift-giving
Dependent Variable:		Number
		(1)
Randomized Explanatory Variable		
Value of Private Cash Prize	β_v	0.0844**
		(0.037)
Value of Public Cash Prize	β_b	0.0519
		(0.033)
Household FE		Yes
Round × Village FE		Yes
Observations		1,561

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

TABLE C.4

PRIZE WINNINGS AND GIFT GIVING - INCLUDING LIVESTOCK WINNINGS

		Gift Giving		
Dependent Variable:		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Value of Private Cash Prize	β_v	0.150^{*}	0.130^{**}	0.167^{**}
		(0.083)	(0.054)	(0.067)
Value of Public Cash Prize	β_b	0.00494	-0.0281	0.0607
		(0.085)	(0.065)	(0.080)
Value of Private Livestock Prize		-0.00135	0.0205	-0.00188
		(0.065)	(0.051)	(0.062)
Value of Public Livestock Prize		-0.0511	-0.0479	-0.0562
		(0.082)	(0.064)	(0.071)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-value $H_0 : \beta_v = \beta_b$		0.23	0.07	0.32
P-value $H_0 : \beta_v <= \beta_b$		0.12	0.04	0.16
Left-censored N		946	946	946
N		1,561	1,561	1,561

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. The 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 and Livestock prize are all divided by $10 \in \{0, 1, 2, 3.5, 5, 7\}$. Tobit estimator used in all columns with a lower bound of zero.

TABLE C.5

Attrition Analysis

	(1) OLS	(2) Logit	(3) Probit
Per Capita Food Consumption (100s GH¢)- Round 1	-0.023	-2.067	-0.704
	(0.022)	(2.165)	(0.882)
Total Log Value of Gifts Given (100s GH¢) - Round 1	-0.009	-4.565	-1.346
	(0.009)	(4.625)	(1.427)
Total Number of Gifts Given - Round 1	0.001	0.217	0.073
	(0.003)	(0.206)	(0.079)
N Adults Present in Household	-0.003	-0.419	-0.124
	(0.002)	(0.378)	(0.105)
N of HH in Network (10s)	-0.001	-0.117	-0.026
	(0.003)	(0.314)	(0.104)
N	1,585	1,585	1,585
Clusters	317	317	317
F Statistic	1.18		
Chi Squared Statistic		2.69	2.35
P-Value H_0: Coefficients Jointly Null	0.32	0.75	0.80

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent variable equals 1 if the household-round observation is missing and zero otherwise. There are 24 instances of missing household-round observations out of 1,585 potential gift observations. Standard errors clustered by household. OLS estimator used in column 1, Logit estimator in column 2 and Probit estimator in column 3. All explanatory variables use values measured in round 1 of data collection when available. In a subset of cases, we use data from the earliest available round for information missing in round 1. This table shows that it is unlikely that our sample exhibits attrition bias — we can not reject the possibility that all of the coefficients are jointly null in any of the models.

TABLE C.6

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 D.5). Coefficient estimates represented graphically in Figure 3.

TABLE C.7

Predictors of Gift Network Size

Dependent Variable:	N in Gift Network	N in Gift Network
	(1)	(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 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 5) in round 1. Coefficients estimated using OLS. Robust standard errors in parentheses.

TABLE C.8

TESTING THE SHUT DOWN HYPOTHESIS — NETWORK DEFINED BY NUMBER OF FAMILY MEMBERS

		Gift Giving		
Dependent Variable:		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Value of Private Cash Prize	$\beta_v > 0$	0.171*	0.159*	0.154*
		(0.102)	(0.081)	(0.084)
Value of Private Cash Prize × N_{FAM}	$\beta_{vg} \leq 0$	-0.001	-0.002	0.000
		(0.004)	(0.003)	(0.003)
Value of Public Cash Prize	$\beta_b > 0$	0.150	0.031	0.299***
		(0.107)	(0.085)	(0.089)
Value of Public Cash Prize × N_{FAM}	$\beta_{bg} < 0$	-0.011*	-0.004	-0.018***
		(0.006)	(0.005)	(0.005)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v = \beta_b$		0.89	0.28	0.23
P-Value $H_0 : \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		0.59	0.15	0.59
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.29	0.07	0.67
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.06	0.03	0.02
N at Shut Down		14.21	7.31	16.83
Left-censored N		946	946	946
N		1,561	1,561	1,561

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 family members (related to head or spouse) in the village sample. Sample average is equal to 17. N_{FAM} at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$.

TABLE C.9

TESTING THE SHUT DOWN HYPOTHESIS — NETWORK DEFINED BY PREDICTED GIFT-NETWORK SIZE

		Gift Giving		
Dependent Variable:		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Value of Private Cash Prize	$\beta_v > 0$	0.228*	0.194*	0.171*
		(0.126)	(0.100)	(0.103)
Value of Private Cash Prize × $\hat{N}$	$\beta_{vg} \leq 0$	-0.007	-0.006	-0.001
		(0.009)	(0.007)	(0.007)
Value of Public Cash Prize	$\beta_b > 0$	0.173	0.023	0.421***
		(0.138)	(0.109)	(0.117)
Value of Public Cash Prize × $\hat{N}$	$\beta_{bg} < 0$	-0.018	-0.006	-0.040***
		(0.013)	(0.010)	(0.012)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v = \beta_b$		0.76	0.25	0.11
P-Value $H_0 : \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		0.37	0.09	0.61
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.09	0.03	0.09
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.14	0.25	0.00
$\hat{N}$ at Shut Down		9.39	4.13	10.43
Left-censored Observations		946	946	946
Observations		1,561	1,561	1,561

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 = ∈ $\{0, 1, 2, 3.5, 5, 7\}$. 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 number of family members (related to head or spouse) in the village sample. Sample average is equal to 17. $\hat{N}$ at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$.

TABLE C.10

TESTING THE COORDINATION HYPOTHESIS

Dependent Variable:		Number of Gifts Given		
		(1)	(2)	(3)
Own Winnings and Network Size Interaction Variables				
Private (Own)	$\beta_v > 0$	0.210**	0.245**	0.218**
		(0.099)	(0.103)	(0.099)
Private (Own) × N	$\beta_{vg} \leq 0$	-0.004	-0.003	-0.005
		(0.007)	(0.007)	(0.007)
Public (Own)	$\beta_b > 0$	0.423***	0.416***	0.418***
		(0.123)	(0.129)	(0.122)
Public (Own) × N	$\beta_{bg} < 0$	-0.042***	-0.042***	-0.041***
		(0.010)	(0.010)	(0.010)
Second Degree Adjusted Network Average Winnings ($\overline{\textbf{Private/Public}}_{it}^{2}{}'$)				
$\overline{\text{Private}}_{it}^{2}{}'$		8.934**	9.775***	12.456**
		(3.569)	(3.645)	(6.043)
$\overline{\text{Public}}_{it}^{2}{}'$		1.814	1.745	-4.592
		(2.099)	(2.223)	(3.490)
$\overline{\text{Private}}_{it}^{2}{}' \times$ Private (Own)			-3.560*	
			(1.949)	
$\overline{\text{Public}}_{it}^{2}{}' \times$ Public (Own)			0.399	
			(0.701)	
$\overline{\text{Private}}_{it}^{2}{}' \times$ N				-0.360
				(0.373)
$\overline{\text{Public}}_{it}^{2}{}' \times$ N				0.728**
				(0.303)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.08	0.04	0.10
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.00	0.00	0.00
Left-censored N		946	946	946
N		1,561	1,561	1,561

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. 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 C.11

THE SOCIAL COST OF SHUTTING DOWN

		Receiving Gifts		
Dependent Variable		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Lagged Randomized Explanatory Variables With Network Size Interaction				
Won Private in Past?	β_v	0.160	0.121	0.020
		(0.274)	(0.224)	(0.222)
Won Private in Past? × N	β_{vg}	-0.011	-0.007	-0.010
		(0.019)	(0.016)	(0.016)
Won Public in Past?	β_b	0.576^{**}	0.415^{*}	0.543^{**}
		(0.282)	(0.232)	(0.223)
Won Public in Past? × N	β_{bg}	-0.040^{*}	-0.030^{*}	-0.034^{**}
		(0.021)	(0.017)	(0.016)
Round × Village FE		Yes	Yes	Yes
N at Shut Down		14.29	13.96	15.84
Left-censored Obs.		1,292	1,292	1,292
Observations		1,556	1,556	1,556

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Dependent Variable equals log total value of (cash) gifts received per adult in household in column 1; log average value of (cash) gifts received per adult in column 2; number of (cash) gifts received per adult in column 3. Won Private/Public in Past $\in \{0, 1\}$ indicates whether household won lottery at any point in current or up to past 3 rounds. Tobit estimator used in all columns. N denotes network size.

TABLE C.12

Testing Whether the Shutdown Effect is Due to Default

		Gift Giving		
Dependent Variable:		Value (Total)	Value (Average)	Number
		(1)	(2)	(3)
Value of Private Cash Prize	$\beta_v > 0$	0.302**	0.199**	0.235**
		(0.134)	(0.093)	(0.100)
Value of Private Cash Prize × N	$\beta_{vg} \leq 0$	-0.012	-0.005	-0.005
		(0.009)	(0.005)	(0.007)
Value of Public Cash Prize	$\beta_b > 0$	0.212	0.039	0.474***
		(0.153)	(0.123)	(0.182)
Value of Public Cash Prize × N	$\beta_{bg} < 0$	-0.022*	-0.008	-0.041***
		(0.011)	(0.009)	(0.013)
Household FE		Yes	Yes	Yes
Round × Village FE		Yes	Yes	Yes
P-Value $H_0 : \beta_v = \beta_b$		0.66	0.31	0.26
P-Value $H_0 : \beta_v + \beta_{vg} \times 5 = \beta_b + \beta_{bg} \times 5$		0.37	0.16	0.71
P-Value $H_0 : \beta_v + \beta_{vg} \times 10 = \beta_b + \beta_{bg} \times 10$		0.15	0.07	0.31
P-Value $H_0 : \beta_v + \beta_{vg} \times 15 = \beta_b + \beta_{bg} \times 15$		0.09	0.06	0.01
P-Value $H_0 : \beta_v + \beta_{vg} \times 20 = \beta_b + \beta_{bg} \times 20$		0.11	0.11	0.00
N at Shut Down		10	5	11
Left-censored N		935	935	935
N		1,545	1,545	1,545

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 $10 = \in \{0, 1, 2, 3.5, 5\}$ — we remove instances of lottery winnings equal to 7. Tobit estimator used in all columns with a lower bound of zero. Wald tests of equivalence specified for network size (N) of 0, 5, 10 and 20. N at Shutdown is equal to $-\frac{\beta_b}{\beta_{bg}}$. This table shows that default is unlikely to be a driver of the shutdown effect since the negative coefficient on β_{bg} persists despite the removal of observations when public lottery winnings are equal to 7.

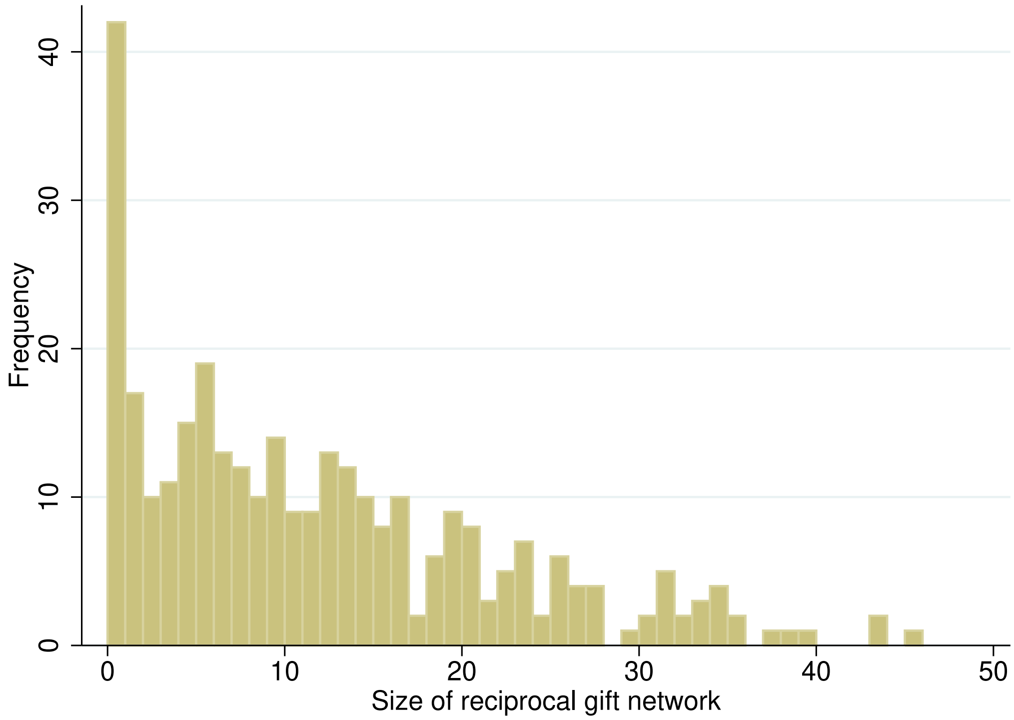

Note: Each bin is one unit wide.

FIGURE D.1: Distribution of Gift Network Size

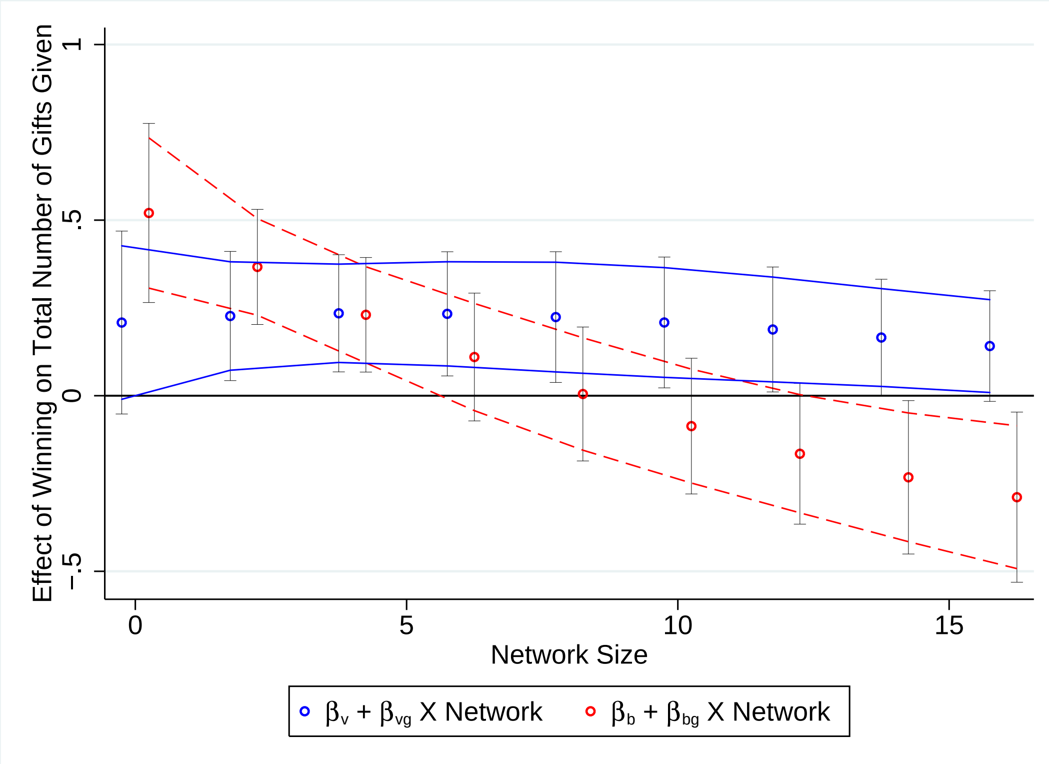

Note: Dependent variable equals number of gifts given. Estimation of Model 17 with the inclusion of 2nd and 3rd order polynomial interactions on network-size variable (with respective coefficients $\beta_{bg^2}, \beta_{vg^2}, \beta_{vg^3}$ and β_{bg^3}). Dots represent point estimates of $\beta_b + \beta_{bg} \times N + \beta_{bg^2} \times N^2 + \beta_{bg^3} \times N^3$ (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 D.2: Shut-down Hypothesis on Number of Gifts Given

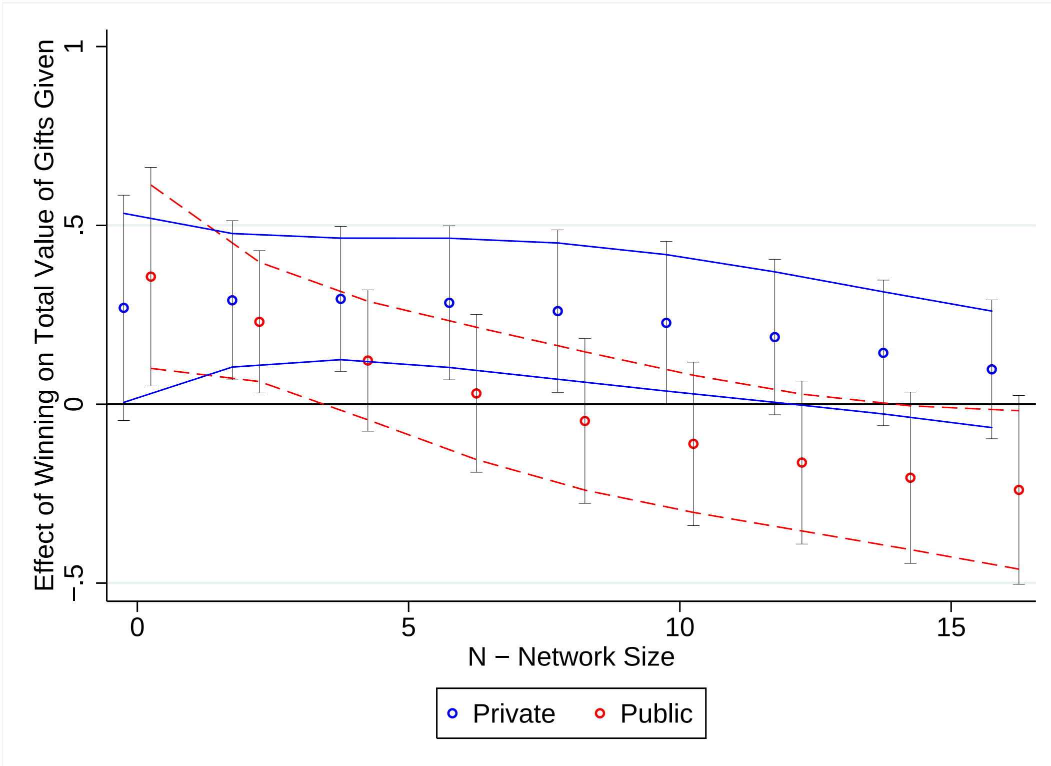

Note: Dependent variable equals log total value of gifts given. Estimation of Model 17 with the inclusion of 2nd and 3rd order polynomial interactions on network-size variable (with respective coefficients $\beta_{bg^2}, \beta_{vg^2}, \beta_{vg^3}$ and β_{bg^3}). Dots represent point estimates of $\beta_b + \beta_{bg} \times N + \beta_{bg^2} \times N^2 + \beta_{bg^3} \times N^3$ (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 D.3: Shut-down Hypothesis on Total Value of Gifts Given

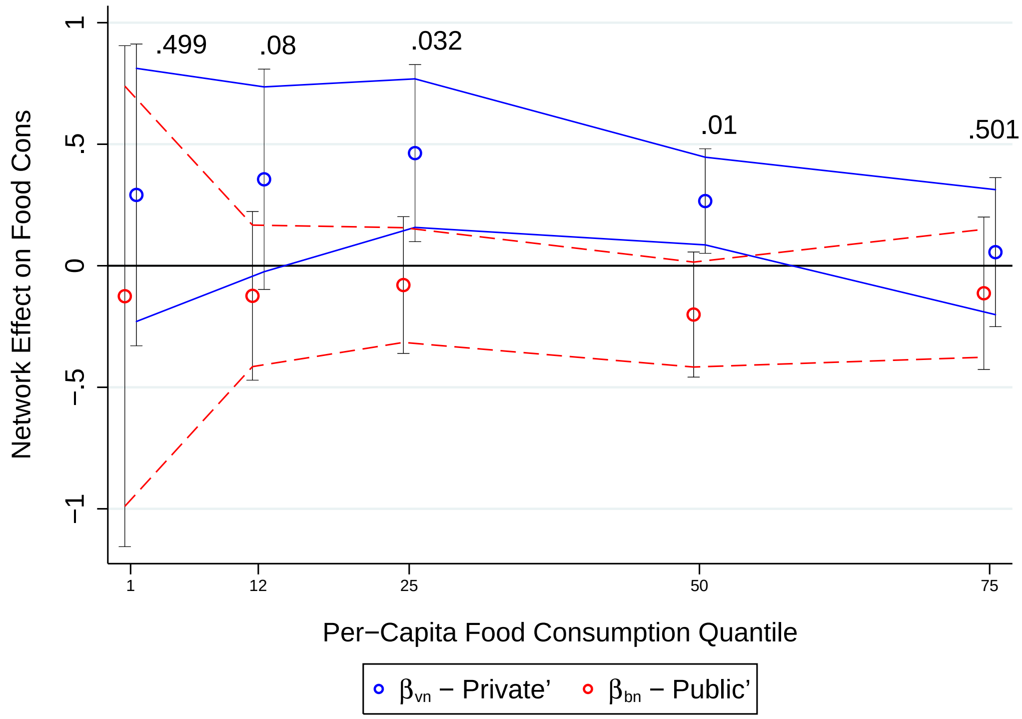

Note: Results of a simultaneous quantile regression at 1st, 12.5th, 25th, 50th, and 75th quantiles bootstrapped over 1,000 iterations. Dependent variable is log home-produced per capita food consumption over the last month. Quantiles represented on the x axis. Blue dots (lines) show the coefficient estimates (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 D.4: Effect of Unadjusted Network Winnings on Food Consumption by Quantile

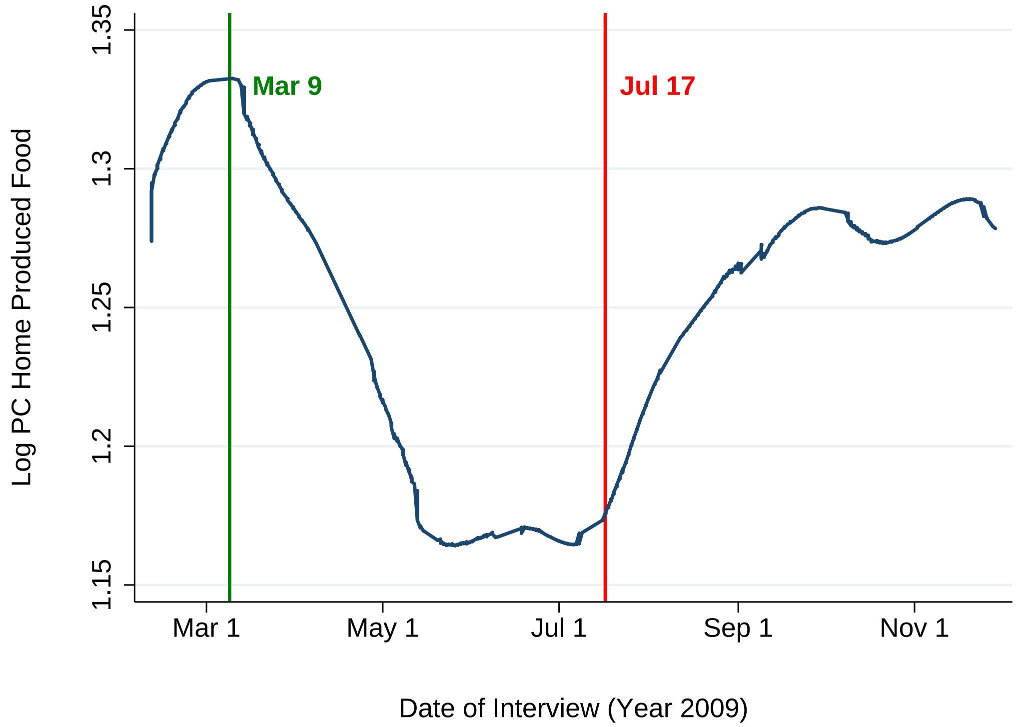

Note: Log home-produced per capita food consumption over the last month on the y axis. Date of interview on the x axis. Blue line shows the lowess smoothed curve by date with a bandwidth of 0.4. The peak of the average home produced food consumption is around March 14. After this point, average home produced food consumption begins to decrease until its nadir on around July 12. We include all observations between the vertical green line and vertical red line in our quantile regression analysis in Section 5. 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.

FIGURE D.5: Home Produced Food Over The Course of the Year