Fraying of the ties that bind: Community-level financial institutions and HIV/AIDS with evidence from KwaZulu Natal, South Africa

Benjamin Linkow*
NORC at the University of Chicago

Lucas Rentschler
Centro Vernon Smith de Economía Experimental
Universidad Francisco Marroquín

Abstract

This paper provides a theoretical and empirical investigation of the effect of the HIV/AIDS pandemic on community-level informal financial institutions such as rotating savings and credit associations. Our theoretical model illustrates that the mortality risk implied by such a mortality shock limits the scope of informal contracts, leading to more exclusive institutions. Using panel data from the high-prevalence area of KwaZulu Natal, South Africa, we find that mortality at the community level substantially reduces the propensity to participate in informal financial institutions in ways that are consistent with the predictions of our theoretical model.

*Corresponding author.

1 Introduction

For many of the world's poor, community-level informal arrangements provide an important source of access to financial services. Where the formal sector is constrained by information or enforcement constraints, informal arrangements often arise to meet unmet demand for savings, credit, and insurance. A vivid illustration of the important role of these institutions and the wide variety of needs they serve for the poor is provided by Collins et al. (2009). Village-based savings clubs and lending groups provide not only credit for investment or other large purchases, but also play a crucial role in facilitating insurance and consumption smoothing in the face of fluctuating incomes for the poorest of the poor.

These types of institutions share a common rationale that explains their existence and sometimes persistence even in the face of formal alternatives. In particular, the informational advantages of the community can serve as a basis for determining membership of such institutions (Ghatak and Guinnane (1999)). Further, community-level punishments such as ostracism or loss of reputation incurred from breach of an informal agreement can serve as powerful deterrents (Besley et al. (1993)). Particularly where the state is weak, these informal sanctions may be more effective than the legal recourse available to formal service providers. Similarly, community members may have access to information (e.g., creditworthiness) about one another through repeated interactions or networks that would be costly or impossible for impersonal entities to obtain.

This paper will argue that a heretofore unexplored impact of the HIV/AIDS pandemic in sub-Saharan Africa is that it reduces access to informal financial services by weakening the basis for community-level contract enforcement. The central argument is that premature mortality constitutes a risk of informal contractual non-compliance; an individual who has agreed to repay a loan or fulfill some other obligation is unable to do so if he or she dies unexpectedly. This risk of default due to mortality is one that the community cannot mitigate via the threat of sanctions. This is of particular concern as the viability of these institutions may in some cases depend on the deterrent power of such sanctions ensuring very high rates of compliance.

We find empirical support for the assertion that mortality weakens community-level financial institutions using panel data from KwaZulu Natal, South Africa, an area of high HIV prevalence. Our data show that high mortality is associated with less participation in the types of community groups that involve informal contracts. However, mortality does not inhibit participation in other types of community groups, which is consistent with the hypothesis that mortality acts on community institutions by weakening informal contract enforcement.

In the next section, we discuss conceptual issues and provide a review of the relevant literature. We then provide a theoretical model that illustrates the effects of community-level mortality on community-level institutions. We present our empirical results in the next section, followed by policy implications and recommendations for further research.

2 Community level contracting and mortality

2.1 Community level institutions and mortality

The informal institutions that are the subject of this paper are underpinned by rules that are made and enforced at the level of the community, as opposed to formal, legal institutions underpinned by rules at the level of the state. As Bowles and Gintis (2002) elaborate, institutions based on community-level rules may be more effective or efficient than formal alternatives in circumstances where formal contracting is difficult or costly, or where private information is important. They identify three characteristics of communities that can put them at an advantage in these situations. First, transactions that take place within communities are likely to involve repeated interactions. This gives rise to a particular set of incentives to cooperate over time that may not exist in a perfectly anonymous market. The potential benefit to future cooperation or the threat of terminating the relationship or retaliation can thus add an element of self-enforcement to contracts. This has been observed by Maher (1997), who notes that among a sample of European firms that extend supplier credit, many report that they would not pursue legal action in the event of default.

Secondly, repeated interactions allow community members to have information about one another that might not be available publicly. For example, community members may be able to assess creditworthiness more effectively than a credit bureau could. Monitoring costs and principal-agent problems can also be reduced by information that may be available to the community, but not the market.

Finally, communities imply a scope for imposing punishments to enforce rules. Ostracism and the threat of social sanctions or loss of social status can be powerful deterrents to opportunistic behavior. Contracts enforced by community level sanctions may thus be advantageous when formal contracting is poorly enforced, costly, or inflexible. The importance of social sanctions has been observed by other authors as well. Gächter and Fehr (2000) find experimental evidence that introducing the ability to punish free riders improves cooperation in coordination games. In a study of rural Kenya, Miguel and Gugerty (2005) find that communities with a greater scope for imposing social sanctions on members experience improved public goods provision.

Where HIV/AIDS introduces a substantial mortality risk to the members of the community, all three of these advantages of community governance structures are undermined. The risk of mortality to one's transaction partner means that repeated interaction is less likely–the expected number of interactions with a given transaction partner is smaller than it would otherwise be. The potential gains that would occur as a result of future interactions are consequently reduced, and hence the incentive to behave opportunistically in the present is increased.

Similarly, the community's advantage in terms of information becomes less valuable. In the context of credit, for example, high mortality introduces a risk of default that is unrelated to trustworthiness or other characteristics that can be assessed by the community. To the extent that mortality is unpredictable, the community's relative advantage over a formal financial institution in assessing creditworthiness is thus reduced.

The sanctioning power of the community is also affected by high levels of mortality. Informal punishments often imply a time dimension–the costs of social ostracism, for example, are experienced over a period of time rather than instantaneously. In addition, in the case of contracts that involve a commitment to take some action in the future, mortality introduces a risk of non-compliance that sanctions cannot deter. This is potentially significant, as in many cases the threat of sanctions is sufficient to ensure high levels of compliance. For example, a number of authors have noted that default in the context of rotating savings and credit associations is exceedingly rare (e.g., Van den Brink and Chavas (1997)). To the extent that the rationale behind these institutions depends on high levels of compliance, their viability could be threatened by introducing a risk of non-compliance where none had existed before.

The presence of HIV/AIDS may also have effects on the demand for participation in informal institutions, as the mortality risk facing an individual may alter their incentives to participate. Affected individuals may have an increased incentive to participate with the intention of defaulting on the incurred obligations, since informal punishments that are experienced over time become less of deterrent. In this way, HIV/AIDS may increase the demand for certain kinds of institutions.

Conversely, to the extent that the benefits of participating in some types of institutions are experienced over time, HIV/AIDS may reduce the demand on the part of the affected individuals. For example, community level financial institutions are often used to finance the purchase of lumpy consumer durables such as washing machines or refrigerators (Besley et al. (1993)). The shorter time horizon implied by higher mortality implies less utility to be derived from these sorts of purchases.

2.2 ASCAs and ROSCAs

These community-level institutions often provide financial services to community members. Two types of such informal financial institutions that have received significant attention in the literature, and are the focus of the empirical analysis in this paper, are Accumulated Savings and Credit Associations (ASCAs), and Rotating Savings and Credit Associations (ROSCAs). In an ASCA, members contribute savings to a common fund, which is then lent out at interest. The responsibility for disbursing and collecting loans typically falls to the individual members, and the proceeds are divided up among the members. Hence, the ASCA acts as an interest-bearing savings vehicle.

ROSCAs are a community-level institution that have been observed in a wide variety of contexts in the developing world (Armendáriz and Morduch (2010)). The basic structure of a ROSCA is that a group of individuals commit to gathering at regular intervals and each contributes a predetermined amount of money into a fund. At each meeting, the fund is allocated to a single member of group, and the meetings continue until each member has obtained the pot once. ROSCAs may repeat over several cycles, and the method of choosing the order of allocation varies. A number of motivations for joining ROSCAs have been noted in the literature, including financing the purchase of lumpy consumer durables (Besley et al. (1993)), shielding savings from claims by relatives (Anderson and Baland (2002)), and as a commitment device to overcome time-inconsistent preferences (Ambec and Treich (2007) and Gugerty (2007)).

Both of these institutions imply a scope for opportunistic behavior and hence an important role for community-level sanctions as a deterrent. In an ASCA, such opportunistic behavior would take the form of members withholding loan repayments they have collected from the rest of the group. Similarly, in a ROSCA, once an individual has been allocated the pot, they can profit by failing to attend and contribute at subsequent meetings of the group. Anderson et al. (2009) demonstrate theoretically that a ROSCA structure cannot be incentive compatible for all members in the absence of some form of sanctioning to deter this behavior. Thus, as a source of default that is immune to sanctioning, mortality threatens these institutions.

2.3 Empirical literature on HIV/AIDS and community spillovers and social capital

The effects of HIV/AIDS on these types of institutions has yet to receive explicit attention in the literature to our knowledge. However, two recent empirical studies find evidence that is consistent with our central thesis. Jayne et al. (2006) consider the impacts of mortality (to which HIV/AIDS is a major contributor) at the level of the community in Zambia. They find that communities with higher rates of mortality exhibit lower productivity, income, and area under cultivation. However, they do not investigate the mechanisms by which this might occur. Intriguingly, they also find that the reduction in income associated with mortality is of greater magnitude in communities that have experienced greater rainfall variability. An interpretation of their results that is consistent with the approach here is that mortality weakens informal risk sharing networks, leading to greater vulnerability to rainfall shocks.

Similarly suggestive cross-country evidence is provided by David (2007). Controlling for a variety of factors, he finds that incidence of HIV/AIDS has a strong inverse relationship with subjective measures of trust. He thus concludes that mortality acts to weaken social capital, and hypothesizes that a mechanism by which this occurs is through the strain on traditional social networks that mitigate economic shocks–i.e., by reducing the strength of informal agreements such as those enforced at the community level.

3 Theoretical model

In this section, we present a model of community-level institutional formation that allows us to incorporate the effects of community level-mortality. There are $N > 3$ members of the community, who are assumed to have homogeneous preferences, with the exception of heterogeneous discount factors. Let $\delta_i \in [0, 1]$ be the discount factor of community member i. Some subset of these individuals may choose to form a group G. Once a group is formed, the members interact in two stages. In the first stage, Each $i \in G$ chooses whether to contribute an amount $A > 0$ or not. This choice is observable, and in the second stage, those group members who chose not to contribute incur an informal punishment $P > 0$.[1]

If there are $n \leq N$ members of a group, and $m \leq n$ members choose to contribute, then each member of the group obtains $mA\alpha$, where $0 < \alpha < 1$ and α is a parameter that captures the return on contributions. Note that in the absence of a punishment mechanism members have an incentive to default on their obligations. Taking punishment into account, note that if

$$A\alpha < A - \delta_i P$$

[1]The magnitude of P is intended to reflect both the costs of incurring social stigma, as well as the costs of being excluded from future financial arrangements. Thus, in practice P may be endogenous to an individual's mortality risk, since higher mortality risk implies that the cost of future exclusion will be lower. However, the main conclusions of the model are unchanged by our simpler specification that P is exogenous, and the effect of mortality is reflected by δ_i.

then group member i will opt to default.[2] This condition can be re-written as

$$\delta_i < \frac{A(1-\alpha)}{P}.$$

Heterogeneous discount factors are independent and identically distributed. This distribution of discount factors is assumed to be such that there is a probability mass of k at 0, and elsewhere the distribution is uniform on $(0, 1]$.[3] Thus, the probability distribution of the discount factors is given by

$$F(\delta_i) = \begin{cases} k & if\ \delta_i = 0 \\ \delta_i (1-k) & if\ \delta_i \in (0,1] \\ 1 & if\ \delta_i > 1. \end{cases}$$

This implies that there is some proportion k of the population who completely discount the future. We take this parameter k to represent the level of mortality within the community. Conceptually, an individual who suffers premature mortality is not able to live up to the terms of her informal contract. From the standpoint of its impact on the welfare of the group, this is equivalent to defaulting.[4]

We assume complete information except with regard to discount factors. While the distribution F is public information, each individual's discount factor is assumed to be private information. Associated with each individual is a publicly observable signal $\hat{\delta}_i$ which is equal to δ_i with probability θ.[5] With probability $1-\theta$, $\hat{\delta}_i$ is equal to a random draw from F. In the case where $\theta = 1$, for example, the signal is perfectly accurate and discount factors are in effect publicly observable. Conversely, when $\theta = 0$, the indicator conveys no information about an individual's discount factor beyond knowledge of F. For intermediate values of θ, we can say that $Pr(\delta_i = a) < Pr\left(\delta_i = \hat{\delta}_i\right) < 1$ for all $a \neq \hat{\delta}_i$.[6] To ensure that the decision

[2]It is important to note that a group member may be able to elude this punishment. For example, a group member could migrate to another community. This would likely reduce or eliminate the severity of the realized punishment. For simplicity, we neglect this possibility in our theoretical model. However, we will address the issue of migration in the empirical analysis.

[3]The use of a uniform distribution with an atom at 0 is for tractability. However, the intuition of the model is unchanged if a different distribution is employed.

[4]The use of a discount factor of zero to model the mortality shock is a simplification. In the interest of clarity we opted to implement the mortality shock in as simple a way as possible. However, our results are unchanged if we model the shock in a different way. In particular, as long as the probability of default is increased for some subset of the community, and there is some information partially indicating who the affected individuals are, our main results will hold.

[5]The precision of information available to the community is modeled by θ. Note that as θ decreases, the available information decreases. There is likely to be observable factors that correlate with an individuals true discount factor and this may introduce heterogeneity in the quality of information across individuals. For simplicity, we neglect this in the model.

[6]If our theoretical model were taken as a stage game and repeated, then past behavior of group members would be observable. Thus, in every round the observable estimates of individual discount factors could be updated. In the limit, this would tend to increase the quality of information regarding individual discount factors. In the event of a mortality shock (such that some subset of the community has their discount factor reduced to zero, this shock is known to have occurred, but the identity of affected individuals is unknown), the expected discount factor of all individuals would be reduced. We are interested in the effect of such a

regarding whether or not to admit a potential member to a group is non-trivial, we assume parameter values such that an individual with $\hat{\delta}_i = 1$ is expected to contribute, and an individual with $\hat{\delta}_i = 0$ is expected not to contribute.[7] Since we are interested in the effect of an increase in the mortality rate on the institution in equilibrium, we will restrict attention to scenarios in which admitting a contributing member into the group is efficient in equilibrium. This ensures that the risk of default is the relevant factor in determining whether or not to admit additional members, rather than concerns about efficiency.[8] To ensure the viability of the institution, we will also restrict attention to scenarios in which P is large enough that an individual with a discount factor equal to the average discount factor in the population would not default. That is, $\frac{A(1-\alpha)}{P} < \left(\frac{1-k}{2}\right)$.

We assume that the group is selective, and will only admit members who, conditional on $\hat{\delta}_i$, are expected to (weakly) increase the expected payoff of other members. This is equivalent to saying that a potential member is admitted if she is expected to choose to contribute.[9] This condition can be stated as

$$\begin{aligned} E\left(\delta_i \mid \hat{\delta}_i\right) &\geq \frac{A(1-\alpha)}{P} \\ \theta\hat{\delta}_i + \frac{(1-\theta)(1-k)}{2} &\geq \frac{A(1-\alpha)}{P}. \end{aligned}$$

This inequality implies a threshold signal $\hat{\delta}^*$such that any individual with $\hat{\delta}_i \geq \hat{\delta}^*$ is accepted into the group, while individuals with $\hat{\delta}_i < \hat{\delta}^*$ are turned away. This threshold is given by

$$\hat{\delta}^* = \frac{A(1-\alpha)}{\theta P} - \frac{(1-\theta)(1-k)}{2\theta}.$$

The expected equilibrium size of a group is determined by

$$n^* = NPr\left(\hat{\delta}_i \geq \hat{\delta}^*\right) = N(1-k)\left[1 - \frac{A(1-\alpha)}{\theta P} + \frac{(1-\theta)(1-k)}{2\theta}\right].$$

Although n^* need not be an integer, the realized equilibrium group size will be, and depends on the realized publicly observable signals of community members.

mortality shock on the equilibrium group size. For simplicity, we consider a one-shot game. Our results would hold in a repeated environment, although the analysis would be substantially complicated.

[7] These assumptions can be restated as $\frac{A(1-\alpha)}{P} < \theta + \frac{(1-\theta)(1-k)}{2}$ and $\frac{A(1-\alpha)}{P} > \frac{(1-\theta)(1-k)}{2}$, respectively.

[8] Denoting the equilibrium group size as n^*, this ensures that $\alpha > \frac{1}{n^*}$. This restriction can also be stated as $\alpha > \frac{\sqrt{(k-1)N\left((k-1)N(P(\theta+(\theta-1)k+1)-2A)^2-16A\theta P\right)}-(k-1)N(P(\theta+(\theta-1)k+1)-2A)}{4A(k-1)N}$.

[9] For simplicity, we only allow for the possibility of excluding an individual. In practice, a group has intermediate options to reduce the risk of default (such as forcing a marginal case towards the end of a ROSCA cycle).

3.1 Comparative statics

We first consider the equilibrium selectivity of the institution. Of particular interest is the effect of an increased mortality rate on this selectivity. Since

$$\frac{\partial\hat{\delta}^*}{\partial k} = \frac{(1-\theta)}{2\theta} > 0,$$

equilibrium selectivity is increasing in the mortality rate. Further, since

$$\frac{\partial\hat{\delta}^*}{\partial\theta} = \frac{-A(1-\alpha)}{P\theta^2} + \frac{2\theta(1-k)}{4\theta^2} + \frac{2(1-\theta)(1-k)}{4\theta^2} < 0,$$

the equilibrium selectivity of a group is decreasing in the accuracy of the observable signals regarding discount factors. Lastly, since

$$\frac{\partial\hat{\delta}^*}{\partial P} = \frac{-A(1-\alpha)}{\theta P^2} < 0,$$

increasing punishment reduces the need for selectivity.

Next, we turn attention to the expected equilibrium size of the group. Note that

$$\frac{\partial n^*}{\partial k} = -N(1-k)\left(\frac{1-\theta}{2\theta}\right) - N\left(1-\left(1-\frac{A(1-\alpha)}{\theta P} + \frac{(1-\theta)(1-k)}{2\theta}\right)\right) < 0.$$

This result follows from the fact that, by assumption, an individual with $\hat{\delta}_i = 0$ is expected not to contribute.[10] Intuitively, this result shows that as mortality increases, the group size will diminish. This comparative static is the main result that we will test. It is important to note that this result shows that an institutional response to increased mortality is to become more exclusive. That, is, holding the rules of the institution constant, the equilibrium size of the groups shrinks as mortality increases. However, in the long run, one might expect that the rules of the institution could evolve in response to increased mortality. For example, institutions could adopt rules such that, in the event of a default, family members of the defaulter are liable. Studying the long run institutional response to changes is a promising avenue for future research.

Next, note that

$$\frac{\partial n^*}{\partial\theta} = \left(\frac{N(1-k)}{\theta^2}\right)\left(\frac{1-k}{2} - \frac{A(1-\alpha)}{P}\right) < 0.$$

Thus, as the accuracy of the signal decreases, so will the equilibrium group size. Lastly, note that

$$\frac{\partial n^*}{\partial P} = N(1-k)\left(\frac{A(1-\alpha)}{\theta P^2}\right) > 0.$$

[10] That is, that $\frac{A(1-\alpha)}{P} > \frac{(1-\theta)(1-k)}{2}$.

Thus, as the level of punishment for defaulting increases, the group size is actually increasing.

4 Data and econometric approach

The results from the previous section present a number of difficulties from the standpoint of generating empirically testable hypotheses. Many of the parameters, as well as $\hat{\delta}^*$, are fundamentally unobservable. However, we can make use of existing data to investigate the behavior of optimal group size n^*, and in particular the finding that n^* is decreasing in the rate of mortality.

To do so, we use household survey data from the KwaZulu Natal Income Dynamics Study (KIDS), a dataset that has a number of desirable features for our purposes. It is a panel survey that spans a timeframe over which the rate of HIV/AIDS increased dramatically, resulting in substantial variation in community mortality rates. In addition, the data contain information on informal institutions at both the household and community levels. We can thus proxy the optimal group size n^* by taking as our dependent variable the probability that a household is a member of an ASCA or ROSCA, and considering the impact of the level of mortality in the community on membership probability.

4.1 Data

The KIDS dataset is a panel study collected in the province of KwaZulu Natal, South Africa over the period 1993-2004. It contains a range of detailed socioeconomic and demographic information intended to facilitate a range of policy-relevant research, particularly in terms of the dynamics of poverty. The initial 1993 round surveyed 1,354 households drawn from 67 communities; representatives of 74% of the original sample were successfully re-interviewed in both 1998 and 2004. Where core members of the 1993 households split off and formed or joined new households, these new households were also tracked and incorporated into the later rounds. The 1998 and 2004 surveys include questions on membership in a variety of community groups.[11] Since we are interested in the impact of community-level mortality and we can only identify households with their original communities, we omit households that have relocated to new communities during the survey. After adjusting for this and other data irregularities, we are left with 673 households distributed over 62 communities.

The location and timeframe are ideal for studying the effects of HIV/AIDS-related mortality on com-

[11] The 1998 survey also obtains retrospective data on group membership in 1993. However, recall bias appears to be a significant problem for investigating the dynamics of group membership here. For example, only 1 household retrospectively reported membership in a stokvel in 1993 that it had left in 1998. Conversely, there were 100 cases of households reporting membership in stokvels in 1998 that they had left by 2004. Hence, we omit the recall data from our analysis.

munity level institutions. The pandemic induced a massive increase in adult mortality over this period; the probability that a 15-year-old in KwaZulu Natal would not survive to age 49 increased from 27.9% in 1998 to 70.8% in 2004 (Dorrington et al. (2002)). It is important to note that we make no claims regarding the randomness of HIV/AIDS incidence. This introduces the possibility of omitted variable bias in our econometric analysis, since there could be an unobserved variable which drives HIV/AIDS. We are unable to definitively rule this possibility out, and this represents a potential weakness in our econometric analysis.

For each community, we calculate the rate of mortality among adults aged 15-49 in the sample over the periods 1993-1998 and 1998-2004. The results confirm a dramatic increase in mortality after 1998, as prime age mortality nearly triples.

Our data contain information on membership in stokvels, a South African term for a general category of informal financial institution. Stokvels can take a variety of forms- ASCAs and ROSCAs are two of the most commonly noted in the literature but the term encompasses a wide variety of arrangements. Stokvels are widespread in the study area; 25.1% of households in our sample reported membership in at least one.

4.2 Household level econometric approach and estimation results

At the level of the household, our empirical strategy is to estimate the probability that at least one member of a household belongs to a particular type of group as a function of the level of mortality in the community. We would like to control for both household and other community-level characteristics to the greatest extent possible, and hence employ panel data methods. This presents an econometric difficulty in terms of how to eliminate unobserved heterogeneity in the context of a binary response dependent variable. Our preferred approach is fixed effects conditional logit proposed by Chamberlain (1980). We prefer the conditional logit model to a standard probit specification because the latter implies a number of assumptions that are problematic given our data.[12]

The empirical approach is to condition the observed pattern of responses over time on the total number of responses within the panel unit. In our case, there are only two time periods, so that the model is:

$$P\left(g_{i2}=1 \mid X, Z, g_{i1}+g_{i2}=1\right)=\frac{exp\left(\beta_1\left(X_{i2}-X_{i1}\right)+\beta_2\left(Z_{j2}-Z_{j1}\right)\right)}{1+exp\left(\beta_1\left(X_{i2}-X_{i1}\right)+\beta_2\left(Z_{j2}-Z_{j1}\right)\right)}$$

[12]Using probit necessitates a random as opposed to fixed effects panel approach, as parameters cannot be estimated consistently under a probit specification with fixed effects. The main disadvantage of the random effects probit model is that it requires the assumption of independence of the unobservables and the explanatory variables (Wooldridge (2010)). In our case, this is particularly problematic. For example, the household level incidence of prime-age mortality variable is almost certainly related to unobserved household characteristics.

and

$$P\left(g_{i1}=1 \mid X, Z, g_{i1}+g_{i2}=1\right)=1-\frac{exp\left(\beta_1\left(X_{i2}-X_{i1}\right)+\beta_2\left(Z_{j2}-Z_{j1}\right)\right)}{1+exp\left(\beta_1\left(X_{i2}-X_{i1}\right)+\beta_2\left(Z_{j2}-Z_{j1}\right)\right)}.$$

Where g_{it} is an indicator of whether household i belonged to group type g at time t, X_{it} is a vector of household-level time varying characteristics, Z_{jt} is a vector of community-level time varying characteristics including mortality, and β_1 and β_2 are vectors of parameters to be estimated.

In effect, the approach is to restrict attention to households that were members of a group in one period but not the other. We then compare households that joined a group between 1998 and 2004 to those that exited groups between 1998 and 2004 in order to see whether community level mortality and our other controls are associated with group exit.

Our household level control variables are expenditure growth, change in household size, and incidence of prime-age mortality within the household. Community level controls are expenditure growth, as well as dummy variables indicating whether the community experienced any of four different types of events. First, the gain or loss of public or NGO-provided services such as schools, health care, etc. Secondly, the gain or loss of a significant source of employment, either in the private sector or in the form of a public works project. Thirdly, reported outbreaks of violent unrest or crime. Fourth, we control for a range of adverse weather and other natural incidents including drought, flooding, hailstorms, and veldt fires. This is summarized in a dummy variable which is equal to one in the event of an adverse weather shock. Summary statistics for all relevant variables (including membership in burial societies and non-financial groups, which are discussed below) appear in Table 1. Note that it is possible that there are additional unobserved variables at the community level. This possibility means that our results may suffer from omitted variable bias. This suggests the need for further analysis of these issues with a richer dataset than in currently available. Such an analysis would be a fruitful avenue for future research.

Results of the conditional logit estimation are contained in Table 2. We report specifications with subsets of the full set of controls as a robustness check.[13] In all specifications the community level mortality coefficient is negative and significant. This finding is consistent with our theoretical prediction that mortality leads to higher default rates and hence greater selectivity on the part of the institutions. The only other control variable that is significant is household level expenditure growth. Since we have controlled for community level expenditure growth, the interpretation is that households that have done well relative to other households in their community are more likely to join stokvels.

Our use of a conditional logit model presents two potential concerns which we address by estimating a

[13]Marginal effects of the conditional logit model with the full set of controls are reported in the Appendix in Table 8.

linear probability model as robustness check. First, the conditional logit model omits households that were members of a group in either both periods or neither period. This gives rise to concern that our results might be biased because we are considering only a subset of the observations. Secondly, our conditional logit model is not able to account for unobserved variables that might vary over time, but not location, since the structure of the model precludes including a time dummy. Though linear probability models are problematic for statistical inference, parameter estimates are nonetheless consistent and unbiased. Thus, a linear probability model presents a useful robustness check on the conditional logit results, as it addresses both of the aforementioned concerns. The results of these estimations, alongside the conditional logit results for comparability, are reported in Table 3. As these results show, our findings are consistent across the two specifications.

4.3 Alternative interpretations and endogeneity concerns

Here, we consider two alternative explanations for the findings in the previous section. First, while our model pertains to the "supply" of institutional arrangements, it could be that demand side factors are in fact driving the results. If living in a high mortality community reduces demand for the services that stokvels provide, this could explain the observed relationship between mortality and participation.

We argue that this explanation is unlikely because of the lack of significance of the household-level prime age mortality coefficient in specification three in Table 2. If mortality were acting on stokvel participation by reducing demand, we would expect to observe a relationship between individual-level mortality risk and propensity to join a stokvel. We cannot observe individual-level mortality risk directly, but past mortality within the household may be a proxy for future risk. Our findings show that individuals whose households have suffered deaths in the past are no more or less likely to join stokvels than other households. To the extent that past incidence of mortality is a valid proxy for future risk, our results thus suggest that our main findings in Table 2 are driven by the mechanisms of the theoretical model, rather than demand side factors.

Secondly, we address the possibility that our findings could be driven by endogeneity bias. The structure of our model allows us to immediately rule out a number of potential sources of endogeneity. Since we are estimating the impact of a community-level characteristic on individual household behavior, reverse causality is implausible. In addition, household fixed effects eliminate the possibility of omitted variable bias due to any time invariant household or community level factor. We control for a number of time varying factors as well. Moreover, in an exercise reported in Appendix A, we show that changes in community-level mortality predicted by baseline levels of observed community-level characteristics do not explain our main result. This suggests that omitted community-level characteristics are less likely to drive this result.

However, a potentially important category of time varying unobservables remains. In particular, high rates of community mortality may lead households to reduce their propensity to participate in groups more generally. For example, mortality in the community may imply an increase in the marginal value of time as community members care for the sick, which could lead to lower rates of participation in group activities. Thus, reduced stokvel participation could be indicative of a reduced propensity to participate in groups more generally, rather than because of the impact on informal contracting that our model hypothesizes.

To investigate this possibility, we estimate models in which the dependent variable is participation in other types of community level groups. Our data include information on participation in a variety of such groups which we divide into two types. First, we consider burial societies. These are mutual insurance groups in which members assist one another in defraying the costs of funerals, which are a major expense in the study area. Membership in burial societies decreased from 38.1% in 1998 to 32.3% in 2004.

Burial societies would be expected to have some of the same incentives at work as the stokvels in our theoretical model, in particular to exclude riskier members of the community in response to higher mortality rates, thus reducing likelihood of participation. However, we would also expect that community mortality would imply an increase in demand for the services that burial societies provide, which would to some extent mitigate the negative impact of mortality on group participation. Thus, the predicted sign on the coefficient is ambiguous, but if the overall effect is negative, we expect the magnitude of the effect to be smaller than for stokvels.

Secondly, our data include membership in a variety of non-financial groups. These include trade associations and farmer's organizations, as well as civic groups such as school, water, and development committees. Also included are groups with a social or recreational purpose such as music and sports clubs. 20.2% of the study households belonged to these types of groups in 1998, and this increased to 26.2% in 2004. In non-financial groups, an increasing mortality rate of other members implies no pecuniary incentive to reduce group size as our model would predict for stokvels or burial societies. Thus, non-negative coefficients on mortality in the non-financial group regressions would support our claim of that the stokvel results are driven by the mechanisms of our theoretical model.

It could also be the case that community mortality has a different effect on non-financial group membership as compared to stokvels and burial societies. If, in addition to the financial services that the informal financial institutions provide, they also allow members to establish valuable ties within the community, then an increase in community mortality that makes stokvels and burial societies more risky may lead individuals to substitute into alternative social-capital building opportunities, such as non-financial groups. Thus, ceteris paribus this substitution effect could cause community-level mortality to lead to an increase in the propensity to join non-financial groups.

Table 4 contains estimates for models with both burial societies and non-financial groups as the dependent variable, respectively. All specifications include the full set of controls available, although results are robust to using subsets of these controls. We report estimates obtained via conditional logit and the linear probability model, both with household level fixed effects.[14] In no case is the coefficient for community-level mortality statistically significant, which suggests that our main findings are not driven by some unobserved variable that reduces propensity to join groups more generally.

While these coefficients are not statistically different from zero, their magnitude is non-trivial. However, their signs and magnitudes are consistent with our expectations. In particular, the coefficients in the burial society model are negative and smaller in magnitude than the coefficient in the stokvel model. This is consistent with the hypothesis described above that similar mechanisms to those that reduce membership in stokvels are also at work in burial societies, but the impact of mortality on burial society membership is attenuated by demand-side factors. Additionally, the non-negative coefficients on non-financial institutions could be the substitution effect related to building social capital described above.

It is important to bear in mind that these results are suggestive only. In order to conclusively establish that the effect of community-level mortality differs between stokvels and other groups, we would need to show that the corresponding coefficients are statistically different. To test such a null hypothesis requires estimates of the covariance between these estimated coefficients. As we are unable to jointly estimate the relevant conditional logits (and jointly estimating the corresponding linear probability models is problematic for the purposes of statistical inference) we are unable to conduct such an analysis. While our results are consistent with the hypothesis that the observed relationship between mortality and stokvel participation is due the effects hypothesized in our theoretical model, the absence of a joint test to this effect remains a limitation of our analysis.[15]

4.4 Further evidence on the relationship between mortality and stokvel participation

In the previous section, we showed that households in high-mortality communities are less likely to be members of stokvels. According to the theoretical model, this could potentially occur in two ways. Stokvels may become gradually more exclusive, as the optimal threshold for inclusion increases with mortality. Alterna-

[14]Marginal effects of the conditional logit models are reported in the Appendix in Table 8.

[15]The fact that community-level mortality does not have a significant effect on membership in burial societies also suggests that our results are not driven by migration out of communities. If an increase in migration were to drive both community-level mortality (via an increase in HIV risk) and a decrease in stokvel membership (via a reduction in the availability of social sanctions) then we would expect to see a similar effect on membership in burial societies. As discussed above, we are unable to provide a formal test of the null hypothesis that the effect of community-level mortality on group membership does not differ between stokvels and burial societies.

tively, at high enough levels of k a cooperative equilibrium may cease to exist and the group may dissolve. The KIDS dataset includes information on the number of groups of various types that serve each of the 62 communities in the survey. We thus estimate the effects of mortality on the number of stokvels at the community level.

The results of our community-level fixed effects regression appear in Table 5. Our sample size is small, as we have only 124 data points upon which to rely. The coefficient on prime age mortality is positive and insignificant. Thus, higher mortality communities do not appear to have fewer stokvels. The implication, then, is that our household-level results are driven by existing stokvels admitting fewer members, rather than dissolving.

5 Conclusions and suggestions for further research

Our findings suggest a heretofore unexplored impact of the HIV/AIDS pandemic that has important implications for policy. In high prevalence areas, programs designed to mitigate the effects of the HIV/AIDS pandemic must consider not only those directly affected by the disease, but the broader community as well. Even for those who are not directly affected by disease, the pandemic may weaken the informal institutional arrangements upon which many poor households rely. Access to credit, insurance, and other financial services may suffer as a result.

An important caveat to our empirical results is that we are unable to definitively rule out the possibility that the observed (negative) relationship between stokvel membership and community-level mortality is driven by a decrease in the demand for such institutions. That is, it is possible that demand for such institutions decreases as community level mortality increases. However, we find that the incidence of past household-level mortality does not significantly affect stokvel membership. Further, we find no evidence that the prevalence of burial societies and other types of informal institutions are affected by community-level mortality.[16] These two results suggest that that, as predicted by our theoretical model, supply side effects drive our findings. However, this issue remains a possible concern, and demonstrates the need for additional research with data at the institutional level.

Our results have implications for the study of institutional change. Social relations in developing countries are complex and interconnected. However, our results demonstrate that an analysis of the underlying incentives of a particular type of institutional arrangement can nonetheless provide useful insights. In particular, our results suggest that in the short run, informal institutions can react quickly to changing circum-

[16] As discussed above, we are unable to test that the effect of community level mortality on institutions differs between stokvels and other types of institutions.

stances. In the long run, we would expect the institutions themselves to evolve.

We identify three directions for future research that would be promising. First, while our empirical evidence has focused on rotating savings and credit associations, a broad range of institutional arrangements such as mutual insurance networks and informal lending are liable to be subject to the same effects. Further research could investigate the effects of HIV/AIDS on other institutions with similar incentives at work. Secondly, a central component of our argument is that mortality risk due to HIV/AIDS serves to shift time preferences in favor of shorter-term outcomes. This hypothesis could be investigated in a more general way, for example by studying the impact of HIV/AIDS on investment behavior. Finally, while this paper has considered the impact of mortality on exclusivity, other changes to institutional arrangements may be possible in response to increased risk, such as restricting new members to later positions in the order of allocation, or introducing joint or family liability. These issues could be explored with more detailed data at the level of the institutions themselves. Further research along these lines would allow for the magnitude and economic implications of the effects of HIV/AIDS on weakening community-level contract enforcement to be more precisely understood.

References

Stefan Ambec and Nicolas Treich. Roscas as financial agreements to cope with self-control problems. *Journal of Development Economics*, 82(1):120–137, 2007.

Siwan Anderson and Jean-Marie Baland. The economics of roscas and intra-household resource allocation. *Quarterly Journal of Economics*, 117(3):963–995, 2002.

Siwan Anderson, Jean-Marie Baland, and Karl Ove Moene. Enforcement in informal saving groups. *Journal of Development Economics*, 90(1):14–23, 2009.

Beatriz Armendáriz and Jonathan Morduch. *The economics of microfinance*. MIT press, Cambridge, Massachusetts, 2010.

Timothy Besley, Stephen Coate, and Glenn Loury. The economics of rotating savings and credit associations. *American Economic Review*, 83(4):792–810, 1993.

Samuel Bowles and Herbert Gintis. Social capital and community governance. *Economic Journal*, 112 (483):F419–F436, 2002.

Gary Chamberlain. Analysis of covariance with qualitative data. *Review of Economic Studies*, 47(1):225–238, 1980.

Daryl Collins, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven. *Portfolios of the poor: How the world's poor live on $2 a day*. Princeton University Press, Princeton, New Jersey, 2009.

Antonio David. Hiv/aids and social capital in a cross-section of countries. *World Bank Policy Research Working Paper*, (4263), 2007.

Rob Dorrington, Debbie Bradshaw, and Debbie Budlender. *HIV/AIDS profile in the provinces of South Africa: Indicators for 2002*. Centre for Actuarial Research, Medical Research Council and the Actuarial Society of South Africa, 2002.

Simon Gächter and Ernst Fehr. Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994, 2000.

Maitreesh Ghatak and Timothy Guinnane. The economics of lending with joint liability: theory and practice. *Journal of Development Economics*, 60(1):195–228, 1999.

Mary Kay Gugerty. You cant save alone: Commitment in rotating savings and credit associations in kenya. *Economic Development and Cultural Change*, 55(2):251–282, 2007.

Thomas Jayne, Antony Chapoto, Elizabeth Byron, Mukelabai Ndiyoi, Petan Hamazakaza, Suneetha Kadiyala, and Stuart Gillespie. Community-level impacts of aids-related mortality: Panel survey evidence from zambia. *Review of Agricultural Economics*, 28(3):440–457, 2006.

Maria Maher. Transaction cost economics and contractual relations. *Cambridge Journal of Economics*, 21 (2):147–170, 1997.

Edward Miguel and Mary Kay Gugerty. Ethnic diversity, social sanctions, and public goods in kenya. *Journal of public Economics*, 89(11):2325–2368, 2005.

Rogier Van den Brink and Jean-Paul Chavas. The microeconomics of an indigenous african institution: The rotating savings and credit association. *Economic Development and Cultural Change*, 45(4):745–772, 1997.

Jeffrey Wooldridge. *Econometric analysis of cross section and panel data*. MIT press, Cambridge, Massachusetts, 2010.

Table 1: Summary statistics

	1993-1998	1998-2004
Community level means		
Prime age mortality incidence	3.3%	7.3%
Per capita income growth	-5.3%	1.3%
Loss of public or NGO-provided services	11.7%	11%
Gain of public or NGO-provided services	41.6%	36.2%
Departure of major employer or public works project	12.9%	43.1%
Arrival of major employer or public works project	21.8%	49.1%
Adverse weather shock	49%	49.5%
Outbreak of violent unrest or crime	28.7%	8.2%
Household level means		
Prime age mortality incidence	12%	29.4%
Household size change	0.93	-0.82
Per capita income growth	-14.1%	5.4%
Stokvel membership	23.1%	22.6%
Burial society membership	38.1%	36.2%
Non-financial group membership	20.2%	26.1%

Table 2: Conditional logit estimates of stokvel membership at the household level

	(1)	(2)	(3)
Community-level mortality	-5.047*	-5.909*	-6.203*
	(2.452)	(2.821)	(3.112)
Household-level income growth	0.662**	0.623**	0.638**
	(0.129)	(0.133)	(0.140)
Community-level average income growth		0.407	0.404
		(1.029)	(1.028)
Loss of public or NGO-provided services		-0.284	-0.250
		(0.387)	(0.378)
Gain of public or NGO-provided services		0.223	0.223
		(0.277)	(0.282)
Departure of major employer or public works project		0.023	0.006
		(0.282)	(0.280)
Arrival of major employer or public works project		0.129	0.116
		(0.225)	(0.225)
Outbreak of violent unrest or crime		-0.325	-0.332
		(0.341)	(0.341)
Adverse weather shock		-0.116	-0.118
		(0.233)	(0.234)
Household-level prime age mortality			0.02
			(0.278)
Household size change			-0.012
			(0.039)
Observations	392	392	392
Pseudo R^2	0.103	0.117	0.118

Estimates are obtained via conditional logits with household-level fixed effects
Standard errors (in parentheses) clustered at the community level.
$^{+}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01$

Table 3: Conditional logit and linear probability estimates of stokvel membership at the household level

	CL	LP
Community-level mortality	-6.203*	-1.019*
	(3.112)	(0.431)
Community-level average income growth	0.404	0.056
	(1.028)	(0.121)
Loss of public or NGO-provided services	-0.250	-0.053
	(0.378)	(0.041)
Gain of public or NGO-provided services	0.223	0.017
	(0.282)	(0.042)
Departure of major employer or public works project	0.006	-0.009
	(0.280)	(0.049)
Arrival of major employer or public works project	0.116	0.028
	(0.225)	(0.031)
Outbreak of violent unrest or crime	-0.332	-0.043
	(0.341)	(0.042)
Adverse weather shock	-0.118	-0.008
	(0.234)	(0.035)
Household-level prime age mortality	0.021	0.021
	(0.278)	(0.036)
Household-level income growth	0.638**	0.080**
	(0.140)	(0.018)
Household size change	-0.012	-0.002
	(0.039)	(0.004)
Year		-0.008
		(0.036)
Constant		0.288**
		(0.038)
Observations	392	1,346
Pseudo or within R^2	0.118	0.049

CL: conditional logit with household-level fixed effects
LP: linear probability model with household-level fixed effects
Standard errors (in parentheses) clustered at the community level.
$^{+}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01$

Table 4: Conditional logit and linear probability estimates of non-stokvel group membership at the household level

	Burial society		Non-financial	
	CL	LP	CL	LPs
	(1)	(2)	(3)	(4)
Community-level mortality	-3.393	-0.192	2.530	0.248
	(2.700)	(0.418)	(2.421)	(0.496)
Community-level average income growth	1.496	0.292	0.909	0.108
	(0.956)	(0.186)	(0.973)	(0.139)
Loss of public or NGO-provided services	0.483	0.022	-0.554	-0.088
	(0.369)	(0.049)	(0.343)	(0.063)
Gain of public or NGO-provided services	0.727**	0.147**	-0.244	-0.051
	(0.268)	(0.049)	(0.321)	(0.052)
Departure of major employer or public works project	0.334^{+}	0.058	0.217	0.024
	(0.194)	(0.037)	(0.301)	(0.063)
Arrival of major employer or public works project	-0.153	-0.011	0.063	0.012
	(0.228)	(0.041)	(0.230)	(0.045)
Outbreak of violent unrest or crime	-0.256	-0.048	-0.416	-0.057
	(0.413)	(0.056)	(0.364)	(0.059)
Adverse weather shock	0.102	0.034	0.134	0.007
	(0.213)	(0.035)	(0.195)	(0.031)
Household-level prime age mortality	-0.322	-0.025	-0.235	-0.054
	(0.236)	(0.038)	(0.268)	(0.042)
Household-level income growth	0.233^{+}	0.035	0.549**	0.074**
	(0.137)	(0.022)	(0.204)	(0.023)
Household size change	0.048^{+}	0.006	0.010	0.000
	(0.025)	(0.004)	(0.025)	(0.004)
Year		-0.033		0.015
		(0.034)		(0.035)
Constant		0.339**		0.255**
		(0.033)		(0.038)
Observations	466	1,346	460	1,346
Pseudo or within R^2	0.110	0.054	0.117	0.052

CL: conditional logit with household-level fixed effects
LP: linear probability model with household-level fixed effects
Standard errors (in parentheses) clustered at the community level.
$^{+}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01$

Table 5: Net change in number of informal financial institutions

Community-level mortality	22.019
	(11.664)
Community-level average income growth	1.979
	(2.849)
Loss of public or NGO-provided services	2.140
	(1.183)
Gain of public or NGO-provided services	0.384
	(0.955)
Departure of major employer or public works project	1.008
	(1.299)
Arrival of major employer or public works project	0.081
	(0.967)
Outbreak of violent unrest or crime	-0.139
	(1.113)
Adverse weather shock	0.802
	(0.789)
Year	-0.722
	(0.887)
Constant	-0.745
	(0.832)
Observations	126
Within R^2	0.181

Estimates are obtained via OLS with community-level fixed effects.
Standard errors in parentheses.
$^{+}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01$

A Predicted changes in community-level mortality

A potential concern with our main finding is that omitted community-level characteristics may drive our observed relationship between community-level mortality and stokvel participation. In particular, if changes in community-level mortality are not random, and are correlated with baseline levels of unobserved community-level characteristics, then our results could suffer from endogeneity bias. While we are unable to evaluate the degree of correlation with unobserved variables, we undertake an exercise in this Appendix using observable community characteristics to assess the potential for such bias. Table 6 demonstrates that changes in community-level mortality are correlated with the observed baseline levels of observed community-level variables.

Table 6: Pairwise correlation coefficients between changes in community-level mortality and baseline community-level observables

Community-level average income growth	0.016
Loss of public or NGO-provided services	0.205*
Gain of public or NGO-provided services	-0.328*
Departure of major employer or public works project	0.018
Arrival of major employer or public works project	-0.219*
Outbreak of violent unrest or crime	0.080*
Adverse weather shock	-0.090*
Household-level prime age mortality	0.025
Household-level income growth	-0.014
Household size change	-0.070

To evaluate whether correlation between changes in community-level mortality and baseline community-level observables can explain our main result, we predict changes in community-level mortality using baseline levels of community-level observables as the independent variables in the regression. We then construct a variable (which we call constructed community-level mortality) in which the second wave of community-level mortality is the fitted values of the aforementioned regression plus the baseline level of community-level mortality. We report estimates of stockvel membership using this constructed community-level mortality in Table 7. Note that the coefficients corresponding to constructed community-level mortality are insignificant in both the conditional logit and linear probability models. In fact, in the linear probability model, the coefficient is positive. In the conditional logit estimation the coefficient is relatively close to the corresponding coefficient in the burial society results reported in Table 4. This shows that correlation changes in community-level mortality and the baseline levels of observed community-level characteristics does not explain our results. It then seems likely that unobserved community-level characteristics are not driving our results.

Table 7: Estimates of stokvel membership where changes in community-level mortality are predicted using baseline community-level observables

	CL	LP
Constructed community-level mortality	-3.941	0.260
	(4.108)	(0.903)
Community-level average income growth	0.356	0.058
	(0.955)	(0.112)
Loss of public or NGO-provided services	-0.373	-0.086^{+}
	(0.357)	(0.048)
Gain of public or NGO-provided services	0.225	0.035
	(0.289)	(0.045)
Departure of major employer or public works project	0.091	0.020
	(0.305)	(0.050)
Arrival of major employer or public works project	0.138	0.045
	(0.226)	(0.033)
Outbreak of violent unrest or crime	-0.193	-0.049
	(0.341)	(0.049)
Adverse weather shock	-0.060	0.005
	(0.235)	(0.035)
Household-level prime age mortality	-0.100	0.008
	(0.271)	(0.036)
Household-level income growth	0.641**	0.081**
	(0.131)	(0.017)
Household size change	-0.000	-0.001
	(0.040)	(0.004)
Year		-0.066
		(0.054)
Constant		0.232**
		(0.045)
Observations	392	1,346
Pseudo or within R^2	0.113	0.042

Constructed community-level mortality is the baseline community-level mortality plus the fitted values of a regression where the dependent variable is the observed change in community-level mortality, and the independent variables are the baseline levels of the other community-level observables.
CL: conditional logit with household-level fixed effects
LP: linear probability model with household-level fixed effects
Standard errors (in parentheses) clustered at the community level.
$^{+}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01$

B Marginal effects

Table 8: Marginal effects of the conditional logit estimates of stokvel, burial society and non-financial group membership

	Stokvel	Burial society	Non-financial
Community-level mortality	-1.487*	-0.847	0.632
	(0.678)	(0.677)	(0.606)
Household-level income growth	0.153**	0.058^{+}	0.137**
	(0.034)	(0.034)	(0.051)
Community-level average income growth	0.0967	0.374	0.227
	(0.247)	(0.239)	(0.243)
Loss of public or NGO-provided services	-0.059	0.119	-0.136^{+}
	(0.088)	(0.088)	(0.081)
Gain of public or NGO-provided services	0.053	0.180**	-0.061
	(0.069)	(0.064)	(0.080)
Departure of major employer or public works project	0.002	0.083^{+}	0.054
	(0.067)	(0.048)	(0.075)
Arrival of major employer or public works project	0.028	-0.038	0.016
	(0.055)	(0.057)	(0.058)
Outbreak of violent unrest or crime	-0.078	-0.064	-0.103
	(0.076)	(0.103)	(0.088)
Adverse weather shock	-0.028	0.025	0.034
	(0.055)	(0.053)	(0.049)
Household-level prime age mortality	0.005	-0.080	-0.059
	(0.067)	(0.058)	(0.067)
Household size change	-0.003	0.012^{+}	0.003
	(0.009)	(0.006)	(0.006)
Observations	392	466	460

Marginal effects are estimated assuming that fixed effects are equal to zero.
Standard errors (in parentheses) clustered at the community level.
$^{+}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01$