Charity and Reciprocity in Mobile Phone-Based Giving: Evidence from Rwanda*

Joshua Blumenstock† — University of California, Berkeley
Nathan Eagle — Santa Fe Institute
Marcel Fafchamps — Oxford University

December 22, 2011

Job Market Paper

Abstract

We provide empirical evidence that an early form of "mobile money" is used to transmit funds to individuals affected by catastrophic shocks. Contrasting two stylized models of prosocial behavior, we further provide insight into why people help each other in times of dire need. Our findings are based on the analysis of billions of mobile phone-based transactions that occur before and after a destructive earthquake in Rwanda. The observed pattern of transfers is not consistent with a model of pure charity or altruism, but better fits a model of instrumental reciprocity. This conclusion is supported by three distinct results. First, earthquake-induced transfers are increasing in the wealth of the recipient, and are not significantly related to the wealth of the sender. Second, transfers sent in response to the earthquake are highly dependent on the prior history of transfers. Third, transfers decrease as the distance between sender and recipient increases, even after controlling for the strength of pairwise relationships. Taken together, the evidence indicates that Rwandans use the mobile phone network to help afflicted friends and family, but that these gifts are motivated, at least in part, by a desire for reciprocity.

JEL Classification: O16, O17, O33

Keywords: Development; Earthquakes; Charity; Information and communications technologies; Mobile Money; Mobile Phones; Reciprocity; Risk Sharing; Rwanda.

*The authors are grateful for thoughtful comments from Alain de Janvry, Stefano DellaVigna, Frederico Finan, Mauricio Larrain, Ethan Ligon, Jeremy Magruder, Edward Miguel, Alex Rothenberg, Elisabeth Sadoulet, and seminar participants at NEUDC, PACDEV, MWIEDC, and the Berkeley Development Lunch. We gratefully acknowledge financial support from the International Growth Center, the National Science Foundation, the Institute for Money, Technology and Financial Inclusion, and the NET Institute. All errors are our own.
†Correspondence: Joshua Blumenstock, jblumenstock@berkeley.edu. Address: 102 South Hall, Berkeley, CA 94720, USA.

1 Introduction

Why do people help each other in times of dire need? Economists typically ascribe two broad motives for prosocial behavior: charity, where a giver gives out of the desire to improve the welfare of his friend; and reciprocity, where gifts are embedded in long-term relationships of reciprocal exchange. While these motives may be reasonable under ordinary circumstances, during times of real crisis it seems likely that charity would prevail. As Adam Smith observed, humans are frequently moved by pity and compassion, by "the emotion which we feel for the misery of others, when we either see it, or are made to conceive it in a very lively manner" (Smith 1759, p.3). When someone really needs help, aren't we all capable of setting aside self-interest and acting purely out of charity or altruism?

Understanding the motives for pro-social behavior is of primary concern in developing economies, where informal gifts and in-kind transfers play a critical role in enabling individuals and households to smooth consumption in the presence of temporary economic shocks (Udry 1994, Townsend 1994, Jalan & Ravallion 1999). Such knowledge not only provides insight into an important feature of many traditional societies, but can inform policies designed to promote sharing and ameliorate risk (cf. Cox 1987, Goetz & Gupta 1996).

In a developing country context, much of the theoretical literature assumes that interpersonal or inter-household transfers are realized in a repeated game of mutual insurance with limited commitment (cf. Coate & Ravallion 1993, Kocherlakota 1996, Ligon et al. 2002). The resultant risk sharing networks, while effective at insuring against uncorrelated, idiosyncratic shocks, are less effective against large, covariate shocks that affect entire communities simultaneously (Townsend 1995).[1] While altruism has long been known to play an important role in facilitating risk sharing (Becker 1991, Foster & Rosenzweig 2001, Fafchamps & Lund 2003, Platteau et al. 2006), it is typically quite difficult to empirically differentiate altruism from other factors that may affect risk sharing arrangements.

In this paper, we examine the motives governing pro-social behavior in response to large, publicly-observable economic shocks. Empirically, we exploit a comprehensive dataset of mobile phone-based activity that we obtained from the primary telecommunications operator in the Rwanda. We observe over 50 billion mobile phone transactions over a four year period, including roughly 10 million person-to-person transfers of mobile airtime, a precursor to the "mobile money" networks that are now quite common in many developing countries.[2] Our results are identified by a major earthquake that occurred in early 2008 and devastated the

[1] In a world of complete information, enforceable contracts, and negligible transaction costs, these informal arrangements could be sustained over long distances and individuals could in principle receive support from outside the community. However, real-world risk sharing is plagued by information asymmetries (Attanasio & Pavoni 2011), problems of limited commitment (Thomas & Worrall 1990), and transaction costs (Jack & Suri 2011). Thus, the overwhelming body of empirical evidence indicates that in-kind and monetary transfers typically occur between friends and family within small, local communities (Udry 1994, Fafchamps & Gubert 2007, Kurosaki & Fafchamps 2002, de Weerdt & Fafchamps 2010).

[2] As of late 2010, "branchless banking" systems had been deployed in over 80 countries worldwide (McKay & Pickens 2010). A common feature of many of these systems is a transfer mechanism that allows subscribers to transfer money or airtime

Western Lake Kivu region of the country.

We begin by providing empirical evidence that, in the immediate aftermath of the earthquake, people from all over Rwandan transferred funds to individuals living close to the epicenter. While the effect is small in absolute terms (we estimate that between \$25,000 and \$33,000 would be sent in response to a current-day earthquake), it is statistically highly significant. Moreover, we suspect that the marginal utility benefit of the transfer is due less to infra-marginal savings on airtime expenditures than it is to enabling a stricken individual to communicate with loved ones or relief workers. Indeed, most Rwandans carry a near-zero balance on their mobile phones, so the infusion of even a small amount of airtime could have a meaningful impact at a time of distress.

In our second set of results, we investigate the motives that cause people to transfer funds to those affected by the earthquake. Here, we seek to differentiate between two stylized models of mobile phone-based giving. In the first "charity-based" model, we assume that the utility function of an individual i is linearly dependent on the utility of his partner j (Becker 1974, Andreoni & Miller 2002). This model – as well as related models of Fehr & Schmidt (1999), Charness & Rabin (2002), and Andreoni (1990) – predicts that transfers increase in the wealth of the sender and decrease in the wealth of the recipient. In the second "reciprocity-based" model, we adopt the framework of risk sharing under dynamic limited commitment first proposed by Ligon et al. (2002). Building on the insights of Foster & Rosenzweig (2001) and Ligon (1998), this model predicts that shock-induced transfers are dependent on the past history of transfers between i and j, and on the costs of monitoring and enforcing contracts. Alternative formulations of reciprocity, and in particular the models of intrinsic, preference-based reciprocity (cf. Rabin 1993, Falk & Fischbacher 2006), yield similar predictions.

We then compare these empirical predictions to the actual patterns of transfers observed in the data. Contrary to our expectation, we find evidence that pure charity is not the sole determinant of earthquake-induced transfers. Instead, the data are more consistent with reciprocal motives. Three pieces of evidence support this interpretation. First, it is the wealthier individuals who receive the largest volume of transfers in the immediate aftermath of the earthquake, not the poorer individuals as the charity model predicts. Second, there is a strong history-dependence of transfers sent in response to large shocks. An individual i is significantly more likely to receive from j in the immediate aftermath of the earthquake if i sent funds to j prior to the earthquake, or if the "net balance" of the past transfers that i received from j is negative, i.e., i has given more to j than he/she received. Third and finally, post-quake transfers decrease with the geographic distance between i and j, even when controlling for the strength of the relationship between i

balance to another subscriber's account instantaneously. In Kenya, over US\$200 million is transferred over the system *per day* (Pulver 2009). In Rwanda, the system is less sophisticated, and initially only permitted interpersonal transfers of airtime. In late 2010, the system was expanded to allow for over-the-counter purchases and other monetary transactions.

and j.[3] While such an effect is expected when geographic distance limits commitment, it is not expected if transfers are purely charitable, given the absence of transaction fees and the publicly observed nature of the earthquake. In all specifications, we account for possible confounding factors by including dyad fixed effects, time dummies, and time-varying controls.

The evidence presented in this paper thus indicates that Rwandans are using the mobile phone network to help each other cope with large covariate shocks, and that these transfers appear to be driven, at least in part, by reciprocal motives. This finding is consistent with recent work by Jack & Suri (2011), who utilize household surveys to show that Kenyans with access to mobile money are better able to smooth consumption than those without.

This research contributes to, and helps synthesize, two well-established literatures on pro-social behavior and risk sharing. The distinction we draw between charitable and reciprocal motives is rather coarse in comparison to recent behavioral experiments that can distinguish between different types of charity, such as pure vs. impure altruism, and different types of reciprocity, such as intrinsic vs. instrumental reciprocity (c.f. Leider et al. 2009, DellaVigna et al. 2011, Ligon & Schechter 2011, Cabral et al. 2011). While our use of observational data limits our ability to perform this decomposition, our approach has the advantage of providing insight into prosocial behavior in response to events of dire consequence. As discussed more extensively by Levitt & List (2007), it can quite be difficult to recreate in an experimental setting the feelings and behaviors elicited by real-world catastrophes. Our *in situ* analysis is more common in the literature on risk sharing, but such empirical analysis is typically constrained by a lack of reliable data on interpersonal transfers. As a result, most studies have either relied on small samples, self-reported behaviors of subjects, or both (e.g. Fafchamps & Gubert 2007, Jack & Suri 2011). In our study, we observe a complete census of millions transfers – some of which are sent in response to exogenous shocks – and can examine the motives behind *in situ* risk sharing.

These findings also contribute to a growing body of research concerned with understanding the economic impact of mobile phones and other information and communication technologies (ICTs) in developing economies. Recent work in this area describes how mobile phones can reduce information asymmetries (Jensen 2007), lower search costs (Aker 2008), lower transaction costs (Jack & Suri 2011), and enhance communication with government agents (Shapiro & Weidmann 2011). Our analysis indicates that, unlike traditional risk sharing networks, where the vast majority of transfers occur within small, local communities, a majority of the mobile-phone based transfers come from outside the region affected by the quake.[4] By

[3] We control for social proximity in two ways. First, we count the total number of phone calls between i and j over a period of time prior to the earthquake. Second, we measure the "network flow" (Karlan et al. 2009) by counting the total number of unique paths between i and j in the call graph.

[4] In traditional risk sharing networks, Udry (1994) observes that 75 percent of surveyed Nigerian households made informal loans, but that almost all loans occurred within a village. Fafchamps & Gubert (2007) similarly observe that geographic

relaxing geographic constraints to risk sharing, mobile money enables a wider set of potential remitters. While this is generally a positive development, we note that there is considerable heterogeneity in who benefits from access to the network: wealthy individuals, and individuals with larger and more geographically disperse social networks, are most likely to receive a transfer after a severe shock.

Finally, we make two methodological contributions that can facilitate the use of similar large-scale, network-based datasets in social science research. First, we develop a novel method for estimating the permanent income of a mobile phone subscriber based solely on the history of phone calls and the structure of the call graph. Second, we develop a locational inference algorithm that allows us to impute the location of an individual based on the routing of his calls through the physical network of mobile phone towers.

The remainder of the paper is organized as follows. Section 2 describes our empirical strategy for measuring the extent to which the mobile phone network in Rwanda is used to transfer funds to people affected by severe economic shocks. In Section 3, we present two stylized models of charity and reciprocity, show that these models produce divergent empirical predictions, and outline the strategy we will use to test these models with the data. The data and the algorithms used to process them are described in Section 4. We present our empirical results in Section 5, along with several robustness checks. Section 6 concludes.

2 The effect of shocks on mobile phone-based transfers

In October of 2006, the monopoly mobile phone provider in Rwanda launched a rudimentary "mobile money" system that allowed mobile subscribers to transfer airtime from one person to another, free of charge. The first objective of this paper is to test whether this network was used to transfer airtime to individuals affected by idiosyncratic economic shocks.[5] To identify our results, we exploit the exogenous variation in transfers driven by unpredictable economic shocks, and in particular a destructive earthquake. Our empirical model estimates the extent to which individuals living in areas unaffected by such shocks send airtime to individuals close to the epicenter.

We measure this response at three levels: at the regional level (district and cell tower); at the level of individual subscribers; and at the level of dyads, where each dyad is formed by a directed pair of two subscribers. From a policy point of view, the regional analysis is perhaps the most relevant: it allows us to quantify the total value of transfers received by the affected region, and thus provide a sense of the aggregate

proximity is a major determinant of sharing patterns: when two households live near each other, it is more likely that the one will help the other. Kurosaki & Fafchamps (2002) and de Weerdt & Fafchamps (2010) obtain similar findings for Pakistan and Tanzania, respectively. See Rosenzweig & Stark (1989) for somewhat contradictory evidence from India.

[5] More recently, the Rwandan telecommunications provider significantly expanded the capabilities of this mobile money system to allow for bill payment, point-of-sale transactions, re-conversion of credit to cash, and (soon) interest-bearing savings accounts. During the period of time we analyze, however, credit could only be used to make phone calls, though it was quite common to exchange airtime for cash at retail locations. In Section 5, we discuss the extent to which these restrictions affect our interpretation of the quantitative results.

welfare benefit that was achieved.

It does, however, matter whether airtime transfers were broadly distributed across the population, or only reached a happy few. For this reason, we disaggregate transfers to the level of individual subscribers in order to analyze heterogeneity of effects. This allows us to ascertain which types of individuals are most likely to receive shock-induced transfers. Finally, because we are interested not just in the types of individuals that receive transfers, but also the types of *relationships* that support interpersonal transfers, we further disaggregate transfers at the level of pairs of users – or dyads. As we discuss in the following section, it is the dyadic- and individual-level analysis that allows us to differentiate between motives for prosocial behavior. Combining these two types of analysis is seldom possible because researchers typically only have either aggregate or survey data. We have a census of all transfers and can thus look at all levels simultaneously.

Formally, let τ_{ijrt} denote the gross transfer of airtime received by an individual i, located in region r at time t, from another individual j. Further define $\tau_{irt} = \sum_j \tau_{ijrt}$ the total gross transfers received by user i in region r at time t, and define $\tau_{rt} = \sum_i \tau_{irt}$ as the total gross transfers received by users in location r at time t. We estimate models of the form:

$$\tau_{rt} = \alpha_1 + \gamma_1 Shock_{rt} + \theta_t + \pi_r + \varepsilon_{rt} \tag{1}$$

$$\tau_{irt} = \alpha_2 + \gamma_2 Shock_{irt} + \phi NearEpicenter_{it} + \theta_t + \pi_i + \varepsilon_{irt} \tag{2}$$

$$\tau_{ijrt} = \alpha_3 + \gamma_3 Shock_{irt} + \phi NearEpicenter_{it} + \theta_t + \pi_{ij} + \varepsilon_{ijrt} \tag{3}$$

where $Shock_{rt}$ is a dummy variable equal to 1 if location r received a shock on day t, and $Shock_{irt}$ equals 1 if i was near the epicenter at the time of the shock.[6] θ_t is a vector of time dummies, and π_r, π_i, and π_{ij} are fixed effects for the region, individual, and dyad, respectively. $NearEpicenter_{it}$, which indicates whether i was in the area affected by a shock on day t (irrespective of a shock occurring), controls for the possibility that individuals might receive transfers when visiting the area affected by the earthquake. In regression (2), individuals i who never receive airtime transfers are excluded since they do not help identify γ_2. In regression (3), pairs of individuals (i,j) that are never observed to transfer money to one another are similarly omitted. Time dummies θ_t control for long-term growth in traffic, as well as day-of-the-week (e.g., week-end) and day-of-the-month (e.g., payday) effects that affect all regions similarly. Location and recipient fixed effects π_r and π_i control for the fact that different locations or users are more likely to receive transfers on average. Dyadic fixed effects π_{ij} control for the average intensity and direction of transfer flows between two users.

[6] Technically, S_{irt} is the interaction between $NearEpicenter_{it}$ and $DayOfShock_t$, a dummy variable that takes the value 1 on the day of the shock (for all i) and zero otherwise. The uninteracted variable $DayOfShock_t$ is observed by the time dummies θ_t and omitted for clarity.

Finally, to minimize the likelihood that our results are driven by differential growth in mobile usage across locations, we restrict the analysis to a specific time window $T_{min} \leq t_s \leq T_{max}$ around the time of the shock t_s.

Identification is achieved as in a difference-in-difference framework: parameters γ_1, γ_2 and γ_3 represent the average treatment effect of the shock on people with access to the mobile money network. The exogeneity of $Shock_{rt}$ is guaranteed since its timing could not have been predicted, i.e., the shock constitutes a natural experiment. If $\gamma_1 > 0, \gamma_2 > 0$ and $\gamma_3 > 0$, this is interpreted as evidence that the shock $Shock_{rt}$ caused an increase in airtime transfers to users in the affected region. We check the robustness of our results in various ways, notably by varying the time window over which the models are estimated, by controlling for several factors that depend on both time and location, and by running a number of falsification and placebo tests. Following Bertrand et al. (2004), in individual and dyadic regressions standard errors are clustered by location (i.e., by the location of the nearest cellular tower).

3 Charity and reciprocity in interpersonal transfers

A primary objective of the paper is to provide insight into the motives behind transfers that are made in response to a publicly observed shock. Since our data are observational, and since the identities of the mobile subscribers are unknown, we are somewhat constrained in our ability to impute what are fundamentally personal decisions. Thus, we stop short of recent experimental work that, through clever manipulation of experimental conditions, can differentiate between instrumental and intrinsic reciprocity (Ligon & Schechter 2011, Cabral et al. 2011, Leider et al. 2009, Charness & Rabin 2002). Instead, we follow Leider et al. (2009) and divide the motivations for prosocial behavior into two rough categories which, for short, we call 'charity' and 'reciprocity.'[7] In this section, we develop stylized models of charity and reciprocity, derive comparative statics, and describe the identification strategy we will employ to differentiate between the two models.

3.1 Theoretical Framework

(i) Charity

By charity we refer to the broad class of motives where a giver gives because he receives direct utility from the act of giving or from increasing the utility of another. The canonical example of this behavior is pure

[7] Our distinction also parallels the distinction Ligon & Schechter (2011) draw between "preference-related" motives and "incentive-related" motives.

altruism (Becker 1976, Andreoni & Miller 2002), where one person's utility depends positively on another's:

$$U_{it} = u_i(x_{it} - \tau_{jit}) + \gamma_{ij} u_{jt}(x_j + \tau_{jit}) \tag{4}$$

As before, we denote by τ_{jit} a transfer sent to i from j at time t. Assuming $u_i(\cdot)$ and $u_j(\cdot)$ are increasing and concave, with x_i representing the income of individual i and γ_{ij} denoting the level of altruism felt by i toward j, it is easily shown that two predictions of such a model are

$$\frac{\partial E[\tau_{ji}|x_i]}{\partial x_i} \geq 0 \quad \text{and} \tag{5}$$

$$\frac{\partial E[\tau_{ji}|x_j]}{\partial x_j} \leq 0 \tag{6}$$

Giving is expected to increase in the income of the sender (as the marginal cost of giving decreases), and decrease in the income of the recipient (as the marginal benefit of a gift decreases). Such predictions are supported by observed patterns of altruism in a variety of contexts, including charitable giving in the United States (Andreoni 2006) and the behavior of "rescuers" in Nazi-occupied Europe during World War II (Hoffman 2010).

While the comparative statics (5) and (6) are most transparent in the model of pure linear altruism specified by equation (4), similar predictions obtain from several related models of charitable behavior, and we make no pretense to be able to distinguish between them. Thus, models of inequity and inequality aversion (where i seeks to minimize $|x_i - x_j|$), social welfare models (where i maximizes $\min\{x_1, ..., x_n\}$), and warm glow giving (where $v(x_j + \tau_{ji})$ in (4) is replaced with τ_{ji}) all yield similar predictions (cf. Fehr & Schmidt 1999, Bolton & Ockenfels 2000, Charness & Rabin 2002, Andreoni 1990, List & Lucking-Reiley 2002).

Defining charity thus broadly, it follows that if the primary motive for transfers is a charitable one, we expect transfers on average to come from richer users and to flow to poorer users. In cases of directed altruism, where γ_{ij} varies across dyads, we expect transfers to decrease with the social distance between i and j, but conditional on social distance, there is no a priori reason to expect a relationship between transfers and geographic distance. Similarly, after controlling for social distance, charitable transfers should be "memory-less" (Fafchamps & Lund 2003), and past transfers should not directly influence transfers sent in response to shocks. To the extent that an association does exist between past transfers and current transfers, we would expect it to be positive, as past transfers may reveal information about how much i cares about j, above and beyond the undirected measures of relationship strength that we employ.

(ii) Reciprocity

By reciprocity, we refer to motives that are embedded in long-term relationships of bilateral exchange. In the discussion that follows, we focus on a particular type of instrumental reciprocity, where mutual exchange is motivated by the expectation of future reciprocation (cf. Coate & Ravallion 1993, Karlan et al. 2009). This model most transparently leads to empirical predictions that can be tested with the data at our disposal. Other models of reciprocity, and most notably the intrinsic, preference-based reciprocity modeled by Rabin (1993) and Falk & Fischbacher (2006), produce similar predictions. Since our intent is not to differentiate between these different types of reciprocity, we present a simple model of dynamic limited commitment that captures many of the central tenets of the wider literature.[8]

We adopt a model of risk sharing under dynamic limited commitment, developed by Ligon et al. (2002) and Foster & Rosenzweig (2001). Following Foster & Rosenzweig (2001), we assume i has stationary, single-period utility specified by (4), but allow for the possibility that i expects to benefit from future interaction with j:

$$U_{it} = u_i(x_{it} - \tau_{jit}) + \gamma_{ij} u_j(x_{jt} + \tau_{jit}) + E \sum_{s=t+1}^{\infty} \delta^{s-t} [u_i(x_{is} - \tau_{jis}) + \gamma_{ij} u_j(x_{js} + \tau_{jis})] \tag{7}$$

The first part of (7) is identical to the altruistic model (4), while the second term captures the discounted expected utility of the relationship.

This formulation produces two key insights relevant to the current analysis. First, when contracts are not fully enforceable ex-post, transfers received in the current period will depend on past transfers given. This property is formally derived in Appendix A, however the intuition is quite simple: In the stationary model of (4), i and j will transfer the necessary τ_{jit} to equate the ratio of their ex post marginal utilities to γ (or $1/\gamma$ if $x_i < x_j$). If γ is sufficiently small, i and j operate in autarky. By contrast, in the dynamic model specified by (7), i and j also derive utility from expected future interactions and transfers, and so may adjust the ratio of marginal utilities at time t to maintain the relationship and avoid reversion to a series of static Nash equilibria. Intuitively, when i suffers a shock in period t, the marginal utility of τ_{ijt} will be quite high and so he will be willing to sacrifice a greater share of his continuation utility in exchange for a larger transfer at t. As a result, we expect transfers from i to j sent in response to a shock to be decreasing in the net balance of transfers previously made from i to j (denoted by T_{jit}^{net}):

$$\frac{\partial E[\tau_{jit} | T_{jit}^{net}]}{\partial T_{jit}^{net}} \leq 0 \quad \text{where} \quad T_{jit}^{net} = \sum_{s=0}^{t-1} \tau_{jis} - \tau_{ijs} \tag{8}$$

[8] For recent experimental work that differentiates between different types of reciprocity, see Leider et al. (2009), Ligon & Schechter (2011), and Cabral et al. (2011). Fehr & Schmidt (2006) and Sobel (2005) provide theoretical overviews.

As robustness checks, we further expect that τ_{jit} will decrease in the gross volume of prior transfers from i to j (denoted by T_{jit}) and increase in gross prior transfers from j to i (denoted by T_{ijt}):

$$\frac{\partial E[\tau_{jit}|T_{jit}]}{\partial T_{jit}} \leq 0 \quad \text{where} \quad T_{jit} = \sum_{s=0}^{t-1} \tau_{jis} \tag{9}$$

$$\frac{\partial E[\tau_{jit}|T_{ijt}]}{\partial T_{ijt}} \geq 0 \tag{10}$$

Empirically, this dynamic model of limited commitment explains the data better than a stationary model with no history-dependence (Ligon et al. 2002, Genicot & Ray 2003, Fafchamps & Lund 2003).

The second insight of our model of risk sharing under limited commitment, also discussed extensively in Ligon (1998) and De Vreyer et al. (2010), is that transfers would be expected to decrease as, ceteris paribus, the cost of monitoring and enforcement increases. In the context of our study, it is natural to assume that monitoring costs increase monotonically with geographic distance D_{ij}. For instance, if j wants to verify the damage to i (e.g., injury, destroyed building), j has to travel to the affected area, and the cost of travel increases with distance. Thus, if the need to monitor constrains transfers, we expect that:

$$\frac{\partial E[\tau_{ij}|d_{ij}, S_{ij}]}{\partial d_{ij}} \leq 0 \tag{11}$$

i.e., the further away j resides from i, the more costly it is to verify the effect of the shock on i, and the harder it is to overcome j's fear of being cheated.

Of course, distance is likely correlated with other factors that influence the decision to give, but which are not related to monitoring and enforcement per se. For instance, transaction costs, which often rise with distance, have received recent attention by Jack & Suri (2011). However, since the cost of remitting over Rwandan mobile is free at all distances, transaction costs are unlikely to drive empirical estimates of (11). More generally, when we operationalize (11) in a regression setting, we will always condition our empirical results on S_{ij}, an undirected measure of the strength of the relationship between i and j, and a dyad fixed effect π_{ij}, so that we estimate the partial effect of geographic distance holding constant other unobservable characteristics of the dyad.

Unlike the model of charity, the model of reciprocity does not make strong predictions regarding the relative wealth of i and j. We might expect transfers to go to wealthier individuals if strategic agents seek to ingratiate themselves for future reciprocation. Such an interpretation is supported by recent work by Schechter & Yuskavage (2011), who find that transfers are likely to flow from more to less wealthy households in unreciprocated relationships, while reciprocated relationships are more likely between wealthier households. Alternatively, we could observe flows from the rich to the poor if the poor reciprocate in ways other than airtime (Fafchamps 1999, Platteau 1995).

3.2 Identification and Estimation

Table 4 summarizes the empirical predictions of the two stylized models of mobile phone-based giving. If, as seems reasonable to expect under such exigent circumstances, transfers given in response to severe shocks are motivated by feelings of charity, they should be increasing in the wealth of the sender and decreasing in the wealth of the recipient. However, conditional on the strength of the relationship between sender and recipient there should be no marginal impact of the geographic distance between partners, or the past history of transfers. On the other hand, if transfers are embedded in relationships of reciprocity, then the relative wealth of the two individuals would have an ambiguous effect on transfers. The key determinants instead are the past history of transfers between individuals and the geographic proximity. Both models predict that transfers will be increasing in social proximity.

To test these predictions empirically, we estimate heterogeneous effect models of the form:

$$\begin{aligned}\tau_{it} = \quad & \alpha_2 + \quad \gamma_2 Shock_{it} + \beta_2 Z_i Shock_{it} + \phi_2 NearEpicenter_{it} + \\ & \qquad \eta_2 Z_i DayOfShock_t + \zeta_2 Z_i NearEpicenter_{it} + \theta_t + \pi_i + \varepsilon_{it} \qquad (12)\\ \tau_{ijt} = \quad & \alpha_3 + \quad \gamma_3 Shock_{it} + \beta_3 Z_{ij} Shock_{it} + \phi_3 NearEpicenter_{it} + \\ & \qquad \eta_3 Z_{ij} DayOfShock_t + \zeta_3 Z_{ij} NearEpicenter_{it} + \theta_t + \pi_{ij} + \varepsilon_{ijt} \qquad (13)\end{aligned}$$

where Z_i and Z_{ij} are vectors of characteristics associated with individuals and pairs of individuals (dyads). The regressions are used to estimate the partial effects described in Table 4. As before, $Shock_{it}$ takes the value one if i is affected by the shock and zero otherwise. $Shock_{it}$ is the product of $DayOfShock_t$, a dummy variable taking the value one on the day of a severe shock and $NearEpicenter_{it}$, a dummy variable indicating whether i was close to the shock on day t. Double interaction terms of the form $Z_i DayOfShock_t$ are included to control for the possibility that, in the country as a whole, variation in Z_i affects transfers on the day of the shock differently from other days.

We are primarily interested in using models (12)-(13) to differentiate between charity and reciprocity by testing the effects of $Z_{ij} = \{x_i, x_j, S_{ij}, D_{ij}, T_{ij}^{net}\}$ on transfers sent in response to severe shocks. In the following section, we describe the data used in the analysis, our technique for measuring the various components of (1)-(3) that are necessary to estimate the average treatment effect, and our method for measuring wealth, past transfer history, geographic distance, and social connectedness.

4 Data

The main dataset used in this paper comes from Rwanda's primary telecommunications operator, which until recently held a near monopoly on mobile telephony in the country.[9] The data contain a comprehensive log of all activity that occurred on this network between early 2005 and late 2008. In total, we observe detailed information on over 50 billion transactions (including calls, text messages, and airtime transfers and purchases), covering 1.5 million users over four years. Summary statistics of this dataset are given in Table 2.

During the four-year period for which we have data, uptake of mobile phones was extremely rapid. In early 2005, only 2.5 percent of Rwandans owned a mobile phone, but by 2010 mobile penetration had increased to 33.4 percent (a compound annual growth rate of roughly 74 percent).[10] Such rapid growth is common in many sub-Saharan Afrian nations, where landlines are rare and the cost of owning a mobile phone is falling quickly.[11]

4.1 Interpersonal Transfers

Our empirical analysis focuses on interpersonal transfers of airtime credit between mobile subscribers. These transfers are made possible by a rudimentary "mobile banking" system that was launched in October 2006. After purchasing airtime from a local vendor, mobile subscribers are able to send that airtime to another subscriber, instantaneously and free of transaction fees. During the period of time we analyze, the system only allowed for transfers of airtime, and there was no official mechanism that allowed for conversion of airtime back into cash. In 2010 the phone company greatly expanded the capabilities of the system to allow for over-the-counter transactions, bill payments, and other debit account-like services. Following the trend in much of the developing world, where over 1.7 billion people own a mobile phone but do not have a bank account (CGAP and GSMA 2009), most Rwandans did not have access to formal financial services in early 2008. As can be seen in Table 3, compared to other available mechanisms for sending funds, the mobile money system is considerably cheaper, faster, and easier to access.

In our dataset, we observe detailed information on roughly 10 million interpersonal transfers of airtime.

[9]During the window of time we examine, the operator we focus on maintains over 90% market share of the mobile market. The company's primary competitor did not gain traction in the market until the end of 2008, and only very recently has the market become competitive. The number of landlines in Rwanda is insignificant (roughly 0.25% penetration).

[10]Source: International Telecommunications Union. `http://www.itu.int/ITU-D/ICTEYE/Reports.aspx`, accessed November 2011.

[11]Though prices are falling, the cost of mobile telephony are still relatively high and represent a significant portion of household expenditures (Ureta 2005). In Rwanda it costs roughly \$50 for the phone, and an additional \$0.20 per minute and \$0.10 per SMS (see Republic of Rwanda (2010) and Donner (2008)). The ITU estimates the monthly "price basket" for mobile service to be \$12.30 per month, which is based on the prepaid price for 25 calls per month spread over the same mobile network, other mobile networks, and mobile to fixed calls and during peak, off-peak, and weekend times. The basket also includes 30 text messages per month (`http://devdata.worldbank.org/ict/rwa_ict.pdf`).

For each transaction, we know the date, time, and value of the transfer, as well as a unique (anonymized) identifier for both the sender and recipient. To estimate the regressions described in sections 2 and 3, we aggregate this raw data on each day for each region (equation 1, individual (equation 2), and dyadic pair of individuals (equation 3). This allows us to measure, at three different levels of aggregation, the net and gross volume of airtime received on each day. Summary statistics of the dataset are provided in Table 2.

4.2 Locational Inference

The identification strategy we employ relies on spatial, as well as temporal, variation in transfers. Therefore, it is important that we be able to assign each individual, on each day, to an approximate geographic location. The mobile phone operator records information on the location of each individual at the moment a call is initiated or terminated, and the only record kept is of the nearest cellular tower, not the actual GPS coordinates of the subscriber. As can be seen in Figure 1, which shows the spatial distribution of cell phone towers in early 2008, towers in rural areas are relatively sparse, so being able to locate individuals between towers can potentially improve the precision of our models.[12]

To solve this problem, we develop a locational inference algorithm that allows us to approximate the continuous trajectory of each user through time and space, based on the intermittent sequence of phone calls logged by the mobile operator. The method utilizes a kernel function $K(\cdot)$ to estimate an unknown location at time t from the kernel-weighted Euclidean centroid of known locations at times in the vicinity of t (cf. Blumenstock 2012).

Formally, we estimate the unknown location $\widehat{r_{it}}$ of individual i at time t as

$$\widehat{r_{it}} = \frac{1}{N_{it}} \sum_{s=T_{min}}^{T_{max}} K(\frac{t-s}{h}) \cdot \widehat{q_{is}}$$

where N_{it} is the total number of phone calls made by i within a window of time $[T_{min}, T_{max}]$ symmetric around t, and $\widehat{q_{is}}$ is the (known) location of the tower used at time s. $K(x)$ is a symmetric function that integrates to one, which specifies the extent to which additional weight is placed on calls close in time to t. In our results we use a uniform kernel such that $K(u) = 1/N_i$, however very little changes if a different kernel is specified.

[12]the median area covered by a single cell phone tower is 72km^2. Note that cellular coverage is not affected by topology to the same extent as radio transmitters as in Yanagizawa-Drott (2010).

4.3 Characteristics of individuals and dyads

To differentiate between the models of charity and reciprocity described in Section 3, we measure characteristics of individuals and dyads $Z_{ij} = \{x_i, x_j, S_{ij}, D_{ijt}, T_{ijt}^{net}\}$ as follows.

Social connectedness (S_{ij}): Both models predict that transfers will increase in the strength of the relationship between i and j, so measurement of this effect does not contribute directly the theory we hope to test. However, since certain components of Z_{ij} may be correlated with S_{ij} (notably D_{ijt} and T_{ij}), controlling for S_{ij} directly can bias. We measure S_{ij} in two ways: first, following Marmaros & Sacerdote (2006) we simply calculate the total number of non-monetary interactions (phone calls and text messages) between i and j in the N days prior to t. Second, building on the insight of Karlan et al. (2009), we additionally measure the *maximum network flow* as the number of distinct paths between i and j in the complete undirected call graph.

Past transfer history (T_{ijt}): Since we observe all transactions between i and j over a 4-year period, it is possible at any point t to compute the gross balance of payment sent from j to i as $T_{ijt} = \sum_{s<t} \tau_{ijs}$ and from i to j as $T_{jit} = \sum_{s<t} \tau_{jis}$. The net balance of transfers T_{ijt}^{net} is simply $T_{ijt} - T_{jit}$.

Geographical distance (D_{ijt}): To estimate the distance between i and j at time t, we first compute the (latitude,longitude) locations $\widehat{r_{it}} = (\phi_i, \lambda_i)$ and $\widehat{r_{jt}} = (\phi_j, \lambda_j)$ using the locational inference algorithm described in Section 4.2. Then, we compute the arc distance D_{ijt} using the haversine formula:

$$D_{ijt} = 2r \cdot \arcsin\left(\sqrt{\sin^2\left(\frac{\Delta\phi}{2}\right) + \cos\phi_i \cos\phi_j \sin^2\left(\frac{\Delta\lambda}{2}\right)}\right)$$

where $\Delta\phi = |\phi_i - \phi_j|$, $\Delta\lambda = |\lambda_i - \lambda_j|$ and $r = 6356.78$ is the radius of the earth.

Wealth (x_i and x_j): While the preceding metrics are straightforward to compute, it is considerably more difficult to measure the wealth of each individual since all users in the dataset appear anonymously and without demographic information such as the age, gender, education level, or direct proxies for socioeconomic status. Instead, we construct a crude proxy of each individual's wealth by combining detailed call data with regionally-tagged data from a household Demographic and Health Survey and follow-up phone interviews conducted in Rwanda in 2009 and 2010. We provide a concise description of this method here; the reader is referred to Appendix C for further details.

1. Using a 10,000-household Demographic and Health Survey (DHS) that contains detailed consumption and expenditure information (Government of Rwanda, 2008), we first estimate a hedonic regression of

annual expenditures Y_{id} of household i in district d on fixed assets A_{id} and housing characteristics H_{id}.

$$Y_{id} = \alpha + \sum_{j}^{h_{max}} \beta_j H_{id} + \sum_{k}^{a_{max}} \delta_k A_{id} + \mu_d + \epsilon_{id} \tag{14}$$

The quantities Y_{id}, A_{id}, and H_{id} are all captured in the DHS, and μ_d is a district fixed effect. Equation (14) predicts the annual expenditure $\widehat{Y}_{id}$ of a household on the basis of easily observable A_{id}, and H_{id}. We regard $\widehat{Y}_{id}$ is as a proxy for permanent income.

2. Next, we use data from a phone survey to relate H_{id} and A_{id} to phone usage. The survey, conducted by the authors, covers a geographically stratified random sample of approximately 2,000 mobile subscribers. For each individual, we collected basic demographic information, together with data on H_{id} and A_{id}. Armed with H_{id} and A_{id}, it is possible to compute predicted annual expenditures $\widehat{Y}_{id}$ for the 2,000 mobile subscribers using the coefficient estimates from equation (14).

3. Finally, we compute, for each phone user, a vector of phone usage variables X_{ir} thought to be correlated with income, such as the total number of calls made and the average amount of airtime purchased over a given time interval (we exclude information on airtime transfers to minimize potential endogeneity). A subset of these variables are summarized in Appendix Table (10). We then fit a flexible model of the form:

$$\widehat{Y}_{id} = f(X_{ir}) \tag{15}$$

and estimate $f(.)$ using data from the phone user survey. Our wealth index is the predicted $\widehat{\widehat{Y}}_{id}$ obtained by applying the estimated flexible function $\widehat{f}(.)$ to the full sample of 1.5 million phone users. This variable is used as proxy for wealth or permanent income when estimating heterogeneous effect equations (12) to (13). Details on $f(\cdot)$, as well as potential endogeneity concerns, are discussed at length in Appendix C.

5 Results

Our identification strategy requires a shock $Shock_{rt}$ that is exogenous to transfers on the mobile phone network. The primary shock that we exploit is a large earthquake that occurred in the Western Rusizi and Nyamasheke districts of Rwanda on February 3, 2008. The magnitude 6 earthquake left 43 dead and 1,090 injured. It destroyed 2,288 houses and caused regional school closures and electrical outages (though only one cell tower of 267 was affected). The effects of the earthquake, though large, were geographically circumscribed. The United States Geographical Survey estimates an impacted radius of approximately 20 kilometers from the epicenter – see Figure 1. Based on news reports and discussions with individuals in Rwanda, it does not appear that any particular demographic subgroup of the population was disproportionately affected by this earthquake, and in particular, rich and poor households appear to have been similarly affected (USGS 2009).[13] This event is ideal for our estimation strategy since the shock is unequivocally exogenous and precisely located in time and space. We later demonstrate that our results are robust to using alternative shock measures including a severe flood that occurred in late 2007.

We begin by estimating models (1)-(3) to measure the causal impact of the earthquake on interpersonal transfers. We then turn to models (12) and (13) to measure heterogeneity and test the models of charity and reciprocity presented in Section 3.1.

5.1 Average effect of the earthquake

(i) Baseline results

We first estimate equation (1) at the regional level, and present the results in Panel A of Table 5. The dependent variable τ_{rt} is the total value of transfers received on day t by region r. Column 1 defines r at the level of the political district (of which there are 30); column 2 defines r at the level of the cell tower (of which there are 267). While the district-level specification may seem more natural, the advantage of the tower-level results is that each observation corresponds to a smaller geographical unit and thereby allows us to more precisely identify the regions affected by the quake. In all specifications, we use data from 30 days before to 30 days after the earthquake, though as demonstrated in (iii) results our results change little if we use a different time window. Region and day fixed effects are included as additional regressors to control for systematic differences across districts and over time. The earthquake shock variable $Shock_{rt}$ equals one on February 3rd 2008, the day of the earthquake, in regions affected by the earthquake.[14] Robust standard errors are reported, clustered at the district level.

We are primarily interested in the coefficient on "Earthquake Shock," which indicates the extent to which an anomalous volume of mobile airtime was sent to regions affected by the earthquake. In all specifications the estimate is highly significant, with T-statistics between 7 and 16.

Columns (3) and (4) repeat the estimation at the level of the individual user and of the dyadic pair of individuals. In column (3), the dependent variable τ_{irt} is the amount of airtime transfered to individual i in location r at time t. Users who never receive airtime transfers are excluded since they do not help identify

[13] Much of the damage was sustained in and around the semi-urban city of Cyangugu, where a relatively representative subset of the population resides. See, for instance, `http://earthquake.usgs.gov/eqcenter/eqinthenews/2008/us2008mzam/`

[14] For the district-level specification, affected regions are the districts of Rusizi and Nyamasheke. For the tower-level specification, it is towers within 20km of the epicenter, though similar point estimates and standard errors are produced if we redefine affected areas as those lying anywhere between 10 to 50 miles of the epicenter.

the effect of the shock, leaving roughly 110,000 unique individuals. Results from the dyad-level regression (3) are presented in column (4) of Table 5. In this regression, pairs in which i never receives airtime from j are ignored from the estimation, leaving roughly 180,000 valid dyads. At both levels, the $Shock_{irt}$ coefficient is positive and statistically significant. The evidence is thus consistent: at all levels of aggregation we observe an increase in gross transfers.

(ii) Interpretation

Do these effects matter? Based on the coefficient estimate in column (2) of Table 5, we observe that the earthquake produced an additional influx of approximately 42,000 RWF (or $84 U.S. Dollars) to the 15 towers within 20km of the epicenter. This amount is small in absolute terms, but Rwanda is a small country of ten million people, and only 50,000 people lived in the region affected by the earthquake, of whom no more than 7 percent owned a mobile phone at the time of the earthquake.[15] Based on the data at our disposal, we know that only 1,400 individuals living in the earthquake region had used the service prior to January 2008. If the net influx were evenly distributed among these eligible users (which it wasn't - most transfers were sent to a lucky few), each individual would have received roughly 30 RWF (about five cents).

If the sole benefit of this transfer were an infra-marginal savings on future airtime expenditures, then the utility gain would indeed be quite small. However, as can be seen in Figure 3, most Rwandans carry very little airtime on their account. At midnight the night before the earthquake, the median balance for the entire population of phone users was only 49.4 RWF, and roughly 32 percent of all subscribers had an airtime balance of less than 5 RWF. Thus, the marginal utility of 30 RWF to someone who had just experienced an earthquake is potentially quite large – it would have been sufficient to make a short phone call or send a text message, to call for help or to simply reassure a loved one.[16] In the immediate aftermath of several recent natural disasters, mobile phones have been instrumental in facilitating rescue attempts, though we have no direct evidence of this occurring after the Lake Kivu earthquake.[17]

To further put these numbers in context, it is worth noting that since the earthquake, service utilization has increased over 400-fold. As of early 2011, there were between 750,000 and 1,000,000 active mobile money users in Rwanda each day. This compares to 2,500 at the time of the earthquake. If we were willing to

[15] It is difficult to estimate the total population affected by the earthquake, however Cyangugu, the largest city in the region, has a population of roughly 20,000. Mobile phone penetration in Rwanda was roughly 7 percent in 2008, but a disproportionate share of owners live in the urban capital of Kigali, and penetration rates outside the capital are significantly lower (Blumenstock et al. 2010).

[16] 30 RWF would also have been sufficient to enable a Rwandan to send a "missed call" (also referred to as a flash or a beep) to a friend, which is a common way of signaling that the caller wishes to talk but does not want to pay for the cost of a call. Sending missed calls in Rwanda in 2008 required that the subscriber have a positive balance on his or her account.

[17] For a recent example, see "In Turkey, Desperate Race to Find Trapped Survivors", *New York Times*, October 25, 2011: "Some dug with their bare hands, while other used heavy machinery to remove chunks of fallen concrete and relied on cellphone calls from the missing in the search for survivors... a 19-year-old in the town survived by using his cellphone to direct teams to the collapsed building where he had been trapped."

assume that airtime transfers following an earthquake would increase proportionally to the number of active users, a similar earthquake today would be expected cause an additional influx of US$25,000 to $33,000 to affected areas.[18] In neighboring Kenya, where the population is much larger and the local mobile money service is more widely used, the daily volume of money transferred over the mobile phone network is in excess of US$200 million, compared to US$1,500 in Rwanda at the time of the quake. If we were willing to assume that emergency transfers would increase proportionally to the total volume of airtime transfers, we would expect an influx of approximately US$11.2 million ($84 ×200M/1500) to affected regions.

(iii) Robustness

As a robustness check, we redo the same analysis using net instead of gross transfers. The concern is that gross transfers may misrepresent the aggregate magnitude of the transfers if individuals who receive airtime pass it on to others in the same region. This could result in double-counting at the district or cell tower level. For the individual user regression (2) we redefine the dependent variable as $\tau'_{irt} = \sum_j \tau_{ijrt} - \sum_j \tau_{jirt}$, that is, the transfers received by i from others minus the transfers given by i to others. At the district and cell tower levels, we proceed as follows. Let $\tau_{r_1 r_2 t} = \sum_{i \in r_1} \sum_{j \in r_2} \tau_{ijrt}$ where r_1 and r_2 are two different locations (e.g., districts or cell tower area); $\tau_{r_1 r_2 t}$ represents the total transfers received by individuals in location r_1 from individuals in location r_2. Summing over all other locations yields the gross transfers from other locations to location r_1. Net inflows to region r_1 are thus $\tau'_{r_1 t} = \sum_{r_2} \tau_{r_1 r_2 t} - \sum_{r_2} \tau_{r_2 r_1 t}$. Results, shown in Panel B of Table 5 are similar in significance and magnitude to those reported in Panel A, implying that the magnitude of our findings is not driven by double counting.

In Appendix B, we present several additional robustness checks to ensure that our results are (i) not sensitive to the econometric specification, (ii) do not depend on the choice of time window, and (iii) are not affected by the structure imposed on the variance matrix. In addition, we run a series of "placebo" tests to demonstrate that similar results do not obtain on days where no earthquake occurs. Finally, to demonstrate that the effects observed in response to the Lake Kivu earthquake are likely to generalize to other severe shocks, we show that a similar, albeit muted, response is observed after a series of large floods that occurred in late 2007. The reader is referred to Appendix B for further details on these tests.

[18] $\frac{750,000}{2,500} * \$84 = \$25,200$. If these transfers are proportional to traffic and traffic increases non-linearly in the number of subscribers, as much of the network literature suggests, the projected amount may even be much larger. It is also conceivable, however, that early adopters are not representative of late adopters and respond more strongly to an earthquake; in this case, transfers need not increase proportionally with traffic or the number of users.

5.2 Differentiating between charity and reciprocity

Our next task is to test whether the pattern of transfers observed following the Lake Kivu earthquake is more consistent with a model of charity or of reciprocity. As a starting point, we note that the vast majority of pairwise relationships we observe do not exhibit a strong reciprocal component. Of the 646,713 dyadic pairs (i, j) for which a transfer is observed in one direction (from either i to j or j to i), transfers are observed in both directions in only 143,394 dyads (22 percent). However, for the subset of dyads that respond to the earthquake, that ratio jumps to 31 percent. This difference, which is statistically significant, suggests that the quake-induced transfers are embedded in long-term reciprocal relationships.

To test these intuitions more formally, we introduce heterogeneous effects into our analysis. As summarized in Table 4, we have proposed three indirect tests: (i) *Wealth*: if transfers are motivated by charity, they are expected to flow from wealthier to poorer individuals; not necessarily so if they follow a reciprocal motive. (ii) *History-dependence*: Ceteris paribus, charitable transfers should not depend on history; however, if transfers are embedded in reciprocal relationships, shock-induced transfers should exhibit history-dependence as characterized by equations (8)-(10). (iii) *Geographic distance:* Transfers motivated by reciprocity are expected to decrease with distance, but conditional on the strength of the relationship between i and j, charitable transfers should exhibit no such dependence.[19]

We focus on the individual and dyadic regression models (12) and (13), where the vector of covariates Z_{ij} includes the measures of x_i, x_j, S_{ij}, D_{ij}, and T_{ij}^{net} discussed in Section 4.[20] Regression results are presented in Tables 6-7. As before, we include data from 30 days before to 30 days after the earthquake (January 4, 2008 through March 3, 2009), and compute Z_{ij} using data from 2007. In all specifications we include a vector of daily fixed effects and interactions between Z_{ij} and $DayOfQuake_t$ and $NearEpicenter_{it}$, though these coefficients are omitted from most tables for clarity of presentation. In our preferred specification (column 1), we additionally include individual and pairwise fixed effects π_{ij} to reduce potential bias from time-invariant omitted variables. However, inclusion of these fixed effects makes it impossible to estimate the unconditional effect of Z_{ij} on τ_{ij}, i.e., whether certain characteristics make transfers more likely on non-shock days. Thus, we follow Wooldridge (2005) and separately recover the average partial effects by obtaining the predicted $\widehat{\tau_{ij}}$ from (13) and then regressing these predicted values on Z_{ij} (column 2). We also include specifications with no fixed effects as a point of reference, though as noted these estimates are likely to be biased (column 1). In all specifications, we include a full set of interactions between S_{ij} and $Shock_{it}$, in order to reduce the potential bias that other factors correlated with Z_{ij} (such as how much i and j like one another) are

[19] Of course, individuals who are themselves affected by the earthquake are less likely to be in a position to assist,

[20] While in principle we could perform similar analysis at the regional (district and tower) level by aggregating $Z_{ir} = \sum_{i \in r} Z_i$, the results do not improve our ability to distinguish between the two models so we omit them for clarity.

spuriously driving the effect we attribute to Z_{ij}. The measure S_{ij} used in our main specifications is simply the total number of phone calls observed between i and j in the year prior to t; Appendix B shows that our results are also robust to a different measure of S_{ij} – the number of shared contacts between i and j – proposed by Karlan et al. (2009).

As summarized in the final column of Table 4, we find that transfers increase in the wealth of the recipient, are not significantly related to the wealth of the sender, are dependent upon the past history of transfers, and decrease nonlinearly with distance. We discuss each of these findings in turn, then summarize the extent to which this evidence informs our understanding of the motives governing transfers sent in response to severe shocks.

(i) Wealth

To measure the marginal effect of the wealth of the sender and recipient on transfers, we use as a wealth proxy the predicted expenditure variable $\widehat{\widehat{Y}}_{id}$ described in Section 4. To avoid the possibility that results are driven by differences between high- and low-usage individuals (i.e. that richer users may receive more airtime but also transfer more to others), we use net transfers as the dependent variable, though we find similar results with respect to gross transfers. Results are presented in Table 6.

The primary coefficient of interest is the interaction between the wealth of the recipient x_i and the $Shock_{irt}$ dummy. The estimates in the second row of Table 6 indicate that wealthier individuals are significantly more likely to receive transfers in the immediate aftermath of the earthquake. This effect exists conditional on the wealth of the sender x_j, and on the normal level of transfers between i and j, as captured in the dyad-specific fixed effects. The former control is important because it limits the possibility that wealthier individuals are receiving more simply because they have wealthier friends, and not because of their own wealth. The latter rules out the possibility that the effect is caused by time-invariant aspects of the i-j relationship, for instance that wealthy i's may always receive more from j, even in transactions that are unrelated to economic shocks.[21]

The fact that earthquake-induced transfers increase in the wealth of the recipient (x_i) but are not significantly correlated with the wealth of the sender (x_j) is difficult to reconcile with a model of pure charity/altruism, which predicts that transfers would decrease in x_i and increase in x_j. If an individual j knows two people equally affected by the earthquake, a model of charity predicts he would give more to the one who has a higher marginal utility of the transfer. In the Rwandan context, it is natural to assume that the marginal utility of a transfer is higher for poorer individuals, who are significantly more likely to carry a

[21] In the final three rows of the table, it is evident that, on non-earthquake days, wealthy individuals are more likely to both send and receive airtime, and that dyads with strong social ties are more likely send money. These results are consistent with expectations, but do not help in differentiating between charity and reciprocity.

near-zero balance on their account.[22]

While the specification with dyadic fixed effects most stringently controls for potential omitted sources of bias, a more parsimonious approach is given in Appendix B.4, which includes *sender* fixed effects to directly test the intuition stated above, that conditional on the identity of the sender, the wealthier recipients receive more on the day of the earthquake. As expected, these results are consistent with the point estimates in Table 6.

(ii) History Dependence

We next investigate whether pairs of individuals with a history of reciprocal transfers are more likely to exchange transfers in response to the earthquake. The theory presented in Section 3.1 suggests that if shock-induced transfers are motivated by pure charity, they should not depend heavily on prior transfers. Or, to the extent that past transfers from i to j signal i's compassion for j (above and beyond the undirected strength of the relationship S_{ij}), transfers sent in response to the earthquake should increase in past transfers from i to j. The model of reciprocity predicts the opposite: transfers from i to j after the earthquake should be driven more by prior transfers from j to i than from i to j.

Following Foster & Rosenzweig (2001), we capture this history-dependence with T_{jit}^{net}, the net balance of payments made from i to j in the periods prior to t. In Appendix B, we separately test the partial effect of the gross volume of past transfers from i to j (T_{ijt}) and from j to i (T_{jit}). Results are presented in Table 7.

We observe that an individual i who has a net positive balance with j (i.e., j "owes" i money), is significantly more likely to receive help from j on the day of the earthquake (row two). Importantly, this effect persists even when controlling for the social proximity of i and j, and is unlikely to be caused by unmodelled correlation between past transfer activity and general characteristics of the dyadic relationship (such as shared ethnicity, family ties, etc.). This finding is consistent with a model of reciprocity, but does not seem natural if transfers are motivated by pure charity.

The positive and significant coefficient on $T_{jit}^{net} * Shock_{irt}$ is particularly striking given that the uninteracted effect of the prior net balance is negative (row four). In other words, on normal days without shocks, transfers are expected to flow primarily in one direction – i.e., if i has transferred more to j than j to i prior to t, it is more likely that another i to j transfer will occur at t. Under normal circumstances, there is a structural dependency where one person consistently gives and the other receives; after the shock, what is becomes important is the past reciprocity of the relationship.

[22] See Figure 3. The correlation between end-of-day balance and wealth is 0.15, with a T-statistic of 58.8.

(iii) Distance

Finally, we test whether transfers come uniformly from other unaffected regions of Rwanda, or whether the distance between individuals affects of the volume of transfers received. Given that the earthquake was a publicly-observed shock about which the whole country was quickly informed, and since there were no fees associated with sending airtime, if transfers follow a charitable motive we would expect all unaffected areas to contribute (conditional on the strength of the social ties between dyads). In reciprocal arrangements, where quid-pro-quo contracts are harder to enforce over long distances, we expect a negative relationship. Appendix Table 1 shows that the data exhibits a significant negative association between distance and transfers.

While this simple linear relationship is consistent with our model of reciprocity, there is strong reason to suspect that the relationship between transfers and geographic distance is non-linear. For instance, when i and j live nearby, they are likely to be similarly affected by large, covariates shocks, and thus able to help each other than they would in response to smaller, idiosyncratic shocks. We therefore additionally estimate the nonparametric relationship between transfers and distance. Figure 4 shows how $\partial\tau_{ij}/\partial D_{ij}$ evolves with D_{ij}.[23] We observe that after the quake, people with many contacts near the epicenter do not receive more transfers, presumably because nearby friends are also affected by the earthquake. People with contacts more than 30 Km away from the epicenter are more likely to receive transfers in the aftermath of the earthquake, but the effect dies down for contacts located more than 100 Km from the epicenter. To provide further intuition, figure 5 shows the distribution of distances over which transfers are sent in the month prior to the Lake Kivu earthquake, for transactions involving at least one user in the earthquake region. While the vast majority of transfers are sent over a short distance, there are a large number of transfers sent to and from the capital of Kigali, which is approximately 150km from the epicenter. After the quake, the distribution shifts toward transfers occurring in an intermediate range of 20-130 kilometers, where the j is likely to be unaffected by the quake, but still lives relatively close to i.

5.3 Limitations and Alternative Explanations

In the results and discussion above, we have used observational data on interpersonal airtime transactions to measure the effect of large shocks on transfers, and to attempt to infer the motives governing prosocial behavior observed in response to natural disasters. We do not want to overstate our ability to identify these motives, since our analysis is fundamentally constrained by the data at our disposal: we observe very little information about the demographic characteristics of users, their consumption, and their vulnerability to risk. Rather, our position is that the sum total of empirical evidence clearly indicates that pure charity is

[23] Specifically, we plot the coefficients that results from interacting the $Shock_{irt}$ variable with disaggregated measures of the size of i's network.

not the only force at play. The data are more consistent with the interpretation that transfers are embedded in long-term relationships of reciprocal exchange. Before concluding, we briefly address two limitations of this analysis.

(i) Airtime vs. Cash

Our analysis focuses on the transfer of mobile airtime, which is different from hard cash. This affects our results in two ways. First, it affects the interpretation of the point estimates presented above. To the extent that the marginal utility of airtime is lower than the marginal utility of cash, it may seem disingenuous to refer to an average treatment effect in dollar terms. While this criticism is legitimate, as we have noted in Section 5.1(ii), there are strong reasons to suspect that the marginal utility of airtime, particular in the immediate aftermath of a severe shock, may be quite high.

The fact that we observe only transfers of airtime also affects the external validity of our results, in the sense that it may seem reasonable to expect a different response when people are able to send real mobile money, as opposed to mobile airtime. At the time of the earthquake in 2008, informal mechanisms existed that allowed individuals to convert airtime to cash (typically for a 20% commission at a local retailer), but qualitative evidence suggests such conversions were rare. In the time that has passed since the earthquake of 2008, the Rwandan operator has upgraded their system to allow for over-the-counter purchases, re-conversion of airtime to cash by authorized agents, and (soon) interest-bearing savings accounts. These services, and evidence from neighboring Kenya, suggest the mobile-based financial services should be expected to play an increasingly prominent economic role in the lives of Rwandans.[24] For these reasons, we believe, if anything, that our estimates represent a lower bound on the quantity of mobile money that would be sent in response to a current-day catastrophe. As mobile money becomes more common and more useful, we expect the mobile network will play an increasingly important role in facilitating risk sharing over distance.

(ii) Intensive vs. Extensive Margin

Since we only observe activity that occurs on the mobile phone networks, we are unable to infer whether the mobile phone-based transfers are substitutes for transfers that would have otherwise been sent using another mechanism, or whether they affect the extensive margin. This latter effect could go both directions: it could increase the total transfers sent, if the advantages of the technology (speed, efficiency, lowered transaction costs, and lowered minimum transaction) induce more people to give. Alternatively, it is possible that mobile phone-based transfers, which tend to be quite small (usually on the order of one dollar), could crowd out other gifts that would otherwise have been sent in a larger denomination.

[24] In Kenya, where the M-PESA mobile money system has been wildly successful, there are over 25,000 mobile money agents.

While these effects are the active focus of future work, it is perhaps useful to provide some qualitative context. Table 3 summarizes the alternative methods for transferring money over long distances, that were available in Rwanda circa 2008. MoneyGram, Western Union, and the Post Office are the other official methods for transferring money, but transaction costs across these services range from 10 - 100% of the value of the money sent. For each of these services, it is impossible to transfer amounts under US$10. In the informal sector, the most common method for transferring money is by bus/taxi, but for that service the driver typically charges 10-20% of the amount transferred, and the availability of the service is contingent on the schedule of busses and the condition of the roads. With the mobile transfer service, by contrast, the transfer of money is instantaneous and has no associated fees or commissions.

6 Conclusion

Using detailed log data from Rwanda, we have tested whether individuals in locations affected by a natural disaster receive transfers from unaffected parts of the country. We find a significant increase in airtime sent to individuals affected by the 2008 earthquake. The impact is robust to a variety of econometric specifications, and does not exist for a large number of "placebo" earthquakes on different dates and in different locations. Based on simple back-of-the-envelope calculations, we estimate that the total response to a similar current-day earthquake in Rwanda would be between $25,000 and $33,000.

We interpret the anomalous transfers observed after the quake as prima facie evidence that people use the mobile network to help each other cope with economic shocks. However, the motives behind these transfers are not clear ex ante. In particular, it is ambiguous whether people give out of purely charitable motives, or whether they are giving out of an expectation of future reciprocity (or as a repayment for past assistance). By contrasting two stylized models of mobile phone-based giving, we show that these two motives for giving produce conflicting empirical hypotheses, in particular with respect to the marginal effect of wealth, distance, and past transfers on the amount transferred following the earthquake. Testing these hypotheses with the data from Rwanda, we find that the giving observed after the earthquake is most consistent with a model of reciprocity.

Given the increasing prominence of mobile phones in the developing world, it is important that we develop a better understanding of the economic impacts that this technology will have on the lives of their users. In this paper, we argue that by allowing for inexpensive interpersonal transfers, mobile phones are providing a new method for risk sharing. Since the alternative mechanisms used for interpersonal transfers are considerably slower and more expensive, this immediate influx of support may be of material consequence. As the capabilities of the mobile money system are further expanded, for instance to allow for purchase of over-the-counter goods with airtime, the potential benefits to users on the networks can be expected to increase.

We also find that the potential benefits of the mobile-based service are not evenly distributed. In prior work, we have shown that there is a sharp divide between people who do and don't own mobile phones: relative to non-owners, phone owners are significantly wealthier, better educated, older, and more likely to be male (Blumenstock & Eagle 2012). In this paper, we have noted that even among mobile phone owners, it is the wealthiest who are most likely to receive transfers – both on normal days and in the period immediately after a large economic shock. Thus, transfers of airtime, or mobile-based transfers of money, may not reach the people who need them most. Such evidence suggests that blanket investment in telecommunications infrastructure may not have the transformative economic impacts envisioned by the popular media. Instead, policies that more actively target poorer segments of the population, and which lower barriers to adoption and use, might better ensure that the potential benefits of mobile phones are realized by those with the greatest need.

References

Aker, J. (2008), 'Does digital divide or provide? the impact of cell phones on grain markets in Niger', *BREAD Working Paper* (177).

Andreoni, J. (1990), 'Impure altruism and donations to public goods: a theory of warm-glow giving', *The Economic Journal* **100**(401), 464477.

Andreoni, J. (2006), 'Philanthropy', *Handbook on the Economics of Giving, Reciprocity and Altruism* **2**, 12011269.

Andreoni, J. & Miller, J. (2002), 'Giving according to GARP: an experimental test of the consistency of preferences for altruism', *Econometrica* **70**(2), 737–753. ArticleType: research-article / Full publication date: Mar., 2002 / Copyright 2002 The Econometric Society.

Attanasio, O. P. & Pavoni, N. (2011), 'Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption', *Econometrica* **79**(4), 1027–1068.

Becker, G. S. (1974), 'A theory of social interactions', *Journal of Political Economy* **82**(6), 1063–1093.

Becker, G. S. (1976), 'Altruism, egoism, and genetic fitness: Economics and sociobiology', *Journal of Economic Literature* **14**(3), 817826.

Becker, G. S. (1991), *A Treatise on the Family*, Harvard Univ Pr.

Bertrand, M., Duflo, E. & Mullainathan, S. (2004), 'How much should we trust Differences-in-Differences estimates?', *Quarterly Journal of Economics* **119**(1), 249275.

Blumenstock, J. (2012), 'Inferring patterns of internal migration from mobile phone call records: Evidence from Rwanda', *Information Technology for Development* **18**(2), forthcoming.

Blumenstock, J. E. & Eagle, N. (2010), 'Mobile divides: Gender, socioeconomic status, and mobile phone use in Rwanda', *4th International IEEE/ACM Conference on Information and Communications Technologies and Development* .

Blumenstock, J. E. & Eagle, N. (2012), 'Divided we call: Disparities in access and use of mobile phones in Rwanda', *Information Technology and International Development* (forthcoming).

Blumenstock, J. E., Gillick, D. & Eagle, N. (2010), 'Whos calling? Demographics of mobile phone use in Rwanda', *AAAI Symposium on Aritificial Intelligence and Development* **18**, 116–117.

Bolton, G. E. & Ockenfels, A. (2000), 'ERC: a theory of equity, reciprocity, and competition', *American Economic Review* p. 166193.

Cabral, L., Ozbay, E. & Schotter, A. (2011), Intrinsic and instrumental reciprocity: An experimental study, Technical report, Working paper). New York University.

CGAP and GSMA (2009), 'Mobile money for the unbanked: Annual report 2009'.

Charness, G. & Rabin, M. (2002), 'Understanding social preferences with simple tests*', *Quarterly journal of Economics* **117**(3), 817869.

Coate, S. & Ravallion, M. (1993), ‘Reciprocity without commitment : Characterization and performance of informal insurance arrangements’, *Journal of Development Economics* **40**(1), 1–24.

Collins, D., Morduch, J., Rutherford, S. & Ruthven, O. (2009), *Portfolios of the Poor: How the World’s Poor Live on $2 a Day*, Princeton University Press.

Cox, D. (1987), ‘Motives for private income transfers’, *Journal of Political Economy* **95**(3), 508–546. ArticleType: research-article / Full publication date: Jun., 1987 / Copyright 1987 The University of Chicago Press.

de la Statistique du Rwanda (INSR), I. N. & Macro, O. (2006), *Rwanda Demographic and Health Survey 2005.*, INSR and ORC Macro, Calverton, Maryland.

De Vreyer, P., Gubert, F. & Roubaud, F. (2010), ‘Migration, self-selection and returns to education in the WAEMU’, *Journal of African Economies* **19**(1), 52 –87.

de Weerdt, J. & Fafchamps, M. (2010), ‘Social identity and the formation of health insurance networks’, *Working Paper* .

Deaton, A. & Muellbauer, J. (1980), *Economics and consumer behavior*, Cambridge Univ Pr.

Deaton, A. & Zaidi, S. (2002), *Guidelines for constructing consumption aggregates for welfare analysis*, World Bank Publications.

DellaVigna, S., List, J. & Malmendier, U. (2011), *Testing for altruism and social pressure in charitable giving*, U.C. Berkeley mimeo.

Donner, J. (2008), ‘The rules of beeping: Exchanging messages via intentional ”Missed calls” on mobile phones’, *Journal of Computer-Mediated Communication* **13**(1), 1–22.

Fafchamps, M. (1992), ‘Solidarity networks in preindustrial societies: Rational peasants with a moral economy’, *Economic Development and Cultural Change* **41**(1), 147–174.

Fafchamps, M. (1999), ‘Risk sharing and quasi-credit’, *Journal of International Trade and Economic Development* **8**, 257278.

Fafchamps, M. & Gubert, F. (2007), ‘The formation of risk sharing networks’, *Journal of Development Economics* **83**(2), 326–350.

Fafchamps, M. & Lund, S. (2003), ‘Risk-sharing networks in rural philippines’, *Journal of Development Economics* **71**(2), 261–287.

Falk, A. & Fischbacher, U. (2006), ‘A theory of reciprocity’, *Games and Economic Behavior* **54**(2), 293–315.

Fehr, E. & Schmidt, K. M. (1999), ‘A theory of fairness, competition, and cooperation*’, *Quarterly journal of Economics* **114**(3), 817868.

Fehr, E. & Schmidt, K. M. (2006), ‘The economics of fairness, reciprocity and altruism-Experimental evidence and new theories’, *Handbook on the economics of giving, reciprocity and altruism* **1**, 615691.

Foster, A. D. & Rosenzweig, M. R. (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.

Fox, J. (1997), *Applied regression analysis, linear models, and related methods*, Sage Publications, Inc.

Genicot, G. & Ray, D. (2003), 'Group formation in RiskSharing arrangements', *Review of Economic Studies* **70**(1), 87–113.

Goetz, A. & Gupta, R. (1996), 'Who takes the credit? gender, power, and control over loan use in rural credit programs in bangladesh', *World development* **24**(1), 4564.

Hoffman, M. (2010), 'Does higher income make you more altruistic? evidence from the holocaust', *Review of Economics and Statistics* **93**(3), 876–887.

Jack, W. & Suri, T. (2011), 'Risk sharing and transaction costs: Evidence from Kenya's mobile money revolution', *Working Paper* .

Jalan, J. & Ravallion, M. (1999), 'Are the poor less well insured? evidence on vulnerability to income risk in rural china', *Journal of development economics* **58**(1), 6181.

Jensen, R. (2007), 'The digital provide: Information (Technology), market performance, and welfare in the south indian fisheries sector', *The Quarterly Journal of Economics* **122**(3), 879–924.

Kabbucho, K., Sander, C. & Mukwana, P. (2003), 'PASSING THE BUCK: money transfer systems: The practice and potential for products in Kenya', *MicroSave White Paper* .

Karlan, D., Mobius, M., Rosenblat, T. & Szeidl, A. (2009), 'Trust and social collateral', *The Quarterly Journal of Economics* **124**(3), 1307 –1361.

Kocherlakota, N. R. (1996), 'Implications of efficient risk sharing without commitment', *The Review of Economic Studies* **63**(4), 595.

Kurosaki, T. & Fafchamps, M. (2002), 'Insurance market efficiency and crop choices in pakistan', *Journal of Development Economics* **67**(2), 419–453.

Leider, S., Mbius, M. M., Rosenblat, T. & Do, Q. (2009), 'Directed altruism and enforced reciprocity in social networks', *The Quarterly Journal of Economics* **124**(4), 1815 –1851.

Levitt, S. D. & List, J. A. (2007), 'What do laboratory experiments measuring social preferences reveal about the real world?', *The Journal of Economic Perspectives* **21**(2), 153–174.

Ligon, E. (1998), 'Risk sharing and information in village economies', *The Review of Economic Studies* **65**(4), 847–864.

Ligon, E. A. & Schechter, L. (2011), 'Motives for sharing in social networks', *SSRN eLibrary* .

Ligon, E., Thomas, J. P. & Worrall, T. (2002), 'Informal insurance arrangements with limited commitment: Theory and evidence from village economies', *The Review of Economic Studies* **69**(1), 209–244.

List, J. A. & Lucking-Reiley, D. (2002), 'The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign', *Journal of Political Economy* **110**(1), 215233.

Marmaros, D. & Sacerdote, B. (2006), 'How do friendships form?', *Quarterly Journal of Economics* **121**(1), 79119.

McKay, C. & Pickens, M. (2010), 'Branchless banking 2010: Whos Served? At What Price? Whats Next?', *Washington, D.C.* .

Montgomery, M. R., Gragnolati, M., Burke, K. A. & Paredes, E. (2000), 'Measuring living standards with proxy variables', *Demography* **37**(2), 155174.

Orozco, M. (2009), 'Emerging markets for Rwanda: remittance transfers, its marketplace and financial intermediation'.

Platteau, J., Kolm, S. & Ythier, J. M. (2006), Solidarity norms and institutions in village societies: Static and dynamic considerations, *in* 'Foundations', Vol. Volume 1, Elsevier, pp. 819–886.

Platteau, J. P. (1995), 'An indian model of aristocratic patronage', *Oxford Economic Papers* p. 636662.

Pulver, C. (2009), 'The performance and impact of MPESA: preliminary evidence from a household survey', *Presentation to FSD-Kenya* .

Rabin, M. (1993), 'Incorporating fairness into game theory and economics', *The American Economic Review* p. 12811302.

Republic of Rwanda (2010), 'Development of the telecom sector for the last three years', *Rwanda Utilities Regulatory Agency (RURA)* .

Rosenzweig, M. R. & Stark, O. (1989), 'Consumption smoothing, migration, and marriage: Evidence from rural india', *Journal of Political Economy* **97**(4), 905–926. ArticleType: research-article / Full publication date: Aug., 1989 / Copyright 1989 The University of Chicago Press.

Schechter, L. & Yuskavage, A. (2011), 'Reciprocated versus unreciprocated sharing in social networks', *SSRN eLibrary* .

Shapiro, J. N. & Weidmann, N. B. (2011), 'Talking about killing: Cell phones, collective action, and insurgent violence in iraq', *Working Paper* .

Smith, A. (1759), *The theory of moral sentiments*, 1853 edn, Henry G Bohn, London.

Sobel, J. (2005), 'Interdependent preferences and reciprocity', *Journal of Economic Literature* **43**(2), 392–436. ArticleType: research-article / Full publication date: Jun., 2005 / Copyright 2005 American Economic Association.

Thomas, J. & Worrall, T. (1990), 'Income fluctuation and asymmetric information: An example of a repeated principal-agent problem', *Journal of Economic Theory* **51**(2), 367–390.

Townsend, R. M. (1994), 'Risk and insurance in village India', *Econometrica: Journal of the Econometric Society* **62**(3), 539–591.

Townsend, R. M. (1995), 'Consumption insurance: An evaluation of Risk-Bearing systems in Low-Income economies', *Journal of Economic Perspectives* **9**, 83–83.

Udry, C. (1994), 'Risk and insurance in a rural credit market: An empirical investigation in northern nigeria', *The Review of Economic Studies* **61**(3), 495–526.

Ureta, S. (2005), 'Variations on expenditure on communications in developing countries: A synthesis of the evidence from albania, mexico, nepal and south africa (2000-2003)', *Diversifying Participation in Network Development* .

Wooldridge, J. M. (2005), 'Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity', *Journal of Applied Econometrics* **20**(1), 39–54.

World Bank Group (2009), 'Remittance prices world wide', *Available at http://remittanceprices.worldbank.org* .

Yanagizawa-Drott, D. (2010), Propaganda and conflict: Theory and evidence from the rwandan genocide, Technical report, Working Paper, Harvard University.

Tables

Table 1: Mobile phone penetration: Number of mobile phones per 100 inhabitants.

	2000	2001	2003	2005	2007	2009	2010	Annual Growth
Rwanda	0.49	0.78	1.49	2.47	6.53	24.3	33.4	73.9%
South Africa	18.28	23.39	35.93	71.60	87.08	92.67	100.48	8.2%
United States	38.53	44.77	54.90	68.63	82.47	89.15	89.86	6.5%

Source: International Telecommunication Union (`http://www.itu.int/`, accessed November 2011)

Table 2: Summary statistics of mobile network data.

Dates covered	All dates	Earthquake window
	10/1/2006-7/1/2008	1/3/2008-3/3/2008
Panel A: Aggregate traffic		
Number of Calls	46,000,000,000	868,786,684
Number of interpersonal transfers	9,202,954	362,053
Number of unique users	1,084,085	119,745
Number of people who send airtime	870,099	48,295
Number of people who receive airtime	946,855	101,351
Number of people who both send and receive	732,869	29,901
Number of unique dyads	646,713	159,204
Panel B: Basic statistics (12/1/2007-4/1/2008)	Mean	S.D.
Transactions per user (send+receive)	6.05	12.05
Average distance per transaction (km)	13.51	27.67
Average transaction value (RWF)	223.58	652.02

Notes: The window 10/1/2006-7/1/2008 encompasses the entire dataset with valid data on interpersonal airtime transfers. The window 1/3/2008-3/3/2008 is the same window used in later regressions. US$1=550RWF.

Table 3: Formal services for transferring money in Rwanda, c. 2008

Service	Estimated Fees (small transfers)	Availability	Source
Mobile Phone	Free	3 million phones	
MoneyGram	7% - 100% ($15 minimum)	5 locations	MoneyGram Website, 2010
Western Union	10%-100% ($10 minimum)	Approx. 50 locations	Western Union Website, 2010
Post Office	8%-50%	19 branches	World Bank Group (2009)
Commercial Bank	6%-40%	Only in urban, semi-urban areas	World Bank Group (2009)
Public Bus	6% - 20%	Populous areas	Avg. for Bus-Star, Scandinavian Transportation

Notes: Data compiled in August 2010 from sources listed in column 4. See Orozco (2009) and Kabbucho et al. (2003) for further quantitative estimates, or Collins et al. (2009) for a more general overview.

Table 4: Summary of predictions and results

		Predicted		
Partial	Interpretation	Charity	Reciprocity	Observed
$\partial\tau_{ijt}/\partial x_i$	Wealth of i (recipient)	Negative	–	Positive
$\partial\tau_{ijt}/\partial x_j$	Wealth of j (sender)	Positive	–	–
$\partial\tau_{ijt}/\partial S_{ij}$	Social connectedness of i and j (recipient)	Positive	Positive	Positive
$\partial\tau_{ijt}/\partial D_{ijt}$	Geographic distance between i and j at time t	–	Negative	Negative
$\partial\tau_{ijt}/\partial T_{ijt}^{net}$	Net balance of transfers from j to i	(Positive)	Negative	Negative

Notes: Summary of predictions of stylized models of charity and reciprocity presented in Section 3.1, and results described in Section 5. Parentheses indicate weak predictions.

Table 5: Average Effect of the Earthquake on Transfers Received

	(1) District	(2) Cell Tower	(3) User	(4) Dyad
Panel A: Gross transfers received (total incoming)				
Earthquake Shock	14,169.2***	2,832***	9.48***	8.14***
	(1,951.3)	(177)	(0.74)	(1.34)
$NearEpicenter_{it}$			1.26***	0.23*
			(0.19)	(0.11)
Unconditional Mean	19,006.9	2,436.2	5.90	3.65
Unconditional Mean (affected area)	6355.9	1245.3	3.8	3.2
Panel B: Net transfers received (incoming - outgoing)				
Earthquake shock	12822.58***	3053***	10.01***	8.816***
	(1599.87)	(116)	(1.08)	(1.65)
$NearEpicenter_{it}$			0.58	0.45^{+}
			0.46	(0.24)
Unconditional Mean	0	0	0	0
Unconditional Mean (affected area)	1398.51	248.38	0.514	0.64
Panel C: Total cost of calls received				
Earthquake Shock	2,501,220**	565,329**	247.95***	–
	(886,529)	(168,349)	(34.21)	–
$NearEpicenter_{it}$			-22.17***	–
			(5.89)	–
Unconditional Mean	3,845,858	425,740	162.91	
Unconditional Mean (affected area)	1,843,455	353,501	154.17	
Number of observations	1,800	16,020	6,619,440‡	10,032,721
Day dummies	Yes	Yes	Yes	Yes
Fixed effects	District	Tower	Individual	Dyad (Directed)

Notes: In Panels A and B, the dependent variable is the gross (Panel A) and net (Panel B) transfers received on a given day by a given district (column 1), cell tower (column 2), individual subscriber (column 3), or by an individual i from a specific individual j (column 4). In Panel C, the dependent variable is the total amount of money spent to call the district/tower/subscriber on each day. The Earthquake Shock variable takes the value one for regions/subscribers who were within 20km of epicenter on the day of the earthquake; $NearEpicenter_{it}$ takes the value one for all observations where the subscriber is within 20km of the epicenter (even when there is no earthquake). The unconditional mean reports the average of the dependent variable across the entire 2-month window (January 4 2008 - March 3 2008), for the country as a whole and for the region affected by the earthquake. ‡ Column (3) of Panel C includes all phone subscribers, not just individuals who had used the mobile money service, and includes 35,539,241 observations. We do not estimate column (4) of Panel C because the roughly 100 billion observations made the computation infeasible. Standard errors, clustered by district, reported in parentheses. * significant at $p < .05$; $^{**}p < .01$;$^{***}p < .001$.

Table 6: Net transfers and wealth

	(1) No Fixed Effects	(2) Avg. Partial Effects	(3) Fixed Effects
Earthquake Shock	10.206*		9.794*
	(4.56)		(4.33)
Recipient Wealth x_i * Shock	5.767***		5.763***
	(1.11)		(1.26)
Sender Wealth x_j * Shock	15.682		14.903
	(14.39)		(13.33)
Social Proximity S_{ij} * Shock	0.101		0.110
	(0.08)		(0.07)
Recipient Wealth (x_i)	0.863***	0.976***	
	(0.17)	(0.09)	
Sender Wealth (x_j)	1.641***	1.618***	
	(0.11)	(0.13)	
Social Proximity (S_{ij})	0.040***	0.040***	
	(0.00)	(0.00)	
Day dummies	Yes	Yes	Yes
Fixed effects	None	None	Dyad (directed)
Number of observations	10032721	174799	10032721

Notes: Outcome is τ_{ijt}, the total airtime received by i from j on day t. Regressions include observations from the period January 4, 2009 through March 3, 2008. Wealth proxies x_i and x_j are computed using equations (14) and (15), as described in the text. S_{ij} is measured by counting the number of phone calls between i and j in the last three months of 2007. All regressions include $NearEpicenter_{it}$ and pairwise interaction terms (e.g. x_i * $NearEpicenter_{it}$, x_i*$DayOfQuake_t$); these coefficients are omitted for clarity. Standard errors, clustered by district, reported in parentheses. *** p<0.001, ** p<0.01, * p<0.05, + p<0.10.

Table 7: Net transfers and History-Dependence

	(1) No Fixed Effects	(2) Avg. Partial Effects	(3) Fixed Effects
Earthquake Shock (to i)	8.866***		8.403***
	(2.11)		(1.87)
Net Bal. Outgoing Airtime T^{net}_{jit} * Shock	0.006*		0.006*
	(0.002)		(0.002)
Social Proximity S_{ij} * Shock	0.192		0.200
	(0.14)		(0.13)
Net Balance of Outgoing Airtime (T^{net}_{jit})	-0.001***	-0.011***	-0.010***
	(0.00)	(0.00)	(0.00)
Social Proximity (S_{ij})	0.051***	0.051***	
	(0.00)	(0.00)	
Day dummies	Yes	Yes	Yes
Fixed effects	None	None	Dyad (directed)
Number of observations	10032721	174799	10032721

Notes: Outcome is τ_{ijt}, the total airtime received by i from j on day t. Regressions include observations from the period January 4, 2009 through March 3, 2008. The net balance of outgoing airtime T^{net}_{jit} is measured as the total volume of airtime sent from i to j minus the total volume of airtime received by i from j prior to t. S_{ij} is measured by counting the number of phone calls between i and j in the last three months of 2007. All regressions include $NearEpicenter_{it}$ and pairwise interaction terms (e.g. T^{net}_{jit} * $NearEpicenter_{it}$, T^{net}_{jit}*$DayOfQuake_t$); these coefficients are omitted for clarity. Standard errors, clustered by district, reported in parentheses. *** p<0.001, ** p<0.01, * p<0.05, + p<0.10.

Appendix Tables

Appendix Table 1: Net transfers and Distance

	(1)	(2)	(3)
	No Fixed Effects	Avg. Partial Effects	Fixed Effects
Shock (recipient)	4.191**		3.901**
	(1.39)		(1.31)
D_{ijt} * Shock	-0.117*		-0.116*
	(0.05)		(0.05)
S_{ij} * Shock	0.178		0.191
	(0.13)		(0.12)
Distance (D_{ijt})	-0.013**		-0.021***
	(0.00)		(0.00)
Social Proximity (S_{ij})	0.050***	0.050***	
	(0.00)	(0.00)	
Day dummies	yes	yes	yes
Fixed effects	None	None	Dyad (directed)
Number of observations	9915422	193256	9915422

Notes: Outcome is τ_{ijt}, the total airtime received by i from j on day t. Regressions include observations from the period January 4, 2009 through March 3, 2008. D_{ijt} measures the distance between i and j in kilometers on day t, using the locational inference algorithm described in Section 4.2. S_{ij} is measured by counting the number of phone calls between i and j in the last three months of 2007. All regressions include $NearEpicenter_{it}$ and pairwise interaction terms (e.g. D_{ijt} * $NearEpicenter_{it}$, D_{ijt}*$DayOfQuake_t$); these coefficients are omitted for clarity. Standard errors, clustered by district, reported in parentheses. *** p<0.001, ** p<0.01, * p<0.05, + p<0.10.

Appendix Table 2: Sensitivity of estimation to function form assumptions

	(1) Pooled OLS	(2) OLS w/Controls	(3) Region FE	(4) Region & Day FE
Earthquake Shock	1793.639***	2819.503***	2787.305***	2710.861***
	(313.08)	(121.33)	(136.18)	(183.53)
$DayOfQuake$	-748.548**	-1375.294***	-1287.145***	
	(212.87)	(111.76)	(119.77)	
$NearEpicenter_{it}$	-2262.728***	-510.751***		
	(577.57)	(73.57)		
Total call volume		0.074***	0.064***	0.103***
		(0.00)	(0.01)	(0.01)
Outgoing transfers		0.677***	0.637***	0.527***
		(0.03)	(0.03)	(0.04)
Tower Fixed Effects	No	No	Yes	Yes
Date Fixed Effects	No	No	No	Yes
R^2	0.008	0.702	0.729	0.753
N	171414	74895	74895	74895

Notes: Outcome is the total amount transferred into a tower on a single day. $NearEpicenter$ defined as those towers within 20 miles of the earthquake epicenter. Columns 2-4 include controls for overall network activity at the tower-day level. Column 3 includes tower-level fixed effects. Column 4 includes daily fixed effects. Estimates made using data from October 1, 2006 through July 1, 2008. Heteroskedasticity-robust SE's in parentheses (clustered at district level). * p<0.05, ** p<0.01, *** p<0.001.

Appendix Table 3: Lagged effects of the earthquake on transfers and calls received.

	(1) Transfers Received	(2) Calls Received	(3) Int'l Calls Received
Shock	13512.649***	14208.656***	142.584*
	(1335.51)	(3753.46)	(56.71)
shock_lag1	-917.294	4594.386***	126.538
	(1330.88)	(499.87)	(76.47)
shock_lag2	1540.204	1639.237	62.719
	(2796.36)	(1026.90)	(49.83)
shock_lag3	830.593	1297.175***	47.690
	(3157.92)	(295.21)	(33.18)
shock_lag4	-189.597	552.066*	-28.472
	(1518.35)	(208.00)	(17.58)
shock_lag5	-40.867	1070.376***	-66.248*
	(3028.17)	(229.54)	(29.46)
shock_lag6	-2648.816	927.869**	-95.259
	(3138.61)	(303.77)	(58.89)
shock_lag7	-335.684	1468.774**	-86.875
	(849.38)	(420.29)	(46.27)
shock_lead1	810.813	228.141	34.601
	(1732.01)	(316.09)	(25.81)
shock_lead2	1341.489	218.922	40.632
	(1124.93)	(387.07)	(44.32)
shock_lead3	-2460.249	-72.909	-24.811
	(2003.26)	(201.42)	(59.38)
Total call volume	0.010		
	(0.01)		
Outgoing transfers	0.876***		
	(0.02)		
Outgoing calls		0.969***	
		(0.00)	
Outgoing int'l calls			0.959***
			(0.02)
Constant	155.928	2069.706	1417.339***
R^2	0.984	1.000	0.943
N	16808	18840	18840

Notes: Outcome specified in column heading. All specifications include daily and district fixed effects. Heteroskedasticity-robust standard errors in parenthesis (clustered at district level).

Appendix Table 4: Placebo Tests - Region

	(1)	(2)	(3)	(4)	(5)
	1 week early	1 month early	2 months early	1 month late	2 months late
placebo	-55.046	-883.510	476.872	422.916	-165.949
	(333.48)	(671.79)	(1098.90)	(424.18)	(356.80)
placebo_lag1	-381.418	-53.947	-618.709	2003.713	11.852
	(618.73)	(217.97)	(612.11)	(1128.16)	(247.78)
placebo_lag2	-984.936	-1168.092*	-26.755	50.986	1436.589***
	(541.80)	(510.72)	(458.54)	(925.05)	(302.94)
placebo_lag3	-961.343	130.801	-1566.041	-2797.677***	-254.537
	(603.55)	(484.92)	(1163.88)	(722.97)	(333.08)
placebo_lag4	-764.067	-828.406*	-535.389	-542.332	662.051
	(465.33)	(349.95)	(895.21)	(609.62)	(401.94)
placebo_lag5	818.791	-1152.675	-388.534	396.936	83.309
	(1436.49)	(747.69)	(1208.20)	(548.55)	(206.67)
placebo_lag6	1032.607	-671.954**	-789.253	-759.333	1191.760
	(880.44)	(182.22)	(529.11)	(1680.09)	(586.25)
placebo_lag7	252.983	88.647	268.176	225.380	835.777**
	(257.85)	(838.98)	(697.38)	(1449.35)	(250.26)
calls_gross	0.103***	0.103***	0.103***	0.103***	0.103***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
transfer_val_out	0.529***	0.529***	0.529***	0.529***	0.529***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
_cons	737.229	737.553	737.288	737.174	736.841
r2	0.754	0.754	0.754	0.754	0.754
rmse	4025.238	4025.264	4025.233	4025.195	4025.252
N	74300.000	74300.000	74300.000	74300.000	74300.000

Outcome: Value of incoming airtime sent to people in district (in RWF; US$1=550RWF). Heteroskedasticity-robust SE's in parentheses (clustered at district level).

Appendix Table 5: Effect of flood on transfers

	(1) Pooled OLS	(2) OLS controls	(3) tower FE	(4) tower/Time FE
Flood Shock	1456.901	933.040**	1029.241**	1068.659**
	(770.84)	(316.98)	(329.36)	(375.45)
$DaysOfFlood_t$	774.798***	952.838***	981.247***	
	(166.92)	(230.79)	(206.75)	
$NearEpicenter_{it}$	263.474	237.740*		
	(919.80)	(88.55)		
Total calls		0.075***	0.065***	0.103***
		(0.00)	(0.01)	(0.01)
Outgoing transfers		0.678***	0.637***	0.527***
		(0.03)	(0.03)	(0.04)
R^2	0.000	0.702	0.729	0.753
N	171414	74895	74895	74895

"In flood region" defined as towers in the two districts affected by the flood. "Days of flood" are 9/12/07 - 9/18/07.

Appendix Table 6: Net transfers and wealth (Robustness to sender FE)

Earthquake Shock	9.727*
	(4.36)
Recipient Wealth x_i * Shock	5.702***
	(1.31)
Sender Wealth (x_j) * Shock	15.299
	(13.93)
Social Proximity S_{ij} * Shock	0.102
	(0.07)
Recipient Wealth (x_i)	0.847***
	(0.17)
Social Proximity (S_{ij})	0.069***
	(0.00)
Day dummies	Yes
Fixed effects	Sender j
Number of observations	10032721

Notes: Outcome is τ_{ijt}, the total airtime received by i from j on day t. Regressions include observations from the period January 4, 2009 through March 3, 2008. Wealth proxies x_i and x_j are computed using equations (14) and (15), as described in the text. S_{ij} is measured by counting the number of phone calls between i and j in the last three months of 2007. All regressions include $NearEpicenter_{it}$ and pairwise interaction terms (e.g. x_i * $NearEpicenter_{it}$, x_i*$DayOfQuake_t$); these coefficients are omitted for clarity. Standard errors, clustered by district, reported in parentheses. *** p<0.001, ** p<0.01, * p<0.05, + p<0.10.

Appendix Table 7: Network Flow as Social Proximity

	(1) No Fixed Effects	(2) Avg. Partial Effects	(3) Fixed Effects
Shock (recipient)	8.935*		8.480+
	(4.33)		(4.16)
Recipient Wealth x_i * Shock	6.791***		6.806***
	(0.75)		(0.95)
Sender Wealth x_j * Shock	16.501		15.766
	(14.74)		(13.66)
Network Flow S_{ij} * Shock	-0.918		-0.846
	(0.62)		(0.51)
Recipient Wealth (x_i)	1.146***	1.266***	
	(0.17)	(0.10)	
Sender Wealth (x_j)	1.915***	1.892***	
	(0.12)	(0.14)	
Network Flow (S_{ij})	0.048	0.054***	
	(0.03)	(0.02)	
Day dummies	Yes	Yes	Yes
Fixed effects	None	None	Dyad (directed)
Number of observations	10032721	10032721	10032721

Notes: Outcome is τ_{ijt}, i.e. airtime received by i from j on day t. x_i and x_j are estimated using equations (14) and (15), described in text. S_{ij} measures the number of distinct paths between i and j, where a path is defined by the call graph. All regressions include $NearEpicenter_i$ and pairwise interaction terms (e.g. x_i * $NearEpicenter_i$, x_i*$DayOfQuake_t$), but results are omitted for clarity. Standard errors, clustered by district, reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 8: Robustness of dyadic results

	(1)	(2)	(3)	(4)	(5)
Clustering	j's District	j's Tower	j	Day t	Unique senders
Earthquake Shock	9.794**	9.792***	9.892***	9.892***	7.591*
	(4.33)	(3.51)	(3.77)	(0.31)	(4.63)
x_i * Shock	5.763***	5.770**	6.011*	6.011***	0.145
	(1.26)	(2.89)	(3.39)	(0.40)	(0.22)
x_j * Shock	14.903	14.833	14.684	14.684***	8.929
	(13.33)	(12.12)	(11.86)	(0.97)	(7.23)
S_{ij} * Shock	0.110	0.111	0.110	0.110***	0.260
	(0.07)	(0.12)	(0.15)	(0.01)	(0.21)
Day fixed effects	Yes	Yes	Yes	Yes	Yes
Fixed effects	Dyad	Dyad	Dyad	Dyad	Dyad
N	10,032,721	10,032,721	10,251,136	10,251,136	4,430,160

Notes: Specification is identical to that used to produce column (3) of Table 6, with standard errors clustered according to column labels. Column (5) clusters by recipient, but restricts sample to allow only one recipient per sender. In cases were a single sender sends to multiple recipients, one recipient is chosen at random and the others are dropped from the analysis. Slightly fewer observations are contained in (1) and (2) because the tower closest to j is unknown for certain days. * p <0.05, ** p<0.01, *** p<0.001.

Appendix Table 9: Regression of Expenditures on Asset Ownership

Outcome	log(Expenditures)		Expenditures	
	β^a	(S.E.)	β^a	(S.E.)
Radio	0.18	(0.02)	40090	(13007)
Television	1.14	(0.01)	2130434	(44048)
Bed	0.24	(0.04)	187061	(8266)
Table	0.13	(0.01)	57601	(9109)
Car/Truck	0.24	(0.01)	1695284	(57718)
Motorcycle	0.65	(0.04)	8229976	(197091)
Bicycle	0.22	(0.11)	138186	(20359)
HH Size	0.09	(0.02)	56168	(3198)
R^2	0.62		0.75	
N	6900		6900	

Notes: Standard errors reported in parentheses.

Appendix Table 10: Summary statistics of phone use as computed from transaction logs

	Average
Panel A: Domestic and International Calls	
Activation date	1/12/08
Days of activity	770.3
Avg. call length	31.7
Calls per day	6.25
Net calls per day (out-in)	0.087
Int'l calls per day	0.084
Net int'l calls (out-in)	-0.014
Panel B: Social Network Structure	
Degree	734
In-degree	488.2
Out-degree	433
Daily degree	3.78
Net daily degree (out-in)	0.00027
Clustering	0.063
Betweenness	2.72
Panel C: Other Behaviors	
Credit used per day	163.5
Max. recharge value	2756.3
Avg. districts per day	1.36
Avg. districts contacted	1.21
N	901

Notes: Mean values reported, weighted by sampling strata to produce averages representative of entire phone population.

Appendix Table 11: Regression of predicted expenditures on phone use

	Coefficient	(S. E.)
Duration (outgoing)	-0.75	(2.55)
Duration (incoming)	7.66***	(2.01)
Degree	-1038.70*	(439.98)
Int'l duration (out)	-17.60	(10.07)
Int'l duration (in)	-10.63	(7.78)
Int'l degree	10534.70***	(3883.26)
Districts called	129304.11**	(42014.00)
Districts received	-103121.07**	(36356.38)
Unique towers	2918.14	(3600.54)
Months	11916.91	(9027.69)
Avg. recharge denomination	602.87	(461.13)
Daily recharge	716.18	(794.48)
N	671	
R^2	0.21	

Outcome is predicted expenditures $\widehat{Y_{id}}$, in RWF. Standard errors in parentheses. * significant at $p < .05$; $^{**}p < .01$;$^{***}p < .001$.

Appendix Table 12: Robustness check: net transfers and wealth

	(1) No Fixed Effects	(2) Avg. Partial Effects	(3) Fixed Effects
Earthquake Shock	9.803*		9.364*
	(4.67)		(4.38)
Recipient Wealth x_i * Shock	8.764**		8.616**
	(2.51)		(2.43)
Sender Wealth x_j * Shock	15.745		14.913
	(14.40)		(13.38)
Social Proximity S_{ij} * Shock	0.118^{+}		0.127^{+}
	(0.07)		(0.06)
Recipient Calling Activity * Shock	-0.004^{+}		-0.003*
	(0.00)		(0.00)
Recipient Wealth (x_i)	1.382***	1.416***	
	(0.21)	(0.12)	
Sender Wealth (x_j)	1.624***	1.622***	
	(0.11)	(0.13)	
Social Proximity (S_{ij})	0.043***	0.043***	
	(0.00)	(0.00)	
Recipient Calling Activity	-0.001***	-0.001***	
	(0.00)	(0.00)	
Day dummies	Yes	Yes	Yes
Fixed effects	None	None	Dyad (directed)
Number of observations	9984468	174657	9984468

Notes: Outcome is τ_{ijt}, i.e. airtime received by i from j on day t. x_i and x_j are estimated using equations (14) and (15), described in text. S_{ij} measures the number of phone calls between i and j in the three months prior to the earthquake. "Recipient Calling Activity" counts the total number of calls made by i in 2007. All regressions include $NearEpicenter_i$ and pairwise interaction terms (e.g. x_i * $NearEpicenter_i$, x_i*$DayOfQuake_t$), but results are omitted for clarity. Standard errors, clustered by district, reported in parentheses. *** p<0.001, ** p<0.01, * p<0.05, + p<0.10.

Figures

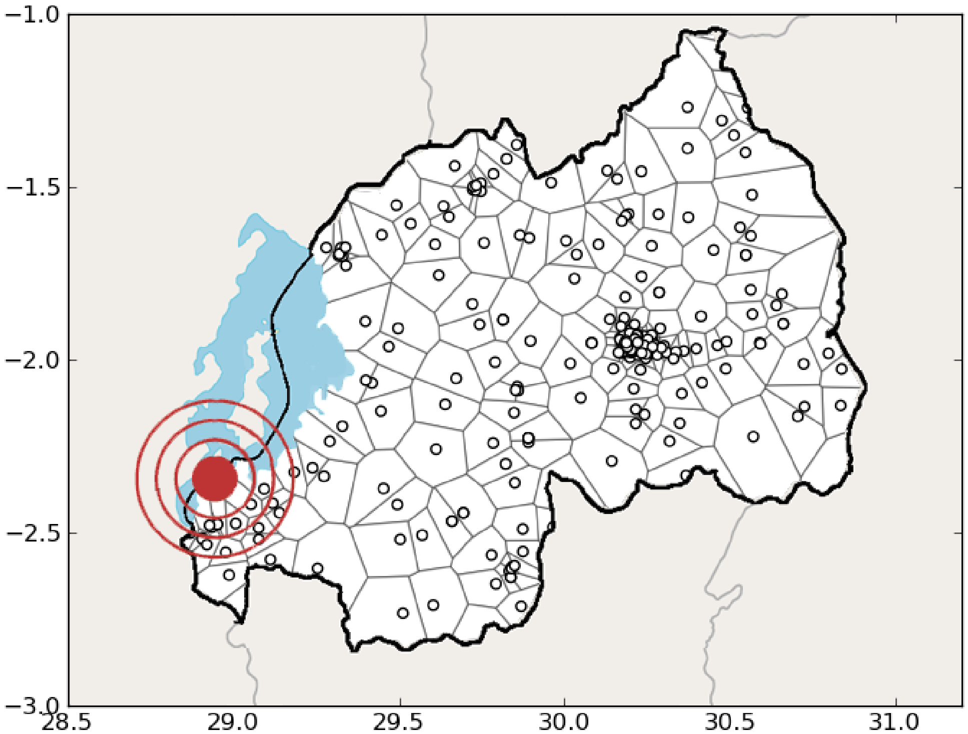

Figure 1: Map of Rwanda showing the location of mobile phone towers (as of February 2008) and the location of the Lake Kivu earthquake of 2008. Each black dot represents a cell tower, with the approximate area covered by the tower demarcated by adjacent Voronoi cells. The epicenter of the earthquake is shown with red concentric circles.

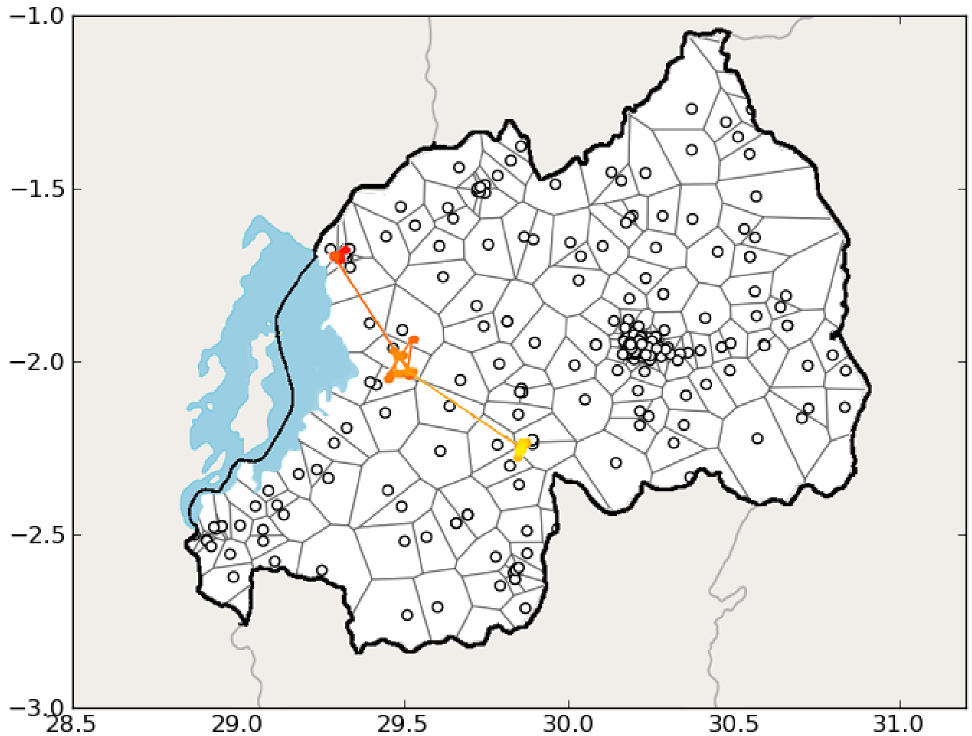

Figure 2: Map of Rwanda showing a single individual's inferred trajectory over a 6 month period. Although the individual only makes a small number of phone calls, the locational inference algorithm is able to roughly assign the user a continuous trajectory through time to places not restricted to the set of known locations of cellular towers. This particular individual is observed to slowly migrate southeastward, with early locations colored dark red and later locations colored yellow.

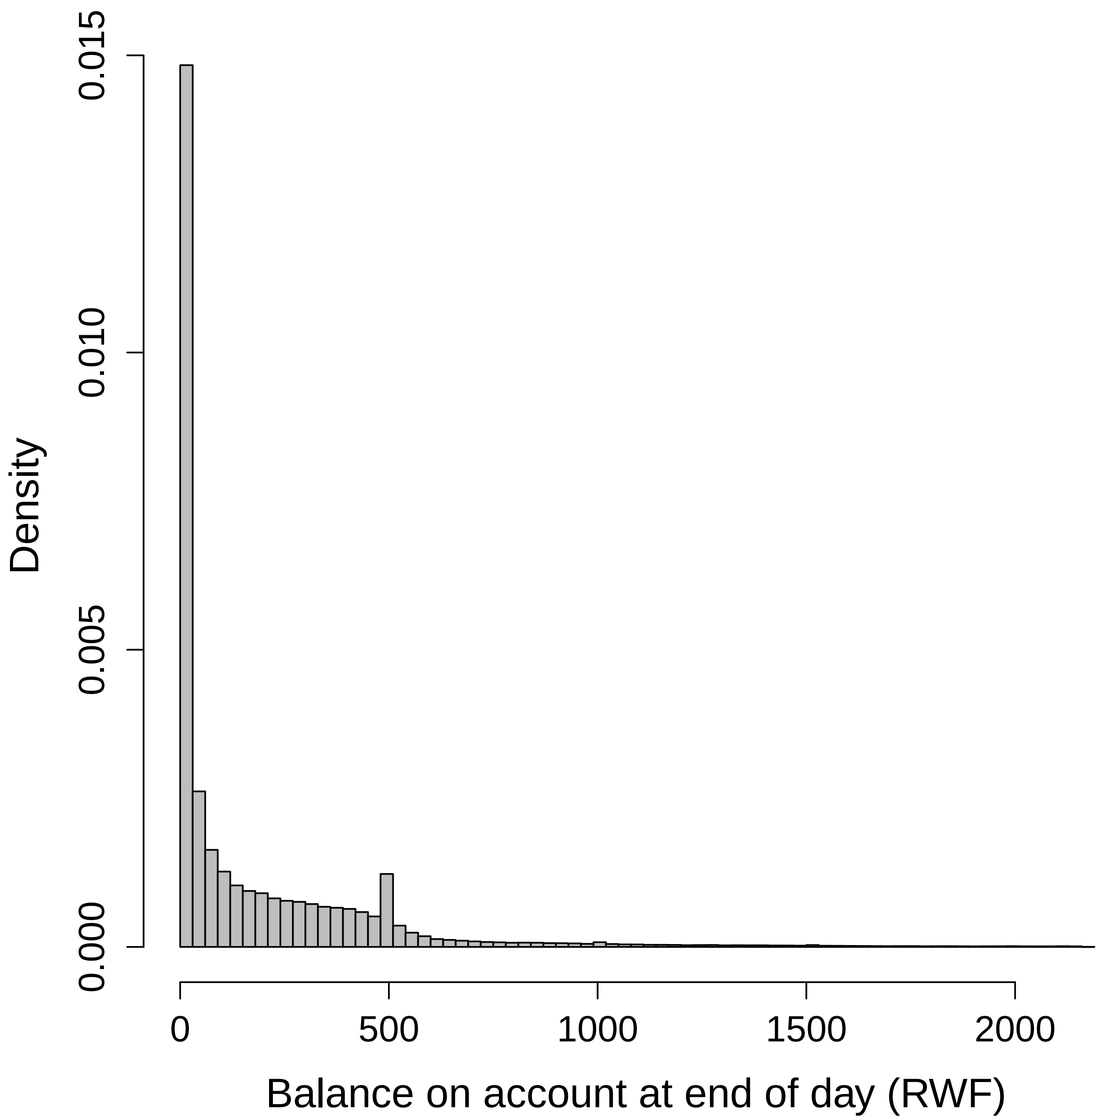

Figure 3: Distribution of end-of-day account balances on February 1, 2008, the day before the Lake Kivu earthquake. Subscribers with balances in the top and bottom 0.5 percent of all users were removed to improve the display of the histogram (one user had a balance of 442500RWF).

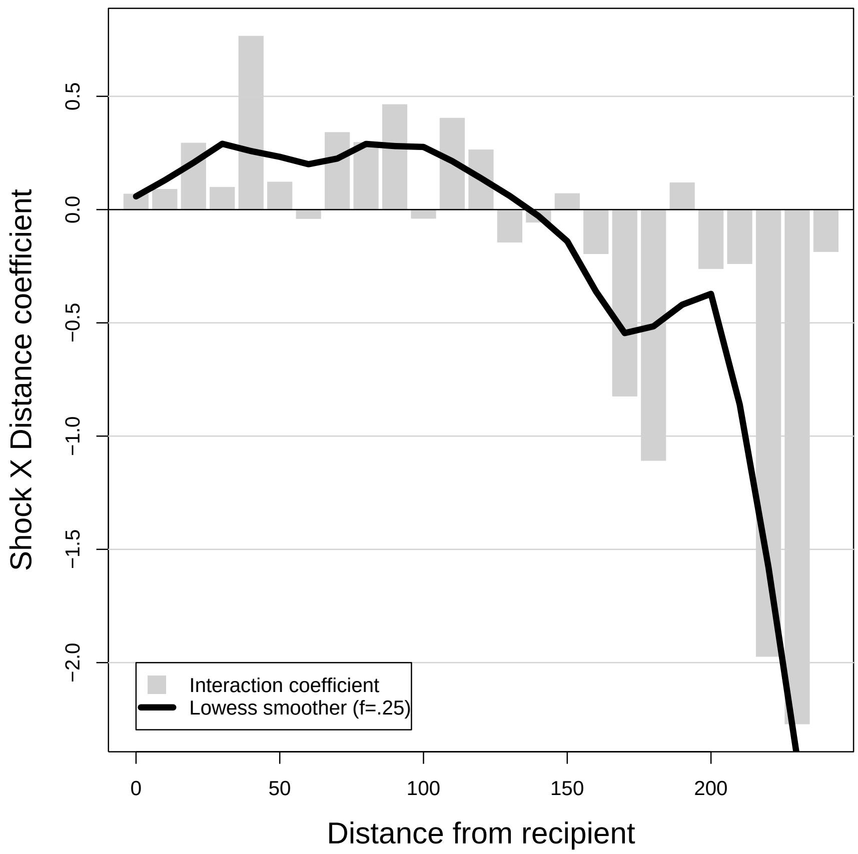

Figure 4: Relationship between the geographic structure of an individual's network and her propensity to receive a transfer after the earthquake.

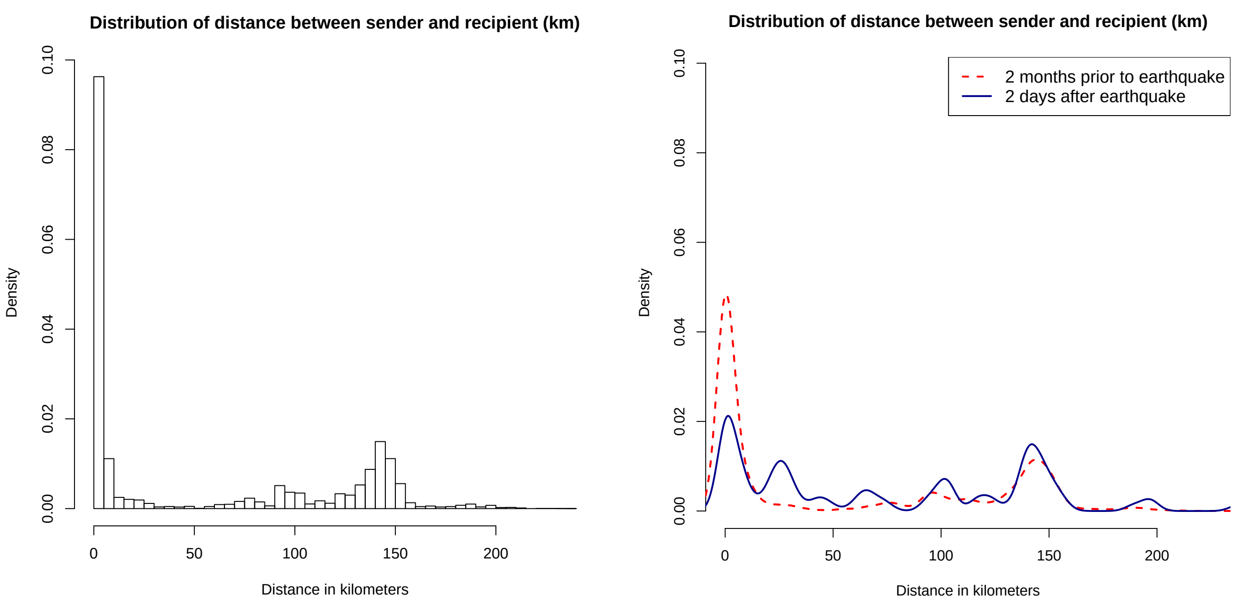

(a) Transfer distances over 4 years

(b) Transfer distances just before/after quake

Figure 5: Distribution of distances over which transfers are sent to and from the earthquake region.

Appendices

A History-Dependence in Reciprocity

[Appendix A not for publication]

We provide a formal derivation of the history-dependence of current period transfers in the model of dynamic limited commitment presented in Section 3.1(ii). The exposition is based on Ligon et al. (2002) and Foster & Rosenzweig (2001), but makes explicit the extension to the case where single-period utility includes a component of altruism. Other models of reciprocity, and in particular the preference-based reciprocity of Rabin (1993) and Falk & Fischbacher (2006), are predicated on fundamentally different "reciprocal" motives, but yield similar empirical predictions. We employ the enforced/instrumental model because the comparative statics that result are most directly testable with the data at our disposal, but do not mean to imply that the behavior we observe is necessarily motivated by this particular type of reciprocity vs. a different type of reciprocity.

We rely on a model with two agents i and j with von Neumann-Morgenstern utility where i's single period utility in increasing in j's single period utility according to (4), and vice versa. In static equilibrium, non-zero transfers occur when i's marginal utility of consumption is less than the γ-weighted utility of j, i.e., when

$$u_i'(x_{it}) < \gamma u_j'(x_{jt}) \tag{A.1}$$

or when the converse applies to j.[25] Call this static transfer, which depends on the state s of the world, τ_{jit}^N, where the superscript N denotes that this is the static Nash equilibrium. It is easily seen that whenever (A.1) holds, i will transfer a non-zero τ_{jit}^N that satisfies

$$\frac{u_i'(x_{it} - \tau_{jit}^N)}{u_j'(x_{jt} + \tau_{jit}^N)} = \gamma \tag{A.2}$$

When the converse of (A.1) applies (i.e. $u_j'(x_{jt} \leq \gamma u_i'(x_{it}))$, the transfer will be negative, and in all other cases the transfer will be zero.

In the repeated-game setting, agents are infinitely lived but are unable to save across periods. Given uncertainty as to the state of the world that will be realized in the future, both i and j can potentially be made better off through (possibly negative) state-contingent transfers τ_{jit}. The current-period utility of i is then the sum of single period utility plus the expected discounted utility of future interaction with j, where

[25] For simplicity we assume that i and j are similarly altruistic, i.e., $\gamma_{ij} = \gamma_{ji} = \gamma$.

δ is the discount factor:

$$U_{it}^T = \underbrace{\underbrace{u_i(x_{it} - \tau_{jit})}_{\text{own consumption}} + \underbrace{\gamma u_j(x_{jt} + \tau_{jit})}_{\text{altruistic benefit}}}_{\text{single period utility}} + \underbrace{E \sum_{s=t+1}^{\infty} \delta^{s-t}[u_i(x_{is} - \tau_{jis}) + \gamma u_j(x_{js} + \tau_{jis})]}_{\text{continuation value of relationship}} \tag{A.3}$$

Though i and j may agree ex ante to a set of state-contingent transfers, imperfect ability to enforce contracts implies limited commitment ex post. After the state of the world is realized, either agent can renege at any time, in which case both i and j will revert to the static Nash equilibrium where $\tau_{jit} = \tau_{jit}^N$ as given by (A.2). Then, i's utility is simply

$$U_{it}^A = u_i(x_{it} - \tau_{jit}^N) + \gamma u_j(x_{jt} + \tau_{jit}^N) + E \sum_{s=t+1}^{\infty} \delta^{s-t}[u_i(x_{is} - \tau_{jis}^N) + \gamma u_j(x_{js} + \tau_{jis}^N)] \tag{A.4}$$

The set of sustainable contracts are then defined by the implementability constraints that require utility for both agents to be (weakly) greater than under the static Nash equilibrium, i.e., $U_{it}^T \geq U_{it}^A$ and $U_{jt}^T \geq U_{jt}^A$. These constraints specify a superset of sustainable contracts, a subset of which are efficient.

In the repeated game, the transition matrix Π specifies the probabilities $p_{sr} = Pr(X_{t+1} = r | X_t = s)$ of transitioning from state s at time t to state r at time $t+1$, where S is the state space of the Markov chain. The set of constrained-efficient contracts along the Pareto frontier can then be characterized as

$$U_{js}(U_{is}) = \max_{\tau_{jit}} \left\{ u_j(x_{jt} + \tau_{jit}) + \gamma u_i(x_{it} - \tau_{jit}) + \delta \sum_{r \in S} p_{sr} U_{jr}(U_{ir}) \right\} \tag{A.5}$$

Where U_{is} denotes the expected discounted utility of i given state s at time t, which j is required to satisfy. The following conditions must be met for (A.5) to be optimal:

$$\begin{aligned}
\lambda: &\quad \left[u_i(x_{it} - \tau_{jit}) - u_i(x_{it} - \tau_{jit}^N)\right] + \gamma\left[u_j(x_{jt} + \tau_{jit}) - u_j(x_{jt} + \tau_{jit}^N)\right] + \delta \sum_{r \in S} p_{sr} U_{ir} \geq 0 \\
\delta p_{sr} \phi_r: &\quad U_{ir} \geq 0 \\
\delta p_{sr} \mu_r: &\quad U_{jr}(U_{ir}) \geq 0 \\
\psi_i: &\quad x_{it} - \tau_{jit} \geq 0 \\
\psi_j: &\quad x_{jt} + \tau_{jit} \geq 0
\end{aligned}$$

With $u_i(\cdot)$ and $u_j(\cdot)$ concave, $U_i(\cdot)$ and $U_j(\cdot)$ are also concave, implying the following first order conditions

(A.6-A.7) and envelope condition (A.8):

$$
\begin{aligned}
\frac{u_j'(x_{jt}+\tau_{jit})+\gamma u_i'(x_{it}-\tau_{jit})}{u_i(x_{it}-\tau_{jit}+\gamma u_j'(x_{jt}+\tau_{jit})} =& \lambda + \frac{\psi_i-\psi_j}{u'(x_{it}-\tau_{jit})+\gamma v'(x_{jt}+\tau_{jit})} && \text{(A.6)}\\
-V_r'(U_r) =& \frac{\lambda+\phi_r}{1+\mu_r} && \text{(A.7)}\\
\lambda =& \frac{u_j'(x_{jt}+\tau_{jit})+\gamma u_i'(x_{it}-\tau_{jit})}{u_i'(x_{it}-\tau_{jit})+\gamma u_j'(x_{jt}+\tau_{jit})}\\
=& -V_s'(U_s) && \text{(A.8)}
\end{aligned}
$$

The optimal contract is thus characterized by the slope of the Pareto frontier, λ, which determines the extent to which i can transfer current-period utility to j. For each state, there exists a set of implementable points on the Pareto frontier between λ_s^{min} and λ_s^{max}.

The history dependence arises as follows. When i and j first enter a relationship at $t=0$, they agree upon a contract that specifies a fixed ratio of marginal utilities λ_0, which is a point on the Pareto frontier. Once the state of nature becomes known at $t=1$, i and j will attempt the τ_{jit}^* that maintains the ratio of marginal utilities at λ_0. For simplicity, assume that at $t=1$, j suffers a negative shock such that $x_{i1} > x_{j1}$, though the same logic applies when $x_{i1} \leq x_{j1}$. Then, in order to maintain λ_0, the agents will attempt $\tau_{jit}^* > 0$. However, if $\tau_{jit}^* > 0$ is large enough to make the implementability constraint on i bind, then i and j will implement $\tau_{jit} < \tau_{jit}^*$ in order to avoid the static Nash equilibruim τ_{jit}^N and preserve the continuation value of the relationship. The implemented τ_{jit} will be just sufficient to relax the constraint on i, and a new constrained-efficient ratio of single-period marginal utilities λ_t will be determined according to (A.7).

If the implementability constraints never bind (intuitively, this occurs when income covariance between i and j is high), then agents will continue to equate marginal utilities at λ_0 and there will effectively be no history dependence in τ_{jit}. However, when a constraint binds, the dynamic model allows the worse-off agent to sacrifice future consumption (by accepting a less favorable ratio of utility λ) in exchange for a transfer in the current period. This recalibration of λ affects all subsequent τ_{jit}.

B Robustness

B.1 Functional form assumptions

We briefly show that our central results are not sensitive to the precise econometric specification, or to the choice of time window (which in most regressions is restricted to the period starting one month before the earthquake and ending one month after the earthquake). Appendix Table 2 presents estimates of the average treatment effect of model (1) using the full dataset from October 2006 until July 2009 under a variety of econometric specifications. Column (1) gives the standard OLS results with no control variables X_{rt}, time fixed effects θ_t, or tower fixed effects π_r. Column (2) includes time-varying controls to account for regional variation in mobile phone use, column (3) adds regional fixed effects, and column (4) adds daily dummy variables. Across all specifications, the estimated effect of the shock remains strong and significant, and of a magnitude similar to that presented in Table 5.

B.2 Placebo tests

As a robustness check of the average treatment effect, we verify that the effects of the earthquake on transfers are unique to the day of the earthquake, and do not generally occur on days without significant economic shocks. We do this first at the district level, following the methodology used to produce Table 5. In Appendix Table 3, we include lag and lead terms to test whether there was a significant effect of the earthquake on transfer patterns in the days immediately before and after the earthquake. To identify these ten additional terms, we include district-level data from the full dataset as in Appendix Table 2. In column 1, we observe that this effect does not exist, and before the earthquake (lead1-lead3) and after the earthquake (lag1-lag7), there was no significant change in transfers to the affected regions. These results hold for lags and leads of up to 10 days. In columns (2) and (3), we see in contrast that national calls to the affected region increase in the days following the earthquake. International calls do not. Critically, there was no anomalous increase in any sort of mobile network traffic in the days prior to the earthquake.

Appendix Table 4 presents results from testing the same specification as in column 4 of Table 5 but with a "placebo" shock at the same location on different dates. Thus, we test for a spurious effect 1 and 2 months before, as well as 1 month after, the actual earthquake. In contrast to the results obtained for the date of the actual earthquake, we observe no significant change in transfers on the day of the placebo earthquakes.

B.3 Other large covariate shocks

The results presented so far provide strong evidence that Rwandans used the mobile phone network to send airtime to friends and families affected by a major earthquake, and that these results are robust to different empirical specifications. We now show that similar transfers are observed following other natural disasters.

During the period for which we have mobile phone data, there were no massive natural disasters on the scale of the Lake Kivu earthquake. However, there were two major floods that severely disrupted the lives of many Rwandans. These floods are not as well suited to our estimation strategy as the earthquake, since floods are less precisely located in space (there is no single epicenter), and the timing is only partially exogenous (prior weather patterns anticipate floods). Therefore, there are a priori reasons to expect that the effect of a flood on transfers would be less pronounced than the effect of an earthquake.

Nonetheless, we do observe a significant increase in transfers on the days following a severe flood. In Appendix Table 5, we estimate equation (1) for the towers in the region of a flood that killed 17 during September 2007. We find a modest but strongly significant increase in airtime sent to regions affected by the flood. In column (4) of Appendix Table 5, the point estimate is roughly half that of the corresponding point estimate of the effect of the earthquake (column 4 of Table 2).

B.4 Robustness of Dyadic Regressions

In the body of the paper, we use dyadic regressions to measure which types of individuals i and j are most likely to receive and send transfers in response to the earthquake. Tables 6 and 7 employ dyad-specific fixed effects to control for unobserved, time-invariant characteristics of the dyad. Thus, the coefficient on the interaction between x_i and $Shock_{it}$ in Table 6 indicates the extent to which wealthier individuals i are more likely to receive a transfer after the earthquake, *in relation to the normal activity observed between i and j.*

This specification, formalized in equation (13), reduces biases resulting from unobserved characteristics of i, j, and the dyad i-j that could be correlated with the error term ϵ_{ijt}. For instance, if i sends a large transfer to j on each day of the year (for reasons unrelated to economic shocks), the dyadic fixed effects will ensure that a large transfer sent from j to i on the day of the earthquake is not mis-attributed to the effect of the earthquake.

As a robustness check, we demonstrate that our key results hold if the regressions are estimated with a more parsimonious model that replaces the dyad-specific fixed effects (for each pair i-j) with sender-specific fixed effects (for each sender j). This model is slightly less restrictive, and more directly corresponds to the intuition that motivates the dyadic results, i.e., that we wish to identify the types of people i (where type is proxied by x_i) are chosen by j to receive a transfer, conditional on j's average behavior. Formally, we estimate the following model:

$$\begin{aligned}\tau_{ijt} \quad = \quad & \delta_0 + \delta_1 Shock_{it} + \delta_2 x_i + \delta_3 x_i Shock_{it} + \delta_4 x_j Shock_{it} + \\ & \delta_5 NearEpicenter_{it} + \delta_6 x_i NearEpicenter_{it} + \delta_7 x_j NearEpicenter_{it} + \\ & \delta_8 x_i DayOfShock_t + \delta_9 x_j DayOfShock_t + \theta_t + \pi_j + \varepsilon_{ijt} \qquad \text{(B.1)}\end{aligned}$$

While the full specification of (B.1) is somewhat unruly, all of the coefficients of interest are contained on the first line, represented by δ_1-δ_4. The remaining δ_5-δ_9 are sub-interaction terms that are included for consistency but which have limited real-world significance. To interpret, δ_1 indicates whether individuals affected by the earthquake (for whom $NearEpicenter_{it} = 1$ and $DayOfShock_t = 1$) are more likely to receive a transfers; δ_2 indicates whether wealthier individuals are more likely to receive more under normal circumstances; δ_3 indicates whether wealthier individuals receive more because of the earthquake; and δ_4 indicates whether wealthier individuals *send* more to friends affected by the earthquake. All estimates are conditional on the average amount sent by j.

Results from estimating model (B.1) are presented in Appendix Table 6. Only very minor differences exist between these results and those presented in Table 6. Transfers sent in response to the earthquake increase significantly in the wealth of the recipient, but are not significantly related to the wealth of the sender. In other words, holding the identity of the sender fixed, it is the wealthier individuals – not the poorer individuals predicted by models of charity – who are most likely to receive a transfer.

B.5 Measuring social proximity

In the regression results presented in the body, we control for the strength of the relationship between i and j with S_{ij}, which is simply the total number of calls made between i and j (in either direction) in the year prior to the window of time used in the regressions. We choose this metric as a simple and easy to compute statistic that is likely to be correlated with the overall social proximity of i and j. However, several other such measures of S_{ij} are also reasonable. In particular, Karlan et al. (2009) and Leider et al. (2009) suggest a related metric, *network flow*, which captures the number of distinct paths between i and j through third parties k. The intuition is that each common friend k increases the shared social collateral between i and j.

In Appendix Table 7, we show that our results are not sensitive to the specific measure of social proximity used. The estimates in Appendix Table 7, which utilize network flow to measure S_{ij}, are quite similar to the estimates in Table 6, which measure S_{ij} as the total number of prior phone calls. Controlling for network flow in the other regressions similarly has no effect.[26]

[26] These results can be provided by the authors on request.

B.6 Standard errors

As discussed in Section 2, for standard error estimates to be consistent in the dyadic regressions, they should ideally be cross-clustered by sender i and recipient j. This is because transfers involving the same individual are likely to be correlated with each other – e.g., if j transfer airtime to i, he is ceteris paribus less able to transfer airtime to others. In the results presented so far we have clustered standard errors by the district in which the recipient resides.

As a robustness check, Appendix Table 8 compares alternative methods of obtaining standard errors using different levels of clustering: no clustering (column 1), by recipient (column 2), by sender (column 3), and by date (column 4). Standard errors are largest when we cluster by recipient, but in all specifications the coefficients of interest are highly significant. In the last column of Appendix Table (8), we drop observations in such a way that each sender j appears only once. More precisely, whenever a sender j appears multiple times, only one dyad involving j is selected at random and kept for estimation purposes. This results in a smaller number of observations but it eliminates the problem of correlation of errors at the source. The standard error is larger – if only because we dropped observations – but the coefficient of interest remains significant.

C Estimating Wealth From Call Records

[Appendix C not for publication]

Our goal is to create a measure of an individual mobile subscriber's wealth or permanent income based on his history of mobile phone use. Intuitively, there is strong reason to suspect that information such as the volume of domestic or international calls, or the pattern of airtime purchases, might be strongly correlated with permanent income. Since our interest here is not in causal inference, we simply want to model a function $g(\cdot)$ that maps observable mobile behavior to unobservable economic status.[27] Although this is conceptually similar to a proxy-means test used for targeting (cf. Montgomery et al. 2000), we are not aware of any prior attempt to predict economic status using communications data.

Our analysis utilizes three different sources of data that are described in greater detail in Blumenstock & Eagle (2010):

1. *Rwanda Demographic and Health Survey (DHS)*: We use a standard Demographic and Health Survey (DHS) conducted by the Rwandan government to explore the relationship between consumption and asset ownership. This survey was conducted in 2005 by the Rwandan government on a large, representative set of 10,272 households. The survey contains roughly five hundred questions typical of Living Standard and Measurement Surveys, with detailed modules on demographic composition and socioeconomic status (de la Statistique du Rwanda (INSR) & Macro 2006). Most relevant to the current analysis, roughly seventy questions were asked about asset ownership and household expenditures, which makes it possible to estimate each household's annual expenditures in a manner following (Deaton & Zaidi 2002).
2. *Phone survey conducted by authors*: In 2009 and 2010, we conducted a phone survey of a geographically stratified group of Rwandan mobile phone users. Using a trained group of enumerators from the Kigali Institute of Science and Technology (KIST), a short, structured interview was administered to roughly 2,200 individuals. In addition to querying basic demographic information, the phone survey collected responses for a small subset of the DHS questions (described above) about household asset ownership and housing characteristics.
3. *Call Detail Records (CDR)*: As described in the main text, this large dataset contains a log of all phone activity by those individuals in the period from January 2005 to December 2008.

Utilizing these data, we follow a 3-step process to model the relationship between phone use and wealth:

[27] If there existed a sample of users for whom we had both wealth information and call history information, this would be a trivial exercise. However, in most developing countries, the phone companies do not collect the demographic or socioeconomic information of their customers.

C.1 Modeling the relationship between assets and expenditures

Given information on assets and housing characteristics, we seek to develop a scalar measure of economic status based on the "basket of goods" owned by the individual. We do this using data from the government Demographic and Health Survey (DHS), which contains detailed information on each household's assets, characteristics, and expenditures. Our general strategy is to create a model that maps the N assets and characteristics $(X_{1i}, ..., X_{Ni})$ of household i to the same household's expenditures Y_i:

$$Y_i = f(X_i^1, ..., X_i^N) \tag{C.1}$$

where $f()$ is a flexible function that can be paramaterized in various ways. We opt for a parsimonious approach similar which models expenditures as a weighted function of $(X_{1i}, ..., X_{Ni})$:

$$Y_{id} = \gamma + \sum_{\alpha} \sum_{A} \beta^a X_{Ai}^{\alpha} + \mu_d + \epsilon_{id} \tag{C.2}$$

where expenditures Y_{id} of household i in district d is a weighted polynomial function of the assets and characteristics X_{Ai} of i, where the weights β^a reflecting each asset's relative contribution to total expenditures. We allow for district-specific intercepts μ_d. To reduce the potential bias of outliers, we remove outliers with abnormally large studentized residuals, following a standard process described in (Fox 1997).[28] Appendix Table 9 gives the coefficients from the linear terms that result from estimating (C.2) It is evident that annual expenditures are heavily correlated with asset ownership.

C.2 Predicting the expenditures of phone survey respondents

After estimating (C.2) on the DHS data, we obtain a vector of coefficients $\widehat{\beta_A}$ that can be used to predict total expenditures given knowledge of assets and housing characteristics X^a. Thus, for any individual in Rwanda, we could in principle predict that individual's annual expenditures, denoted by $\widehat{Y_{id}}$, by asking that individual a small number of questions about his household.[29] Through our phone-based surveys, we collect information on the assets and housing characteristics (i.e., the X_A) which maximize the predictive power of Equation C.2. To determine which X_A to collect in the phone surveys, we ran feature selection algorithms on the DHS data – in rough terms, this process identifies the characteristics which will maximize the R^2 of $g(\cdot)$.[30]

[28] Our results change very little if we use an alternate technique for removing outliers, such as removing the top 1% or 5% of extreme values.

[29] In practice, there are some outlandish assumptions that must be made to justify this approximation, but we will defer discussion of these and other limitations until the end of the Appendix.

[30] Full details on the process of feature selection, and the characteristics found to most predictive of annualized expenditures, are available from the authors.

C.3 Relating call histories to predicted expenditures

Using the above technique, it is possible to obtain the *predicted expenditures* $\widehat{Y_{id}}$ for each of the individuals contacted in the phone survey. This gives us a total of roughly 2,200 individuals for whom we have a measure of predicted expenditures and detailed call history information (obtained from the mobile operator). For these individuals, it is then possible to directly evaluate the relationship between phone use and economic status:

$$\widehat{Y_{id}} = g(CDR_i) \tag{C.3}$$

Finding the optimal form of $g(\cdot)$ is an important research topic, but is not our primary focus. Again, we opt for a parsimonious approach and estimate a fixed-effects regression of predicted expenditures on a large number of aggregate statistics of mobile phone use in a manner similar to (C.2). We compute a large number of metrics of phone use CDR_i. To minimize potential endogeneity of τ_{ijt}, we exclude all data on interpersonal transfers, and instead focus on each subscribers history of phone calls, text messages, recharge ("top-up") purchases, and a large number of social network metrics. Some of the characteristics include:

- *Activation date*: The date on which the phone first appears in the transaction logs.
- *Days of activity*: The number of different days on which the phone was used.
- *Net calls*: Number of outgoing calls minus the number of incoming calls.
- *Degree*: Number of unique contacts with whom the person communicated (called or received a call).
- *Daily degree*: Average number of unique people contacted on any given day, conditional on phone use.
- *Recharge*: Monetary value deposited on SIM card.
- *In/Out-degree*: Number of different people to whom/from whom, calls were made/received.
- *Clustering*: Percentage of first-degree contacts that have contacted each other.
- *Betweenness*: Average shortest path between the user and 50 randomly sampled numbers.
- *Interpersonal transfers*: Total airtime transfers (number sent + number received).
- *Districts*: Number of political districts in which the phone was used. Rwanda has 30 districts.

Appendix Table 10 summarizes these metrics for a representative random sample of mobile subscribers, and Appendix Table 11 presents the results from regressing predicted expenditures on twelve different metrics of phone use (polynomial terms excluded). The R^2 of 0.21 – which increases to 0.39 after including polynomial terms, district fixed effects, and several other metrics of phone use – indicates a strong relationship between expenditures and phone use.

C.4 Potential endogeneity and other limitations

The identification strategy used in Section 5.2 (i) relies on a proxy for permanent income, x_i, to measure the relative socioeconomic status of the population of mobile phone subscribers. In utilizing the $\widehat{\widehat{Y_{id}}}$ as computed above, one potential concern is that our x_i is more a proxy of aggregate phone use or technological sophistication than actual economic wealth. We take three steps to minimize this possibility. First, as noted above, in estimating (C.3) we exclude all metrics related to use of the airtime transfer service. Second, we use a highly non-parametric model to estimate (C.3) that includes district fixed effects and second- and third-order polynomial terms. As a result, there are strong non-monotonicities in (C.3) such that $\widehat{\widehat{Y_{id}}}$ is increasing in some measures of phone use, decreasing in others, and in general is highly non-linear. This ameliorates the concern that $\widehat{\widehat{Y_{id}}}$ simply captures aggregate phone use. Finally, we can explicitly control for different measures of aggregate phone use in our regressions that include $\widehat{\widehat{Y_{id}}}$.

Appendix Table 12 presents the results of re-estimating the partial effect of wealth on transfers received (Table 6 in the main text), while additionally controlling for "Recipient Calling Activity," which is simply the total number of calls made by the recipient in 2007. We assume that this measure is likely to be more directly correlated with the recipient's general propensity to utilize the phone than our non-parametric measure of wealth $x_i = \widehat{\widehat{Y_{id}}}$. Comparing the results of Appendix Table 12 with Table 6, we note that all coefficients have the same sign and statistical significance. In particular, the correlation between recipient wealth and transfers received remains positive and significant.[31]

Apart from the above concern, there are several limitations worth highlighting in the approach we have taken to estimating wealth from call records. We have glossed over distinctions between income, expenditures, permanent income, and wealth, which are of material consequence in evaluating poverty and development (cf. Deaton & Muellbauer 1980). Also problematic is the possibility that asset-based proxies for expenditures may provide biased estimates of the expenditures of certain types of individuals. For instance, if a strong correlation is found between television ownership and assets among the aggregate population, but a small subgroup of the population has a distaste for television, this simple method would systematically underestimate the expenditures of that subgroup. Finally, one particularly troubling assumption we must make is that the relationship between assets and expenditures identified with the function $f(\cdot)$ in the 2005 DHS data will remain constant when applied to phone survey data collected in 2009 and 2010. This assumption is unjustified for at least two distinct reasons. First, the data for the two populations was collected using very different methodologies, and respondents may respond differently to questions about assets depending on whether they are asked in person or over the phone. Second, the data was collected in different years, and it is possible that the relationship between assets and expenditures would evolve over such a long interval. For instance, the strong relationship observed in 2005 between television ownership and wealth may be weaker in 2010, as electricity becomes more available and used televisions saturate the market.

[31] In separate results, available from the authors, we control for other measures of calling activity (such as the sum of calls made and received). These alternate specifications produce very similar point estimates and standard errors to those presented in Appendix Table 12.

In noting these weaknesses and still using $\widehat{\widehat{Y_{id}}}$ as a proxy for x_i, we do not mean to imply that this method will provide an accurate or unbiased measure of an individual's wealth. However, we hope that amidst all of the noise there is a decipherable signal that lets us infer the relative wealth of individuals in our mobile phone dataset, for whom we only have anonymous records of phone use.