

ALTRUISM, INSURANCE, AND COSTLY SOLIDARITY COMMITMENTS¹

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
MIT

Christopher B. Barrett
Charles H. Dyson School
Cornell University

Eleonora Patacchini
Department of Economics
Cornell University

Thomas Walker
World Bank

¹Corresponding Author: vnourani@mit.edu. Seminar participants at Cornell University, Einaudi Institute for Economics and Finance, London School of Economics, the University of Copenhagen, Massachusetts Institute of Technology, and the AAEA annual meetings provided excellent feedback during early presentations 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, commonly framed as self-enforcing informal insurance contracts (Fafchamps, 1992; Coate and Ravallion, 1993; Townsend, 1994; Ambrus, Mobius, and Szeidl, 2014). This contractual framing of inter-household transfers helps explain risk pooling among households whose objective is to smooth consumption (Coate and Ravallion, 1993; Townsend, 1994). An important implication of this framework for public policy is that social networks should (at least partially) correct targeting errors in publicly observable transfer programs, as non-recipients who have suffered adverse shocks will enforce their claims on recipients within their network to share any windfall gains (Angelucci and De Giorgi, 2009).

A related but distinct literature, however, 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 (Platteau, 2000; Sen and Hoff, 2006; Jakiela and Ozier, 2016; Squires, 2017). This contrary perspective raises important questions about prospective limits to the value of extensive social networks.

The first step towards reconciling the informal insurance and social taxation literatures is the recognition that both depend fundamentally on the observability of income, or at least of shocks to income.² The motives for sharing are different under these frames. The effectiveness of self-enforcing insurance among purely self-interested agents depends upon

²More specifically, they rely on non-uniform shocks across households within the network so that the 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.

each party’s ability to monitor others’ income shocks so as to enforce the contract. Similarly, people can only tax the public portion of others’ income streams. In this paper we take a step towards reconciling these literatures by teasing out the implications of publicly observable (hereafter ‘public’) versus unobservable (i.e., ‘private’) income shocks.

Both frameworks’ dependence on public observability of income shocks implies two maintained, but testable hypotheses. First, public income shocks should lead to inter-household transfers, whether due to social taxation, informal insurance contracts, or both. Therefore, one should be able to reject the null hypothesis that public income shocks have no effect on inter-household transfers in favor of the one-sided alternate hypothesis of positive impacts. Second, private 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 hypothesis that private income shocks have zero effect is a low-power test of the foundational public observability hypothesis.

We show that neither of the two maintained hypotheses hold 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. Contrary to the central predictions of standard informal insurance or social taxation models, 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 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.

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 from publicly unobservable income, 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.

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. These two modest, realistic tweaks let us also address the first hypothesis, concerning giving from publicly observable income windfalls. They also allow us to draw out several other, more subtle, testable hypotheses that match our data.

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 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. 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, altruistic individuals would like to target their giving toward the neediest members of their social network. But when stochastic income realizations are publicly observable, the 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. This reinforces the shutdown hypothesis. And it implies that in the presence of altruism, private rather than public giving 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.

³This really just requires that the marginal returns to giving diminish faster than the marginal costs of giving, not that the costs be strictly constant.

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) altruistic 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. The social solidarity network is multi-functional, (incompletely) pooling income risk across a network so as to (partially) smooth consumption as an insurance contract would, while also accommodating the social taxation pressures of network members, and at the same time mediating altruistic transfers towards the least fortunate members of the network.

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. This indicates more targeted giving when altruistic behavior dominates because the unobservability of one's winnings attenuates network demand due to social taxation and/or informal insurance contract enforcement. Furthermore, this provides strong support for the existence of altruistic motives in social solidarity networks. In the absence of altruistic preferences, one is hard-pressed to provide reasonable motives for sharing unobservable, private winnings (we consider, and refute, some of these alternative motives in section 6).

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 larger for public than for 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.

Finally, we show that, within these gift networks, limited risk pooling holds. We can easily reject both the null hypothesis of full risk pooling and the null of no risk pooling. Cumulatively, these results indicate that inter-household transfers reflect a blend of insurance and altruism motives, mediated by social taxation pressures.

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. [Angelucci, De Giorgi, and Rasul \(2017\)](#) show that Progresa transfers in Mexico are pooled by family networks to finance consumption and investment and [Advani \(2017\)](#) shows using experimental data from Pakistan that poverty traps can exist at the network level. [Simons \(2016\)](#) shows that community targeting of a social safety net program is pro-poor relative to centralized targeting. These results suggest that communities in many parts of the world have intimate knowledge of their members' needs and can potentially allocate resources more efficiently than state institutions ([Alderman, 2002](#); [Bowles and Gintis, 2002](#)). Although observability of income is essential in informal insurance arrangements among purely self-interested agents, observability may impede altruistic agents' ability to focus their giving on the most needy as they are compelled to respond to social taxation or informal insurance demands from the less needy within their network.

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

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 Accra for work. The sample consists of approximately 70 households from each of the four communities.⁴ Individuals in the sample include the household head and his spouse.⁵ 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 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. Approximately 42 percent of individuals and 62 percent of households won at least one prize over the course of the year. Within each village and round, ten of the prizes were cash; the other ten were in the form of livestock. 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. The values of the prizes varied from GH¢10 to GH¢70 as described in Figure 1.⁶ The prizes were of a substantial size - the largest prize is equivalent

⁴The survey was part of a three-wave panel, the first two waves having been conducted in 1997-98 (e.g., in Conley and Udry (2010)) and 2004 (Vanderpuye-Orgle and Barrett, 2009). Slightly more than half of the 70 households were part of the initial 1997-98 sample, and the rest were recruited in January 2009 using stratified random sampling by the age of the household head: 18-29, 30-64, 64+. the shares of households whose head was in each of these age categories corresponded to the community’s population shares. In the original sample, and in the 2009 re-sampling, we selected only from the pool of households headed by a resident married couple. However, we retained households from the 1997-98 sample even if only one of the spouses remained.

⁵Some 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.

⁶During 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

to a month’s worth of food consumption for an average household with five 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). A village meeting was held in a central area of the community, and all respondents were invited to attend. A small amount of free food and drink was provided as an incentive to come. Attendance at the meetings was generally around 100 people; roughly half of the respondents appeared for each public meeting.⁷ There were usually a number of non-respondents at these meetings as well. At each gathering we thanked the participants for their continued support. We explained that respondents had a chance to win one of 20 prizes that day, framing the prizes as a gratuity for their participation in the survey.⁸ We then proceeded to draw winners for the ten public prizes (without replacement) from a bucket containing the names of the survey respondents. A village member not in the sample was chosen by the villagers to do the draw, in order to emphasize that the outcomes were random. Each winner was announced to the group, and asked to come forward to receive their prize. The prizes were announced and displayed clearly before being awarded. Respondents who were absent at the time of drawing were called to pick up their prize in person, if possible. Unclaimed prizes were delivered in person to the winner after the lottery.

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 prize differed according to the type of livestock: 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.

⁷Around 125 of the 150 respondents in each community appeared for the privately revealed lottery, some of them arriving before or after the public meeting.

⁸Following 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).




FIGURE 1: Experimental Data: Lottery Payouts

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 1.⁹ Those who drew winning tokens were informed immediately that they had won a prize, which was identified to them, and were told that they did not have to tell anyone else that they had won. We emphasized that the survey team would not divulge the identities of winners who won in private. Cash prizes were given to the winners immediately and winners 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

⁹Care was taken to shuffle the bottle caps after each draw, and to prevent respondents from seeing into the bag. If a respondent drew more than one bottle cap, those caps were shuffled and the respondent was asked to blindly select one of them. Respondents were shown a sheet relating the tokens to the prizes (See Walker (2011)). At the conclusion of the day, tokens that had not been drawn were counted and the remaining prizes allocated randomly among the non-attending respondents using a computer. There were usually 25-30 non-attendees and less than three prizes remaining.

between February and November.¹⁰ 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.¹¹ 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 location and relationship to the respondent. The value of the gift given and an estimated value for in-kind gifts were also recorded. We focus 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 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 in [Castilla and Walker \(2013\)](#). With respect to gift received, we are interested in gifts from others who are potential winners of 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: 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 roughly five 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

¹⁰For details regarding interview timing and survey instruments, see [Walker \(2011\)](#).

¹¹There 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.

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	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.

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 B.2 presents balance tests conducted on variables collected at baseline according to whether one member of the household won any of the public or private lottery at any point over the course of the year. 119 of the households in the study are thus in our “treatment” group while the remaining 190 did not win a cash prize. We also separate the test according 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 5 percent level and another is significant at the 10 percent level. For the others, balance cannot be rejected at the 10 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: either the number of round t gifts distributed by household i , the average amount per gift given, or the total amount given, which is simply the product of the first two outcomes. The two core hypotheses can be tested using the following regression:

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

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.¹²

Table 2 reports the estimation results of model 1 with three different outcome variables:

¹²The number of gifts given is integer-valued, so we also estimate a Poisson count data estimator to estimate the coefficients of interest using this dependent variable. The results are reported in appendix Table B.3 and they remain qualitatively unchanged.

TABLE 2
PRIZE WINNINGS AND GIFT GIVING

Dependent Variable:	Gift Giving		
	Value (Total) (1)	Value (Average) (2)	Number (3)
Randomized Explanatory Variables			
Value of Private Cash Prize β_v	0.149** (0.069)	0.129** (0.055)	0.166*** (0.057)
Value of Public Cash Prize β_b	0.00789 (0.071)	-0.0265 (0.057)	0.0639 (0.058)
Household FE	Yes	Yes	Yes
Round \times Village FE	Yes	Yes	Yes
P-value: $\beta_v = \beta_b$	0.15	0.05	0.21
P-value: $\beta_v \leq \beta_b$	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 B.3 reports estimates of the number of gifts given using a Poisson estimator with qualitatively similar results as those in column 3.

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.

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 the canonical models 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 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) altering the number of gift requests one receives when one’s income is publicly vs. privately revealed, reflecting social taxation pressures. These seemingly innocuous adjustments, grounded in our observation of solidarity network activity and our field research subjects’ descriptions, generate more nuanced predictions that do not imply 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 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, $i = \{1, 2\}$ receiving stochastic incomes, $y_i(s_t) \geq 0$ that depend on the state, s_t , realized in period t — a sequence of the state history is characterized by $h_t = \{s_1, s_2, \dots, s_t\}$.¹³ 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.

¹³The assumption of stochastic exogenous income is reasonable in our empirical context since we distribute cash prizes randomly across the sample.

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 are not modeling network size as a choice variable in this paper. We acknowledge, however, that there are implications for endogenous network choice that emerge from the principles reflected in our enhanced model. We 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 household consumption. Single-period utility for household 1 is reflected in the following equation:

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

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

We 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 in the following manner:

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

where $\mathbb{1}(\cdot)$ is an indicator function equal to one when there is a transfer between households 1 and 2, and $\bar{\gamma}_1$ 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 is a decreasing function of the total number of household 1’s period t gift obligations, $g_1 \cdot p_1(s_t)$. This reflects the idea that warm glow 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. 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.¹⁴ To set up the household’s problem, we define $U_1(h_t)$ as 1’s expected discounted utility gain from the

¹⁴Households in [Foster and Rosenzweig \(2001\)](#) revert to a sequence of history-dependent Nash equilibria (SHDNE) in which transfers are maintained even when a household defaults from the contract. Such an environment is not crucial for the type of analysis we conduct in our study. Nevertheless, appendix section [A](#) shows how one can adapt our model to reflect such SHDNE default transfers.

risk-sharing contract with 2 relative to a no-transfer rule after history h_t :

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

where δ represents the dynamic discounting factor. $\alpha_1(g_1)$ 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: $U_1(h_t) \geq 0$ and $U_2(h_t) \geq 0$. Together, the economic environment, payoffs and transfer decision represent a simultaneous game in which agents seek to find a contract that can be implemented in the presence of limited commitment and no external enforcement mechanism.

Limited Commitment Contract Solution. Following [Foster and Rosenzweig \(2001\)](#) and [Ligon, Thomas, and Worrall \(2002\)](#), the solution to the utility maximization problem will be a dynamic program in which the current state is given by s out of the set of all states ($s \in \{1, 2, \dots, S\}$), and targeted discounted utility gain for household 2, U_2^s , is given.¹⁵ 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

¹⁵ U_2^s is defined by equation 21 when all subscripts with 1 are replaced with a 2 and vice versa.

programming problem as

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

subject to

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

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

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

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

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

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 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 op-

timal contract:

$$(11) \quad \frac{u'_1(y_1(s_t) - \tau(h_t)) + \gamma_1(g_1(h_t))u'_2(y_2(s_t) + \tau(h_t))}{u'_2(y_2(s_t) + \tau(h_t)) + \gamma_2(g_2(h_t))u'_1(y_1(s_t) - \tau(h_t))} = \lambda + \frac{\psi_2 - \psi_1}{u'_2(y_2(s_t) - \tau(h_t))}$$

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

$$(13) \quad \lambda = -U_1^{s'}(U_2^s).$$

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.¹⁶ For each state s , there is a history independent interval $[\underline{\lambda}_s, \bar{\lambda}_s]$ 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


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

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

Figure 2 describes the intuition behind this contract structure using a stylized example. Suppose that in an initial period, t , a state is realized in which household 1 receives income $y_1(s_t) = 2$ and household 2 receives $y_2(s_t) = 1$.¹⁷ If the two households follow the contract structure in equation 14, then each household will weigh participation in risk sharing against

¹⁶For 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).

¹⁷In later simulations, this income combination will be referred to as state zv



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

FIGURE 2: Contract Intuition

the payoff received when they default from such a contract. Household 2 will only consider this contract if $\lambda(h_t)$ is greater than $\underline{\lambda}_{zv}$ — the point at which household 2 is indifferent between defaulting and participating in the risk-sharing contract (discounted utility surplus equal to zero). Household 1 will have a similar payoff structure when $\lambda(h_t) = \bar{\lambda}_{zv}$. Both households will prefer risk sharing if they can settle on a dynamic contract between these two numbers. Suppose the realized state in period $t + 1$ is zz , where $y_1(zz) = y_2(zz) = 1$. If altruistic preferences (and discount rates) are such that the contract intervals for the realized state in $t + 1$ does not overlap with the state in t (left panel in Figure 2), the surplus will be divided according to $\lambda(h_{t+1}) = \underline{\lambda}_{zz}$. If the contract intervals do overlap, then, $\lambda(h_{t+1}) = \bar{\lambda}_{zv}$. Notice that this results in a division of the surplus in which both households strictly benefit relative to default (e.g., greater consumption smoothing through risk pooling).

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).¹⁸ Explicitly, here we are assuming that the prize-winning household receives a higher income than the non-prize winning household and the prizes are equal in value:

Assumption 1 (Prize-winners Have Higher Incomes)

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

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

Assumption 2 (Observability of Income)

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

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

¹⁸There 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.

the amount transferred.


4.1. Model Implications

Given the complexity of the state space, it is not possible 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.1](#). 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.

These simulations lead to a set of formal empirical predictions. The first is the “Shut down Hypothesis.”

Prediction 1 (The Shut down 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.*

Figure [3](#) 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 . However, as the network size increases, transfer amounts start to decrease until they fall to zero at



Note: This figure represents transfer amounts τ^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 3: Amount of Transfer by Network Size

the shutdown threshold and beyond.¹⁹ This relationship leads to two additional empirical implications:

Prediction 2 (Privately Revealed Prize \rightarrow 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.*

¹⁹Note 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 arises from the insurance contract between agents. Likewise, our two agent model differs from network models that predict that larger networks negatively affect outcomes because network size is negatively associated with network closure, and thus with trust that enhances cooperative behavior (Coleman, 1990; Allcott, Karlan, Mobius, Rosenblat, and Szeidl, 2007)

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

The above two predictions also imply that the total value of gifts out of households who win publicly revealed prizes are higher than the total value of gifts given from other households prior to the household reaching its shut down threshold. This is easily shown by multiplying the average transfer value by the number of gift obligations in period t (see appendix Figure C.3 for a graphical representation). The prediction can be stated as:

Prediction 4 (Prior to Shut Down \rightarrow 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 higher when the peer household wins a private lottery, it is likely that the effect will only be observed in such a state. Furthermore, since household 1's marginal utility is decreasing in household 2's consumption, we should see stronger and more progressive patterns of gift giving through the private lottery when the income gap between households 1 and 2 is large. It is straightforward to show via simulation that average transfer sizes increase as the gap between 1 and 2's per-period income increases.²⁰ This leads to the final prediction:

Prediction 5 (Consumption Increasing in Others' Winnings) *A household's per capita consumption increases in its network's average private lottery winnings. It may be an increasing function of its network's public lottery winnings if its peers do not experience a shut down*

²⁰Similarly, 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.

in giving (i.e., peers have sufficiently small gift giving networks).

5. EMPIRICAL INVESTIGATION

The model implications 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 in turn below. Then, we describe the estimation procedures used to test the predictions of the enhanced model.

5.1. Additional Data

Social Networks. 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.

In our model, we assume that instances of bidirectional inter-household transfers are largely motivated by altruism. 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.²¹

Consumption. The expenditure module asked detailed information on the quantities and values purchased of a long list of items with broad categories 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 B.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.

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.²² We do this by taking the household sum of all expenditures reported by the individuals who incurred the expenditure.²³ 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 expendi-

²¹Consider 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. Formally, and abusing notation slightly, 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).

²²For 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.

²³If 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.

tures 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.

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

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

where $\text{Private}_j \in \{0, 10, 20, 35, 50, 70\}$ 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. The theoretical model suggests that the degree of giving between, say, household 1, the one that receives the positive income shock, and household 2, the household receiving the transfer, also depends on household 1’s network size. The above definition of network average, however is calculated only using household 2’s network. 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

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

The fraction in the numerator represents the weight placed on each household j ’s lottery

TABLE 3
HOUSEHOLD SUMMARY STATISTICS FOR THE ENHANCED MODEL

				Percentile	
	N	Mean	Sd	5th	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.

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. 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.

5.2. Analysis

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:

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

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.²⁴ Time-varying unobservable characteristics of household i are represented by ϵ_{it} .

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 Model 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

²⁴Network 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.

TABLE 4
TESTING THE SHUT DOWN HYPOTHESIS

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

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 give-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.

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 3.²⁵

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.²⁶

Together, the results associated with the first four predictions suggest a clear pattern of behaviors that emerge following private vs. 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

²⁵This relationship is demonstrated graphically with the aid of a third-order polynomial estimation of Model 18 in Figure C.4.

²⁶This relationship is also demonstrated graphically in Figure C.5.

similar network size who win privately revealed transfers. Once network size reaches around 10, one unit larger than the median network size, publicly revealed prizes no longer have any effect whatsoever on giving. Figures C.4 and C.5 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.²⁷ 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 capita food consumption, our preferred proxy for consumption in these data. We test this using the following equation:

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

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.²⁸ 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

²⁷The number approximates 15, which translates to the 70th percentile of gift-network size across all villages.


²⁸We repeat the analysis for $\overline{\text{Private}}_{it}$ and $\overline{\text{Public}}_{it}$ in appendix Figure C.7 with qualitatively similar, but noisier, results.

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.²⁹

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 4 (appendix Table B.4 shows estimation results for each quantile). The lower the per capita food consumption, the larger is the adjusted network average effect of private lottery winnings on food expenditures. In Figure 4, 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.

²⁹Appendix Figure C.6 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 $H_0 : \beta_{vn} = \beta_{bn}$.

FIGURE 4: Effect of Network Winnings on Food Consumption by Quantile

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:

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

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). In other words, we predict β_{vF} to be positive.

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 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 winnings. However, we do

TABLE 5
DYADIC REGRESSIONS

Dependent Variable:		Gift Giving Within Dyad: From i to j			
		Amount (1)	Number (2)	Amount (3)	Amount (4)
(Food _{it} – Food _{jt})	γ_F	0.073 (0.204)	0.029 (0.106)		
Network Size	γ_g			-0.036 (0.027)	-0.017 (0.018)
Randomized Explanatory Variables With Interactions					
Value in Private	β_v	0.182 (0.153)	0.136* (0.078)	0.318 (0.235)	0.239 (0.157)
Value in Private \times (Food _{it} – Food _{jt})	β_{vF}	0.305** (0.127)	0.117** (0.058)		
Value in Private \times N	β_{vg}			-0.005 (0.009)	-0.007 (0.009)
Value in Public	β_b	-0.286 (0.265)	-0.234 (0.166)	0.177 (0.399)	0.341** (0.164)
Value in Public \times (Food _{it} – Food _{jt})	β_{bF}	-0.098 (0.064)	-0.055 (0.042)		
Value in Public \times N	β_{bg}			-0.034 (0.025)	-0.044*** (0.016)
Round \times Village FE		Yes	Yes	Yes	Yes
All Dyads Included		No	No	No	Yes
P-value: $\beta_v = \beta_b$		0.12	0.05	0.76	0.64
P-value: $\beta_{vF} = \beta_{bF}$		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 $\in \{0, 1, 2, 3.5, 5, 7\}$. Food_{it} – 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.

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 — they differ in that column 4 includes all out-of-network households while column 4 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.

The Social Cost of Shutting Down. The shut-down condition implies that households are choosing 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 punishment for 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 6 tests this hypothesis by estimating a variation of Model 17 in which the dependent variable is gifts received. Here, the independent variable we are regressing against is 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 treatment subsequently. Strikingly, the shut down point ($-\frac{\beta_b}{\beta_{bg}}$) is between 14 and 16 across the three columns. In figure C.4, 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 house-

TABLE 6
THE SOCIAL COST OF SHUTTING DOWN

		Receiving Gifts		
Dependent Variable		Value (Total) (1)	Value (Average) (2)	Number (3)
Lagged Randomized Explanatory Variables With Network Size Interaction				
Won Private in Past?	β_v	0.160 (0.274)	0.121 (0.224)	0.020 (0.222)
Won Private in Past? \times N	β_{vg}	-0.011 (0.019)	-0.007 (0.016)	-0.010 (0.016)
Won Public in Past?	β_b	0.576** (0.282)	0.415* (0.232)	0.543** (0.223)
Won Public in Past? \times N	β_{bg}	-0.040* (0.021)	-0.030* (0.017)	-0.034** (0.016)
Round \times 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.

holds 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.

Information Hypothesis. One competing explanation for our results is that households who win the private prize cannot conceal this fact from those who are close to them, such as non-co-resident family members within the village. It is unlikely this is the case since within-family food consumption is likely to be correlated (and hence Prediction 5 would not

TABLE 7
GIVING PRIVATE LOTTERY WINNINGS TO FRIENDS, NOT FAMILY

Dependent Variable:		Value of Gifts Given (Average)		
Gifts directed to:		All Family	Direct Family	Village Friends
		(1)	(2)	(3)
Randomized Explanatory Variable With Network Size Interaction				
Won Private Cash Prize	β_v	-0.298 (0.726)	-1.065 (0.828)	0.875** (0.431)
Won Public Cash Prize	β_b	1.912*** (0.686)	2.029*** (0.652)	1.287*** (0.491)
Won Private Cash Prize \times N	β_{vg}	0.0237 (0.044)	0.0442 (0.046)	-0.0157 (0.029)
Won Public Cash Prize \times N	β_{bg}	-0.120** (0.051)	-0.101** (0.049)	-0.118** (0.048)
Round \times 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.

have been confirmed).³⁰

Nevertheless, we explore this possibility in Table 7, differentiating gifts-given according to links with varying likely quality of information about recipient households. We again estimate Model 17 where the dependent variable is the log value of gifts given to all 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

³⁰Furthermore, using the same experiment, [Castilla and Walker \(2013\)](#) show that even spouses did not necessarily know whether the other won a private prize.

no information story to explain the patterns we observe in the data.

Precautionary Savings. Another competing 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 – essentially as callable deposits that can be withdrawn in future periods. If this is the 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.

Test of Full Risk Pooling Having established that social solidarity networks seem to serve altruistic purposes and be subject to social taxation pressures, we conclude this section by testing whether they 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 8
TESTS OF RISK-SHARING

Dependent Variable:	$\Delta \log$ (PC Food)					
	G (1)	F (2)	$G \not\subseteq F$ (3)	$F \not\subseteq G$ (4)	$G \cap F$ (5)	$\not\subseteq (G \cup F)$ (6)
First Difference of Network Average Per Capita Food Consumption						
$\Delta \log(\text{Network PC Food})_{it}$	0.306*** (0.087)	0.328*** (0.098)	0.102 (0.077)	0.034 (0.063)	0.257*** (0.078)	0.022 (0.224)
Randomized Explanatory Variables						
Value of Private Cash Prize	-0.001 (0.010)	0.011 (0.015)	0.002 (0.011)	0.013 (0.014)	0.002 (0.010)	0.007 (0.013)
Value of Public Cash Prize	0.006 (0.012)	0.007 (0.011)	0.014 (0.013)	0.004 (0.011)	0.008 (0.013)	0.004 (0.011)
$\overline{\text{Private}} \text{Network}_{it}$	0.005 (0.027)	0.057 (0.043)	-0.012 (0.030)	0.025 (0.021)	0.014 (0.023)	-0.320** (0.156)
$\overline{\text{Public}} \text{Network}_{it}$	-0.006 (0.032)	-0.001 (0.021)	0.016 (0.022)	0.006 (0.019)	-0.038 (0.031)	-0.077 (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.

Table 8 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 7 strongly suggest that gifts to village friends - rather than to family - are driven primarily by altruistic motives. Columns 4 through 6 look at three more combinations of gift vs. family networks and conclude that the network with the highest degree of insurance-related sharing corresponds to those networks that include family members with whom one has a prior gift exchange relationship.

Meanwhile, the respondent's own winnings, whether private or public, and the average winnings within one's network are statistically insignificantly related to a respondent's consumption volatility once one controls for consumption volatility within one's network, consistent with the altruism in networks model of [Bourlés, Bramoullé, and PerezRichet \(2017\)](#). 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. CONCLUSION

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. Yet existing models do not fully accommodate that multi-functionality and thereby fail to reconcile with the richness of empirical observations of

giving among households. We study patterns of inter-household transfers in four villages in southern Ghana in which we combined repeated bimonthly surveys over the course of a year, with randomized positive income shocks, some of them publicly observable, others not. We use this design because standard models of informal insurance or social taxation that assume purely self-interested behavior imply a strong, but as-yet-untested public observability hypothesis. 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, the public observability hypothesis. 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.

This strong empirical finding motivates us to refine an otherwise-standard model of dynamic risk sharing under imperfect commitment so as to allow for impure altruism, wherein the marginal gains from giving to others diminish with the transfers one makes. Giving is costly, and stochastic income has both publicly observable and unobservable components, with the public components inducing added demand for transfers from network members, following the insights of the social taxation literature. Contrary to the canonical informal insurance model, in which bigger networks and observable income are preferable, our model predicts that unobservable income shocks may facilitate altruistic giving that better targets the least well off within one's network, and that too large a network can cause even an altruistic agent to cease giving. Risk pooling is maintained, but is likely to prove incomplete.

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.

These results highlight the multiple roles played by social solidarity networks within village economies. As economists have over time increasingly tended to model networks as motivated by self-interested dynamic behavior, the altruistic function of networks has been increasingly overlooked. But that function has implications for the limits to social solidarity networks as channels for managing income shocks as well as for the tradeoffs inherent to transparency in transfer programs. Although observability of income is essential in informal insurance arrangements among purely self-interested agents, observability may impede more altruistic agents' ability to focus their giving on the most needy as they are compelled to respond to demands for assistance from the less needy within their network. Furthermore, in such networks, bigger may not always be better, in that too large a social network can induce shutdown in giving. And when people cease giving from publicly observable windfalls, they become less likely to receive transfers in the future, reinforcing induced social isolation.

The results also 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, 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 have on each other's incomes, and what does that mean for the enforceability of a risk-sharing contract?

Given these theoretical insights and empirical corroboration of the model's predictions, we feel it important to 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 the social context and local norms so as to better understand the mechanisms through which a reliance on incentives might lead to socially harmful outcomes. Our 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

- ADVANI, A. (2017): “Insurance Networks and Poverty Traps,” *Working Paper*.
- 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.
- ANGELUCCI, M., G. DE GIORGI, AND I. RASUL (2017): “Consumption and Investment in Resource Pooling Family Networks,” *Economic Journal*, 0(0).
- ASHRAF, N., AND O. BANDIERA (2017): “Altruistic Capital,” *American Economic Review*, 107(5), 70–75.
- BOURLÉS, R., Y. BRAMOULLÉ, AND E. PEREZRICHE (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.
- 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.

- 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.
- 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.
- 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.
- 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.
- SIMONS, A. M. (2016): “What is the Optimal Locus of Control for Social Assistance Programs? Evidence from the Productive Safety Net Programme in Ethiopia,” *ESSP working papers*, (86).
- SQUIRES, M. (2017): “Kinship Taxation as an Impediment to Growth: Experimental Evidence from Kenyan Microenterprises,” *Working Paper*.
- 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-ORGLÉ, 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: ADDING A SEQUENCE OF HISTORY-DEPENDENT NASH EQUILIBRIA
(SHDNE) TRANSFERS TO OUR MODEL

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

$$(20) \quad \tau^D(h_t) = \begin{cases} r \text{ s.t. } u'_1(y_1(s_t) - r)/u'_2(y_2(s_t) + r) = \gamma_1(g_1(h_t)) \\ \quad \text{if } u'_1(y_1(s_t))/u'_2(y_2(s_t)) < \gamma_1(g_1(h_t)) \\ r \text{ s.t. } u'_1(y_1(s_t) - r)/u'_2(y_2(s_t) + r) = 1/\gamma_2(g_1(h_t)) \\ \quad \text{if } u'_1(y_1(s_t))/u'_2(y_2(s_t)) > 1/\gamma_2(g_1(h_t)) \\ 0 \text{ otherwise.} \end{cases}$$

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:

$$(21) \quad \begin{aligned} U_1(h_t) = & u_1(y_1(s_t) - \tau(h_t)) - u_1(y_1(s_t) - \tau^D(h_t)) \\ & + \gamma_1(g_1(h_t))u_2(y_2(s_t) + \tau(h_t)) - \gamma_1(g_1(h_t))u_2(y_2(s_t) + \tau^D(h_t)) \\ & + \mathbb{E} \sum_{k=t+1}^{\infty} \delta^{k-t} \left\{ \begin{aligned} & u_1(y_1(s_k) - \tau(h_k)) - u_1(y_1(s_k) - \tau^D(h_t)) \\ & + \gamma_1(g_1(h_t))u_2(y_2(s_k) - \tau(h_k)) - \gamma_1(g_1(h_t))u_2(y_2(s_k) - \tau^D(h_t)) \end{aligned} \right\} \\ & - \alpha_1(g_1^D(h_t)) \end{aligned}$$

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.1. 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 $\bar{\gamma}_1^W = \bar{\gamma}_2^W = 2.5$ for both households. Transition probabilities are $\pi_{zz} = 0.3$, $\pi_{zv} = \pi_{zb} = \pi_{vz} = \pi_{bz} = 0.175$, which reflect that the most probable outcome is the case in which neither household wins a prize (zz) — all other states transpire with equal probability. When a household receives a publicly revealed prize, it will receive gift requests from all network members, i.e., $p_1(zb) = p_2(bz) = 1$. Otherwise, the probability that any given gift-network household requests a gift is $p_1(zz) = p_2(zz) = p_1(bz) = p_1(vz) = p_1(zv) = p_2(zb) = p_2(zv) = p_2(vz) = 0.2$. Finally, the discount rate is set to $\delta = 0.65$ for both households.

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

Figure A.2 shows the resulting discounted lifetime expected utility of such a contract when the initial state is either zv or bz and when household 1 extracts all the possible surplus — in other words, in the initial state, we select $\lambda(h_1) = \underline{\lambda}(s_1)$ since household 1 extracts the highest surplus when household 2's surplus is set to zero. Here, we see that discounted utility




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

FIGURE A.1: Contract Intervals

in state zb is less than discounted utility in state zv throughout the domain — this is due to the lower warm glow altruism one experiences when encumbered with a higher number of gift requests. Additionally, discounted utility decreases at a faster rate in the zb state until the zz and zb contract intervals cease to overlap — at this point, there is a slight jump in discounted utility in the zb state. This reflects a reversal in the directionality of gift-giving when the state space changes from zb to zv : all transfers are directed towards household 1 in zb (λ_{zb}) while all transfers are directed towards household 2 in zv ($\bar{\lambda}_{zv}$).³¹ However, after this

³¹i.e., household 2 would not consent to giving household 1 the entire surplus in zb if it will not receive the entire surplus in zv . Nevertheless, the one-period boost in surplus experienced by household 1 will cause its expected discounted utility to jump at this point.



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 are universally smaller in state zb and decrease at faster rates than state zv throughout. Utility spikes for a single period ($10 < g_1 < 11$), which coincides with the zb contract interval no longer overlapping with zv (see Figure A.1). This reflects a reversal in the directionality of gift-giving when the state space changes from zb to zv : all transfers are directed towards household 1 in zb ($\bar{\lambda}_{zb}$) while all transfers are directed towards household 2 in zv ($\bar{\lambda}_{zv}$). The cost of maintaining each network tie, arbitrarily set to $\alpha(g_1) = .1 + .001g_1^{1.2}$ is increasing in network size and intersects with \bar{U}_1^{zb} at a threshold of $g_1 = 15$. Beyond this point, household 1 shuts down all gift transactions when it reaches the zb state. We plot \bar{U}_1^{zb} without the possibility of shutdown; however, utility is $\bar{U}_1^{zb} = 0$ whenever $g_1 > 15$.

FIGURE A.2: Discounted Lifetime Expected Utility

jump, utility in the zb state continues to decrease monotonically. 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 shut down all giving to other households when state zb is realized.

APPENDIX B: APPENDIX TABLES

TABLE B.1
INDIVIDUAL SUMMARY STATISTICS

				Percentile	
	N	Mean	Sd	5	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/loans 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 the spouse, 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 B.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	190	119	0.31	0.22	0.56	0.06*	-0.04	0.91
N Mutual Gifts	190	119	-0.52	0.66	-0.78	0.57	0.24	0.86
Gifts and Loans (last 2 months):								
N Gifts Given	190	119	0.14	0.64	0.07	0.85	0.34	0.32
N Gifts Received	190	119	0.11	0.46	0.18	0.28	0.11	0.52
Food Consumption (last month):								
PC Food Consumption	187	117	-0.75	0.79	-4.40	0.18	2.49	0.44
PC Purchased Food	187	117	-0.02	0.99	-1.97	0.49	2.04	0.47
PC Home-produced Food	187	117	-0.73	0.49	-2.43	0.05**	0.45	0.71

Note: Balance test of round one observations. N Winners separates the sample according to those households that won any type of lottery over rounds two through five and those who did not win a lottery. “Win-at-all” presents the difference (“diff”) in the average round one responses of these two categories of households. We test whether observable characteristics are different across groups — P-values represent outcomes of t-tests. 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 B.3
PRIZE WINNINGS INFLUENCE GIFT-GIVING - COUNT DATA

Dependent Variable:	Gift-giving
	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 \times 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 \times Village FE.


TABLE B.4
QUANTILE REGRESSION ESTIMATES

Dependent Variable:	Log PC Food Consumption (1)
<hr/>	
q01	
Adjusted Network Private ($\overline{\text{Private}}'_{it}$)	3.998*** (1.047)
Adjusted Network Public ($\overline{\text{Public}}'_{it}$)	-0.669 (2.094)
Value of Private Cash Prize	0.111* (0.063)
Value of Public Cash Prize	0.102** (0.045)
<hr/>	
q12	
Adjusted Network Private ($\overline{\text{Private}}'_{it}$)	2.178** (0.874)
Adjusted Network Public ($\overline{\text{Public}}'_{it}$)	-0.707 (0.829)
Value of Private Cash Prize	-0.009 (0.043)
Value of Public Cash Prize	0.053* (0.031)
<hr/>	
q25	
Adjusted Network Private ($\overline{\text{Private}}'_{it}$)	1.725** (0.672)
Adjusted Network Public ($\overline{\text{Public}}'_{it}$)	0.085 (0.746)
Value of Private Cash Prize	-0.034 (0.033)
Value of Public Cash Prize	0.032 (0.026)
<hr/>	
q50	
Adjusted Network Private ($\overline{\text{Private}}'_{it}$)	0.671 (0.944)
Adjusted Network Public ($\overline{\text{Public}}'_{it}$)	-0.081 (0.953)
Value of Private Cash Prize	-0.026 (0.043)
Value of Public Cash Prize	0.034 (0.026)
<hr/>	
q75	
Adjusted Network Private ($\overline{\text{Private}}'_{it}$)	0.804 (0.750)
Adjusted Network Public ($\overline{\text{Public}}'_{it}$)	-0.367 (0.740)
Value of Private Cash Prize	-0.019 (0.025)
Value of Public Cash Prize	-0.008 (0.022)
Round \times 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 C.6).

APPENDIX C: APPENDIX FIGURES

FIGURE C.3: Amount of Total Transfers by Network Size




Note: This figure represents the total value of transfers τ^s from household 1. We multiply τ^s by the number of households 1 is obliged to give to in state s . In this case household 1 takes the entire share of the surplus (U_1^s is set to zero) in all its dealings and also wins a cash prize. The total transfer amount is larger when household 1 wins the publicly revealed prize (zb) relative to when it wins the privately revealed prize (zv) prior to the shutdown point ($g_1 = 15$).




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 β_{bg^2} , β_{vg^2} , β_{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 combination of public coefficients. Bars represent 95% confidence intervals. Plots of public coefficients offset by one for ease of viewing.

FIGURE C.4: 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 combination of public coefficients. Bars represent 95% confidence intervals. Plots of public coefficients offset by one for ease of viewing.

FIGURE C.5: Shut-down Hypothesis on Total Value of Gifts Given



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 C.6: Home Produced Food Over The Course of the Year



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 $H_0 : \beta_{vn} = \beta_{bn}$.

FIGURE C.7: Effect of Unadjusted Network Winnings on Food Consumption by Quantile