UC Berkeley
CUDARE Working Papers

Title
Inferring informal risk-sharing regimes: Evidence from rural Tanzania

Permalink
https://escholarship.org/uc/item/50f6t3fh

Authors
Li, Zhimin
Ligon, Ethan

Publication Date
2020-09-01

DOI
10.1016/j.jebo.2020.06.002

Data Availability
The data associated with this publication are available at:
http://surveys.worldbank.org/lsms/programs/integrated-surveys-agriculture-ISA/tanzania

Peer reviewed

Inferring Informal Risk-Sharing Regimes: Evidence from Rural Tanzania*

Zhimin Li†

Ethan Ligon‡

August 2020

Abstract

This paper studies informal risk-sharing regimes in a unified framework by examining households' intertemporal consumption behavior. We exploit a theoretically-consistent link between interest rates and cross-sectional consumption moments to test alternative risk-sharing models without requiring data on interest rates or assuming a restriction to eliminate the need for such data, which are often unavailable in developing economies. We specify tests that allow us to distinguish among models even with temporal dependence in income shocks. Using data from rural Tanzania we find that the consumption pattern is consistent with the self-insurance regime, and that risk aversion varies substantially across districts. Imposing a strict condition on interest rates, as often done in prior literature, misses their intertemporal heterogeneity and biases the estimation of risk aversion.

Keywords: Risk Sharing; Full Insurance; Self-Insurance; Private Information
JEL Classification: D82, D91, E21, O12

†Corresponding author. HSBC Business School, Peking University; email: zhimin.li@pku.edu.cn.
‡University of California, Berkeley.

1 Introduction

Informal insurance, such as payments from neighbors when crops fail, enables poor households in developing countries to share risk. But this insurance is far from complete: The poor still face substantial fluctuations in consumption due to idiosyncratic shocks to income (Gertler and Gruber, 2002; Banerjee and Duflo, 2007). Understanding the source of frictions that limit risk sharing is critical in the design of development policies since different interventions interact with existing informal insurance regimes differently.

Different models have been proposed to explain why risk sharing is incomplete. The first is the self-insurance model, in which households insure against income risks through borrowing and saving in formal or informal credit markets (Bewley, 1977; Hall, 1978). Another is the private information model, in which households pool income, but inter-household transfers are constrained because households do not perfectly observe others' costly actions or information (Rogerson, 1985; Kocherlakota, 2005). Two special cases of the private information model are worth noting: One involves hidden actions (Spear and Srivastava, 1987; Ligon, 1998), while the second involves hidden information (Rogerson, 1985), including hidden information about income realizations (Townsend, 1982; Kinnan, 2019). The moment restrictions we test in this paper allow both actions and information to be private.

Previous studies have exploited the idea that different risk-sharing models imply different restrictions on the evolution of households' marginal utility of expenditures (MUE). These papers adopt a common set of assumptions regarding preferences, taking them to be intertemporally separable with exponential discounting at a common discount factor; separable also in consumption and leisure; and having utility from consumption feature constant relative risk aversion (CRRA).

This paper *extends* earlier efforts by developing a unified framework to empirically distinguish among models with full insurance, self-insurance, and private information. We *generalize* previous tests in three important ways: (i) relaxing assumptions related to the temporal independence of income; (ii) imposing no parametric restrictions aside from as-

suming CRRA utility over consumption with exponential discounting; and (iii) allowing for time-varying interest rates and variation in aggregate consumption.

Most closely related to this paper is Ligon (1998), which exploits the fact that with CRRA preferences and interest rates equal to the rate of time preference, self-insurance via credit markets implies a martingale restriction on the MUE, while private information implies a similar restriction on the *reciprocal* of the MUE. Using panel data from the Indian ICRISAT villages, he finds that the private information model is consistent with the data for at least some villages, and rejects the self-insurance model for these villages. To perform these tests, however, one needs to either have data on market interest rates or impose a restriction that interest rates are constant and equal to the rate of time preference (as assumed by Ligon), which mechanically eliminates the need for such data. This is because the Euler equation characterizing consumption patterns under alternative risk-sharing models depends on interest rates. However, interest rate data are often unavailable or unreliable in developing countries. Assuming constant interest rates, on the other hand, places a priori restrictions on intertemporal consumption patterns, permitting no heterogeneity of interest rates over time or across space, which potentially biases risk aversion estimation and welfare calculations.

Another paper with similar aims is Kinnan (2019), which adopts the same set of assumptions regarding the preferences described above, and assumes also that the distribution of income is independent and identically distributed (i.i.d.) over households and time conditional on current and lagged effort. By imposing this condition Kinnan is able to add a model of limited commitment to the list of models that have testable restrictions on the evolution of households' MUEs. Using the Townsend Thai panel she rejects models of full insurance, hidden action, and limited commitment, but her tests depend critically on the assumption of i.i.d. income processes, which runs counter to evidence in the data (Meghir and Pistaferri, 2004).

This paper proposes a method that seeks to tackle limitations in the tests of Ligon (1998) and Kinnan (2019). Relative to the tests implemented by Ligon (1998), our chief

methodological contribution is that we relax the assumption that interest rates are constant, which allows for shocks to aggregate consumption in a way that Ligon implicitly rules out. Relative to the tests of Kinnan (2019), we relax the assumption that income shocks are independently distributed across time, and our tests impose less ad hoc structure. Kinnan's assumption of independence is critical to the validity of her test of the limited commitment model, so that we are unable to test that model in the framework of this paper. However, we test a model of private information that nests the other two models that she considers— hidden action (moral hazard) and hidden income—and also test a model of self-insurance, all permitting arbitrary temporal dependence in income shocks.

We apply our method to study informal insurance regimes in Tanzanian villages using a longitudinal dataset from rural household surveys. We test full risk sharing as well as alternative models of partial risk sharing such as self-insurance and private information, including models featuring hidden actions and hidden information. The empirical analysis shows that household consumption is only partially insured against income risks and that the consumption pattern is consistent with the self-insurance regime. Moreover, estimated risk aversion varies substantially across districts. Imposing an a priori restriction on interest rates, as often done in previous literature, misses their intertemporal heterogeneity and biases the estimation of risk aversion and welfare.[1]

This paper contributes to the literature on discriminating among different models of risk sharing. A variety of other papers attempt to do so using tests that do not rely primarily on restrictions on households' MUEs. Dubois, Jullien, and Magnac (2008) and Karaivanov and Townsend (2014) develop full-blown structural models with different frictions, which they estimate using maximum likelihood. Ligon and Schechter (2019) conduct experimental games in rural Paraguay to test risk-sharing models with frictions, and find that outcomes differ across villages. Other studies use related restrictions in high-income countries to test models with private information, for example, Kocherlakota and Pistaferri (2009) and Attanasio

---

[1] The degree of risk aversion is key to the estimation of welfare and vulnerability in developing countries (Ligon and Schechter, 2003; Chetty and Looney, 2006).

and Pavoni (2011), but their focus on asset pricing consequences makes them somewhat less relevant to risk sharing in rural areas of low-income countries.

The rest of the paper proceeds as follows. We first describe the theoretical models of risk sharing and their first-order necessary conditions in section 2. These equilibrium conditions form the basis of the empirical tests detailed in section 3. We then describe the data from rural Tanzania in section 4 and test different risk-sharing regimes in section 5. Section 6 concludes.

2 Theoretical Framework

We begin with a basic framework of alternative models of risk sharing. As these models are reasonably standard, we will focus on their equilibrium conditions and relate them to the empirical strategy in this paper. The notation is adapted from Kocherlakota and Pistaferri (2009) and Attanasio and Pavoni (2011).

2.1 Common Environment

The following description of the economic environment is common across all of the models we consider. This environment is a rural village consisting of N households (indexed by i) that exist for $T + 1$ periods (indexed by t; T may be infinite).

Preferences

Each household has identical von Neumann-Morgenstern preferences that are separable in time and consumption-effort decisions. Each household i chooses a stream of consumption and effort by maximizing

$$\mathbb{E}_0 \sum_{t=0}^{T} \beta^t [u(c_{it}) - v(e_{it})],$$

where $\mathbb{E}_0$ indicates the expectation conditional on the information set as of time $t = 0$. The function u is the flow utility function of consumption, v is the flow disutility function of

5

effort, and c_{it} and e_{it} denote, respectively, consumption and effort of household i at time t.

We assume that households exhibit constant relative risk aversion (CRRA) preferences over consumption:

$$u(c) = \begin{cases} \frac{c^{1-\eta}-1}{1-\eta} & \text{if } \eta \neq 1, \\ \log c & \text{if } \eta = 1, \end{cases}$$

where $\eta > 0$ (which implies that u is strictly increasing and strictly concave). These assumptions are standard, but it may be worth pointing out their limitations. In particular, relative risk aversion (η) must be constant, ruling out non-homothetic preferences, and permissible heterogeneity in preferences is very limited.[2]

Shocks

As in Kocherlakota and Pistaferri (2009) there are two sources of shocks. The first is an aggregate or public shock $z_t \in Z$, Z finite, realized at time t with the history of these shocks up to time t written as z^t. We allow for arbitrary dependence in these aggregate shocks.

The second is an idiosyncratic shock realized at time t by each household i, $\theta_{it} \in \Theta$, Θ finite; the history of such shocks for each household i up to time t is written as $\theta_i^t \in \Theta^t$. Arbitrary temporal dependence is again permitted, but dependence in realizations of θ_{it} across households is governed by the realization of the history z^t; that is, draws of θ_{it} are independent across households conditional on z^t and the history of idiosyncratic shocks. Initial shocks (z_0, θ_{i0}) are given.

Household i's output at time t, denoted as y_{it}, depends not only on contemporaneous effort but also on its history of idiosyncratic shocks θ_i^t and the history of aggregate shocks z^t:

$$y_{it} = F(z^t, \theta_i^t, e_{it}),$$

where F is the production function. In our context idiosyncratic shocks θ_{it} might be illness

---

[2] It would not affect our tests at all were we to allow for household-specific preference shocks provided that the process governing these shocks was a martingale difference, but the additional complexity does not seem worth the trouble.

or crop failure, for example, while aggregate shocks z_t might be related to weather. Different assumptions on the observability of idiosyncratic shocks, effort, and output give rise to different risk-sharing regimes.

Financial Intermediary

We assume that for each village there is some intermediary or intermediaries who after a history of aggregate shocks z^t can borrow or save at an interest rate $r_t \equiv r(z^t)$. This intermediary can also make contingent transfers to households inside the village. Within the class of models we consider below there is no loss of generality in assuming that contingent transfers can *only* be made by the intermediary.

There are different ways to interpret this intermediary: For example, Thomas and Worrall (1990) and Ligon (1998) interpret it as a risk-neutral principal, but it would also be natural to interpret it as a bank or risk-averse moneylender.

2.2 The Full Insurance Model

We next consider a sequence of models with differing frictions, and derive restrictions from these models that allow us to distinguish among them. We start with the benchmark "full insurance" model having complete Arrow-Debreu markets, with risk-sharing properties explored by Townsend (1994), among others.

Assume for this model that z^t, θ_i^t, e_i^t (i.e., histories of shocks and effort) are publicly observable. In equilibrium the economy achieves full risk sharing, with the hallmark result that ratios of marginal utilities are equated across households in all states of the world, or

$$\left(\frac{c_{it}(z^t, \theta_i^t)}{c_{jt}(z^t, \theta_j^t)} \right)^{-\eta} = \frac{\lambda_j}{\lambda_i},$$

where λ_i is the Pareto weight for household i prescribed by the social planner, which is a

constant. Note that without loss of generality we can normalize $\sum_{i=1}^{N} \lambda_i^{1/\eta} = 1$.[3] Then the household's consumption is proportional to aggregate consumption in the village economy:

$$c_{it}(z^t, \theta_i^t) = \lambda_i^{1/\eta} C_t,$$

where $C_t \equiv \sum_j c_{jt}(z^t, \theta_j^t)$ denotes aggregate consumption. Taking the logarithmic transformation yields

$$\log c_{it}(z^t, \theta_i^t) = \log \lambda_i^{1/\eta} + \log C_t. \tag{1}$$

This condition leads to a strong prediction: Under full insurance, each household's consumption is independent of idiosyncratic income and varies over time only with aggregate resources.

2.3   The Self-Insurance (Permanent Income) Model

The self-insurance model is distinguished from the full insurance model by an inability for households to make state-contingent transfers. This gives us a situation in which households cannot engage in mutual insurance with other households but are able to self-insure against shocks by borrowing and lending at a common interest rate via the intermediary (Hall, 1978; Bewley, 1977). A variety of frictions (including private information) might limit state-contingent transfers, and we do not distinguish among these different frictions.

In contrast to Ligon (1998), we allow the interest rate faced by households to vary, depending on the history of aggregate shocks, so that the common interest rate faced by households at time t is $r_t(z^t)$. In contrast to Kinnan (2019), we allow income shocks to have unrestricted temporal dependence—in Kinnan's case all shocks are assumed to be independent, or "transitory" in the language of the permanent income hypothesis. Assuming independence of income shocks seems particularly problematical when one seeks to test this

---

[3] As will be clear shortly, $\lambda_i^{1/\eta}$ can be interpreted as household i's consumption share in terms of aggregate resources.

model, since the distinction between permanent and transitory income shocks was central to Friedman's original conception of the permanent income hypothesis (Friedman, 1957). Meghir and Pistaferri (2004) establish the importance of allowing for temporal dependence in evaluating the permanent income hypothesis in the US, while Paxson (1992) finds strong evidence that transitory and non-transitory income shocks have different effects on savings behavior, as predicted by the permanent income model.

The household chooses consumption and effort bundles by maximizing life-time utility subject to an intertemporal budget constraint

$$A_{it+1} = (1 + r(z^t))(A_{it} + F(z^t, \theta_i^t, e_{it}) - c_{it}(z^t, \theta_i^t))$$

for all (θ_i^t, z^t), where A_{it} is the stock of the asset held by household i at time t, and that household i begins with finite holdings $A_{i0} \in \theta_{i0}$.

In this environment households' consumption will satisfy the usual Euler condition

$$(c_{it}(z^t, \theta_i^t))^{-\eta} = \mathbb{E}_t[\beta(1 + r(z^t))(c_{it+1}(z^{t+1}, \theta_i^{t+1}))^{-\eta}],$$

or, upon rearranging terms,

$$\mathbb{E}_t\left[\left(\frac{c_{it+1}(z^{t+1}, \theta_i^{t+1})}{c_{it}(z^t, \theta_i^t)}\right)^{-\eta} - \frac{1}{\beta(1 + r(z^t))}\right] = 0, \tag{2}$$

where the expectation operator is over the information set at time t.

2.4   The Private Information Model

Suppose now that the households we observe possess private information in the form of hidden actions or hidden information. This is a critical difference from the full insurance model. Relative to both the full- and self-insurance models, another difference is that households do not have (unobserved) access to credit markets.

Though they cannot access credit markets, households can write long-term contracts with the (unobserved) intermediary, so that the intertemporal marginal rate of substitution of the intermediary is pinned down by interest rates.

For households with private information, histories of idiosyncratic shocks (θ_i^t) and effort (e_i^t) may not be observed, but aggregate shocks (z^t) are public. Let the household choose among reporting strategies, denoted by $\sigma : Z^t \times \Theta^t \to \Theta^t$ (Golosov, Kocherlakota, and Tsyvinski, 2003). The household's choice of effort affects the distribution of output (realizations of which may themselves be private information), and reports the household makes about its hidden information may affect the transfers it receives from the intermediary.

In the private information regime, incentive compability typically rules out full insurance, while a result known as the "inverse Euler equation" holds (Rogerson, 1985; Kocherlakota, 2005). With CRRA utility from consumption, this takes the form (see the Appendix):

$$(c_{it}(z^t, \theta_i^t))^\eta = \mathbb{E}_t \left[\frac{(c_{it+1}(z^{t+1}, \theta_i^{t+1}))^\eta}{\beta(1 + r(z^t))} \right],$$

or, upon rearranging terms,

$$\mathbb{E}_t \left[\left(\frac{c_{it+1}(z^{t+1}, \theta_i^{t+1})}{c_{it}(z^t, \theta_i^t)} \right)^\eta - \beta(1 + r(z^t)) \right] = 0. \tag{3}$$

While we think of y_{it} as being some valuable output realized at t, for our present purpose of characterizing household consumption when there is private information output does not need to be treated as a separate stochastic object, as it is completely determined by $(z^t, \theta_i^t, e_{it})$. In the case in which output is "observed" we can think of y_{it} being an element of the public history z^t; in the case in which output is "hidden" it is in the information set generated by (z^t, θ_i^t). Note also that the fact that the probability of some realization of y_{it} depends on the entire history of shocks up to t means that we treat both the case in which agents choose an effort *before* observing shocks that determine output and the case in which they choose the effort *after* observing the shocks as special cases.

2.5 Inferring Risk-Sharing Regimes Without Interest Rate Data

The moment conditions implied by the Euler conditions in models of self-insurance (equation (2)) and private information (equation (3)) provide a unified framework for distinguishing among risk-sharing regimes.

Before describing the empirical method, we need to resolve an important issue: Interest rate data in rural economies are often unavailable. Fortunately, following Kocherlakota and Pistaferri (2009) one can use cross-sectional consumption moments to infer information about interest rates. Consider the self-insurance model. We sum across households on both sides of the Euler equation preceding equation (2),

$$\sum_j (c_{jt}(z^t, \theta_j^t))^{-\eta} = \beta(1 + r(z^t)) \mathbb{E}_t \left[\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{-\eta} \right]$$

and rearrange terms to obtain

$$\frac{1}{\beta(1 + r(z^t))} = \mathbb{E}_t \left[\frac{\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{-\eta}}{\sum_j (c_{jt}(z^t, \theta_j^t))^{-\eta}} \right]. \tag{4}$$

For the private information regime we obtain an analogous condition:

$$\beta(1 + r(z^t)) = \mathbb{E}_t \left[\frac{\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{\eta}}{\sum_j (c_{jt}(z^t, \theta_j^t))^{\eta}} \right]. \tag{5}$$

It is worth noting that these are *equilibrium* conditions that determine the relationship between interest rates and cross-sectional moments of the conditional consumption distribution. If the village is small and has access to outside credit markets, then prevailing interest rates will determine how the distribution of consumption within the village changes over time, consistent with either equation (4) or equation (5). Conversely, if the village is closed, with the intermediary providing credit or transfers *only* to people within the village, then the right-hand side of these equations will instead endogenously determine market-clearing interest rates.

We substitute the two equations above into (2) and (3) to obtain key moment conditions for the self-insurance model,

$$\mathbb{E}_t \left[\left(\frac{c_{it+1}(z^{t+1}, \theta_i^{t+1})}{c_{it}(z^t, \theta_i^t)} \right)^{-\eta} - \frac{\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{-\eta}}{\sum_j (c_{jt}(z^t, \theta_j^t))^{-\eta}} \right] = 0, \tag{6}$$

and for the private information model,

$$\mathbb{E}_t \left[\left(\frac{c_{it+1}(z^{t+1}, \theta_i^{t+1})}{c_{it}(z^t, \theta_i^t)} \right)^{\eta} - \frac{\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{\eta}}{\sum_j (c_{jt}(z^t, \theta_j^t))^{\eta}} \right] = 0. \tag{7}$$

Equations (6) and (7) offer a unified framework for testing partial risk-sharing regimes without requiring data on interest rates. Despite their similar appearance, the two conditions are quite different from each other. One does not imply the other because the inverse function cannot pass through the expectation operator.

These expressions are useful for at least two reasons. First, they provide a theoretically consistent way to link cross-sectional consumption moments to aggregate shocks facing households. Kocherlakota and Pistaferri (2009) and Ligon (2010) also exploit this advantage, although they analyze vastly different contexts from that in this paper.[4] Second, the restrictions above allow one to estimate risk aversion parameters without having to observe interest rates. This is an important methodological difference between the tests in this paper and those in Ligon (1998). Ligon (1998) imposes a restriction on the risk-free return $(1 + r_t)$ such that it is equal to the reciprocal of the discount factor in every period; that is, $\beta(1 + r_t) = 1$ or $r_t = 1/\beta - 1$ for all t, where β is the discount factor. This condition assumes away the need to collect interest rate data by construction, a decision driven by a lack of information on interest rates in Indian villages. The downside, however, is that it rules out aggregate shocks implicitly. In contrast, the approach we develop in this paper does not impose a priori restrictions on the relationship between the discount factor and interest rates, thus allowing

---

[4] Kocherlakota and Pistaferri (2009) attempt to explain the risk premium puzzle using data on asset prices and repeated cross-sectional consumption, and Ligon (2010) measures risks by looking at cross-sectional consumption inequality.

for aggregate shocks to interest rates. These advantages make the approach well suited for studying rural economies.

2.6 Relaxing the Restriction on Interest Rates

A key improvement of the proposed method over prior literature is the relaxation of the restriction imposed on interest rates, and this allows the models to match intertemporal consumption data better and has implications for the inference of risk aversion and the risk-sharing regime.[5] The model-consistent restrictions on interest rates across risk-sharing regimes, shown in equations (4) and (5), produce implied interest rates facing households. With an estimated risk aversion coefficient, we can employ cross-sectional moments of consumption to infer the model-consistent interest rates (for a given discount factor) and examine how they change over time. If the implied value of $\beta(1 + r_t)$ is far from unity, then imposing the restriction that it equates unity may give wrong inference on the risk-sharing regime and/or the degree of risk aversion.

Consider the moment condition for the self-insurance model. The left-hand side of equation (6) depends on the distribution of $\left(\frac{c_{it+1}(z^{t+1}, \theta_i^{t+1})}{c_{it}(z^t, \theta_i^t)} \right)^{-\eta} - \frac{\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{-\eta}}{\sum_j (c_{jt}(z^t, \theta_j^t))^{-\eta}}$, which in turn depends on the joint distribution of both within-household consumption growth (the first term) and cross-household consumption moments (the second term). In the strict version of the model tested by Ligon (1998) interest rates are taken to be equal to the rate of time preference, consistent with his assumption that intermediaries are risk neutral and thus bear all aggregate risk. In this case the moment condition becomes

$$\mathbb{E}_t \left[\left(\frac{c_{it+1}(z^{t+1}, \theta_i^{t+1})}{c_{it}(z^t, \theta_i^t)} \right)^{-\eta} - 1 \right] = 0,$$

which depends only on the distribution of within-household consumption growth, and which

---

[5] Strictly speaking, the restriction we relax in our methodology is $\beta_t(1 + r_t) = 1$ or $r_t = 1/\beta_t - 1$, which would also allow the discount factor to vary over time. Following most of the literature on risk sharing, we restrict attention to the conventional case with a constant discount factor (i.e., exponential discounting).

may differ substantially from equation (6) if $\sum_j (c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{-\eta} \neq \sum_j (c_{jt}(z^t, \theta_j^t))^{-\eta}$. This is the case, for example, if either aggregate consumption or its distribution changes over time. Whether the strict and proposed versions of the moment conditions differ from each other and yield different inference on the risk-sharing regime or the coefficient of risk aversion is an empirical question.

To provide some intuition on how allowing for temporal variation in interest rates may matter for the inference of risk-sharing regimes or risk aversion, consider a simple case of optimal consumption in the self-insurance model without uncertainty, but with time-varying interest rates. Having time-varying interest rates is dual to having aggregate variation in intertemporal marginal rates of substitution, which we would expect if we simply had, for example, deterministic changes in aggregate income. The intertemporal marginal rate of substitution is equated to the risk-free rate of return, i.e., $\frac{1}{\beta} \left(\frac{c_{it+1}}{c_{it}} \right)^{\eta} = 1 + r_t$. A strict restriction that $\beta(1 + r_t) = 1$ for all t implies that consumption is constant over time, but with aggregate variation in income this implies that credit markets are not clearing. Relaxing that restriction allows the consumption path to tilt, so consumption need not be constant over time. This enables the model to match intertemporal consumption patterns in the data better. We can also see why allowing for aggregate shocks to interest rates is crucial for the inference of risk aversion and the risk-sharing regime. If interest rates vary over time, there exist correlated shocks that have not been entirely smoothed away, which leads to variation in household consumption over time even when risk aversion is high (or its reciprocal, the elasticity of intertemporal substitution, is low). However, the more restrictive model prohibits interest rates from varying, and hence concludes incorrectly that the high intertemporal variation in consumption must be attributed to low risk aversion, causing bias (toward zero) in the estimation of the risk aversion parameter.

Heterogeneity in Interest Rates and Risk Aversion

Our proposed approach can account for potential heterogeneity in interest rates and risk aversion across geographical locations (e.g., villages), though not across households within locations. Suppose interest rates and/or the discount factor vary across time (indexed by t) and regions (indexed by v). By examining risk sharing across regions, the proposed method relaxes the restriction imposed by Ligon (1998), $\beta_v(1 + r_{vt}) = 1$, into $\beta_v(1 + r_{vt}) = k_{vt}$, where k_{vt} may differ from unity and can vary across time and regions; it also permits spatial heterogeneity in risk aversion. Specifically, the model-consistent restriction on interest rates in the self-insurance regime is

$$\frac{1}{\beta_v(1 + r_{vt})} = \mathbb{E}_t \left[\frac{\sum_{j \in v}(c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{-\eta_v}}{\sum_{j \in v}(c_{jt}(z^t, \theta_j^t))^{-\eta_v}} \right], \tag{8}$$

and in the private information regime is

$$\beta_v(1 + r_{vt}) = \mathbb{E}_t \left[\frac{\sum_{j \in v}(c_{jt+1}(z^{t+1}, \theta_j^{t+1}))^{\eta_v}}{\sum_{j \in v}(c_{jt}(z^t, \theta_j^t))^{\eta_v}} \right]. \tag{9}$$

Equations (8) and (9) allow not only interest rates to vary over time and space (because cross-sectional moments in consumption can differ along both dimensions), but also risk aversion to vary across regions. In the empirical analysis, we test the models using both the pooled sample (assuming homogeneous interest rates and risk aversion) and subsamples across districts (allowing for possible heterogeneity in interest rates and risk aversion parameters). We also examine how test results change if the restriction $\beta_v(1 + r_{vt}) = 1$ is imposed instead.

A caveat is that the method cannot accommodate interest rate heterogeneity across individuals within the same geographic location. For example, interest rates faced by individuals in the same village may differ due to different household asset levels, or people may display different preferences as a result.[6] This type of heterogeneity may matter for risk-sharing

---

[6] See Mazzocco and Saini (2012) for a test of heterogeneous preferences across individuals.

arrangements, but the method is not readily applicable to such situations.

3 Empirical Implementation

Our estimation problem is based on the Euler-type conditions in equations (6) and (7). Let household i's information set at time t be denoted by I_{it}. By the properties of conditional expectations, elements of I_{it} are orthogonal to the Euler forecast error, so one can transform conditional moments into unconditional ones. This transformation allows us to apply the generalized method of moments (GMM) estimator to estimate risk aversion coefficients (Hansen and Singleton, 1982). A primary advantage of this approach is that there is no need to solve a fully-specified dynamic model, which would involve parametric assumptions and complicated computation. Specifically, let

$$\mathbb{E}[\zeta_{it+1}(\varphi) \cdot x_{it}] = 0, \tag{10}$$

where $\zeta_{it+1}(\varphi) \equiv \left(\frac{c_{it+1}(z^{t+1},\theta_i^{t+1})}{c_{it}(z^t,\theta_i^t)} \right)^{-\varphi} - \frac{\sum_j (c_{jt+1}(z^{t+1},\theta_j^{t+1}))^{-\varphi}}{\sum_j (c_{jt}(z^t,\theta_j^t))^{-\varphi}}$ and $x_{it} \in I_{it}$. Notice that $\zeta_{it+1}(\varphi)$ expresses the terms inside the brackets of equations (6) and (7) for informal risk-sharing models in a unified manner: $\varphi = \eta$ if the economy is operating according to the self-insurance regime, and $\varphi = -\eta$ if it is the private information regime.

The parameter φ is the key object to estimate for inference of risk-sharing regimes. Following Ligon (1998), we infer the insurance regime from the sign of the parameter estimate $\widehat{\varphi}$. If $\widehat{\varphi} > 0$, then one rejects the private information model; if $\widehat{\varphi} < 0$, then one rejects the self-insurance model.[7] Table 1 summarizes the model testing strategy.

---

[7] If all households are risk neutral, i.e., $\varphi = 0$, or if there is in fact full insurance then this case is degenerate because the Euler moment condition always holds. This degeneracy is not economically meaningful. As explained in section 5, this issue motivates the choice of a continuously-updating GMM estimator.

Robustness to Measurement Error in the Data

The proposed method is robust to a relatively broad class of measurement error processes in the data—either in consumption or instrumental variables—due to employment of cross-sectional consumption moments to infer interest rates.

Consider first measurement error in consumption. Suppose that the observed consumption contains multiplicative measurement error, $\widetilde{c}_{it}(z^t, \theta_i^t) = c_{it}^*(z^t, \theta_i^t) \exp(\nu_{it})$, where $\widetilde{c}_{it}$ and c_{it}^* are, respectively, measured and true consumption levels, and ν_{it} is the measurement error. One may deal with measurement error by imposing either a parametric or a non-parametric structure for the error process (Ventura, 1994; Hong and Tamer, 2003; Chioda, 2004; Alan, Attanasio, and Browning, 2009; Kocherlakota and Pistaferri, 2009). Suppose ν_{it} is independent from idiosyncratic shocks and independent across households, and takes a stationary structure such that $\kappa \equiv \mathbb{E}\left[\frac{\exp(-\varphi\nu_{it+1})}{\exp(-\varphi\nu_{it})}\right] < \infty$ for all t.[8]

We show that the moment condition using observed consumption will still be valid. Notice that

$$
\begin{aligned}
\mathbb{E}_t\left[\left(\frac{\widetilde{c}_{it+1}(z^{t+1}, \theta_i^{t+1})}{\widetilde{c}_{it}(z^t, \theta_i^t)}\right)^{-\varphi}\right] &= \mathbb{E}_t\left[\left(\frac{c_{it+1}^*(z^{t+1}, \theta_i^{t+1})}{c_{it}^*(z^t, \theta_i^t)}\right)^{-\varphi} \frac{\exp(-\varphi\nu_{it+1})}{\exp(-\varphi\nu_{it})}\right] \\
&= \mathbb{E}_t\left[\left(\frac{c_{it+1}^*(z^{t+1}, \theta_i^{t+1})}{c_{it}^*(z^t, \theta_i^t)}\right)^{-\varphi}\right] \mathbb{E}\left[\frac{\exp(-\varphi\nu_{it+1})}{\exp(-\varphi\nu_{it})}\right] \\
&= \mathbb{E}_t\left[\left(\frac{c_{it+1}^*(z^{t+1}, \theta_i^{t+1})}{c_{it}^*(z^t, \theta_i^t)}\right)^{-\varphi}\right] \kappa,
\end{aligned}
$$

where the independence assumption of the measurement error process is invoked. Similarly,

---

[8] If $\nu_{it} \sim \mathcal{N}(0, \sigma_\nu^2)$ i.i.d., using a property of the log normal distribution (i.e., if $m \sim \mathcal{N}(\mu_m, \sigma_m^2)$, $\mathbb{E}[\exp(m)] = \exp\left(\mu_m + \sigma_m^2/2\right)$), we can show that $\kappa = \exp(\varphi^2\sigma_\nu^2)$. A similar result would hold were we to assume a more flexible parametric distribution of measurement error such as Laplace (Hong and Tamer, 2003; Chioda, 2004).

the cross-sectional component

$$\mathbb{E}_t\left[\frac{\sum_j(\widetilde{c}_{jt+1}(z^{t+1},\theta_j^{t+1}))^{-\varphi}}{\sum_j(\widetilde{c}_{jt}(z^t,\theta_j^t))^{-\varphi}}\right] = \mathbb{E}_t\left[\frac{\sum_j(c_{jt+1}^*(z^{t+1},\theta_j^{t+1}))^{-\varphi}}{\sum_j(c_{jt}^*(z^t,\theta_j^t))^{-\varphi}}\right]\mathbb{E}\left[\frac{\exp(-\varphi\nu_{jt+1})}{\exp(-\varphi\nu_{jt})}\right]$$

$$= \mathbb{E}_t\left[\frac{\sum_j(c_{jt+1}^*(z^{t+1},\theta_j^{t+1}))^{-\varphi}}{\sum_j(c_{jt}^*(z^t,\theta_j^t))^{-\varphi}}\right]\kappa.$$

Therefore, the left-hand side of the key moment condition for model testing is

$$\mathbb{E}_t\left[\left(\frac{\widetilde{c}_{it+1}(z^{t+1},\theta_i^{t+1})}{\widetilde{c}_{it}(z^t,\theta_i^t)}\right)^{-\varphi} - \frac{\sum_j(\widetilde{c}_{jt+1}(z^{t+1},\theta_j^{t+1}))^{-\varphi}}{\sum_j(\widetilde{c}_{jt}(z^t,\theta_j^t))^{-\varphi}}\right]$$

$$= \mathbb{E}_t\left[\left(\frac{c_{it+1}^*(z^{t+1},\theta_i^{t+1})}{c_{it}^*(z^t,\theta_i^t)}\right)^{-\varphi} - \frac{\sum_j(c_{jt+1}^*(z^{t+1},\theta_j^{t+1}))^{-\varphi}}{\sum_j(c_{jt}^*(z^t,\theta_j^t))^{-\varphi}}\right]\kappa.$$

so setting the empirical moment using observed, error-ridden consumption to zero is equivalent to setting an analogous condition using true consumption. These assumptions about the nature of the measurement error process are not completely innocuous, of course, but there is no need to impose any restrictions on the functional form of the measurement error, its magnitude, or its autocorrelation structure, beyond assuming that a particular moment is finite (Kocherlakota and Pistaferri, 2009).

We can similarly show that the test method is robust to classical measurement error in variables in the instrument set as well. Suppose the error takes a multiplicative form: $\widetilde{x}_{it} = x_{it}^*\exp(\epsilon_{it})$, where $\widetilde{x}_{it}$ and x_{it}^* are, respectively, measured and true levels of an instrumental variable, and ϵ_{it} is the measurement error. Assume that the error process is independent from idiosyncratic shocks, independent across households, and stationary with a finite moment such that $\gamma \equiv \mathbb{E}[\exp(\epsilon_{it})] < \infty$.[9] When the measured variable $\widetilde{x}_{it}$ is used as an instrument,

---

[9] If $\epsilon_{it} \sim \mathcal{N}(0,\sigma_\epsilon^2)$ i.i.d., then $\gamma = \exp\left(\sigma_\epsilon^2/2\right)$.

the left-hand side of the moment condition of the test restriction is

$$\mathbb{E}_t \left[\left(\left(\frac{c_{it+1}}{c_{it}} \right)^{-\varphi} - \frac{\sum_j c_{jt+1}^{-\varphi}}{\sum_j c_{jt}^{-\varphi}} \right) \widetilde{x}_{it} \right] = \mathbb{E}_t \left[\left(\left(\frac{c_{it+1}}{c_{it}} \right)^{-\varphi} - \frac{\sum_j c_{jt+1}^{-\varphi}}{\sum_j c_{jt}^{-\varphi}} \right) x_{it}^* \right] \mathbb{E}[\exp(\epsilon_{it})]$$

$$= \mathbb{E}_t \left[\left(\left(\frac{c_{it+1}}{c_{it}} \right)^{-\varphi} - \frac{\sum_j c_{jt+1}^{-\varphi}}{\sum_j c_{jt}^{-\varphi}} \right) x_{it}^* \right] \gamma,$$

so setting the empirical moment that uses observed data for the instrumental variable to zero is equivalent to setting an analogous condition using the true measure.

4 Data

We use a household-level panel dataset from the Kagera Health and Development Survey (KHDS) conducted in the Kagera region in Tanzania between 1991 and 1994, and apply the proposed method discussed above to study risk sharing in Tanzanian villages. This region (with a population of about two million in 2004) is primarily rural with the north mainly producing bananas and coffee and the south planting rain-fed crops such as maize, sorghum, and cotton. The data come from a longitudinal survey jointly conducted by the Population and Human Resources Department and the Africa Technical Department of the World Bank, which interviewed about 800 households from nearly 50 communities in all five districts of Kagera: Bukoba Urban, Bukoba Rural, Muleba, Biharamulo, Ngara, and Karagwe. Although the KHDS questionnaires were adapted from those in the World Banks Living Standards Measurement Study, the dataset was unique due to its panel structure (for detailed information on the sampling design, see Ainsworth (2004)).

Because the main objective of the survey was to estimate the economic impact of adult mortality and morbidity on surviving household members, it contained detailed questions about household consumption, income, transfers, and demographic information such as the gender and schooling of the household head, offering researchers the opportunity to study how households cope with risks. This paper uses the three waves of data for 1992, 1993,

and 1994 because the recall periods of purchased and home-produced food items in the consumption and expenditure data were 12 months in the first wave in 1991, rather than six months as in 1992–1994. This sample selection is made to ensure sampling consistency.[10]

As households in the region predominantly engage in agricultural production, rainfall is a key determinant of income. To measure this important source of aggregate risk, we have obtained monthly rainfall data from Tanzania Meteorological Agency for years 1980–2004 and matched them with the survey data based on the nearest weather station according to direct-line estimates from the Geographic Information System data on village centers and rainfall stations. The region has two rainy seasons, a long season usually between March and May and a short season between October and December. We construct a rainfall measure using average monthly z-score deviations of rainfall during two most recent rainy seasons preceding the interviews.

This dataset is well suited for studying consumption fluctuations and informal risk sharing. The income of households in Kagera is subject to substantial variation arising from uncertainty in rainfall. We find that a one standard deviation increase in average rainfall during rainy seasons is associated with an average increase of 38.3% in household income. The high degree of income variation implies potentially large gains from risk-sharing arrangements between households.

Table 2 shows the summary statistics of key variables. Consumption, transfers, and income are expressed in 2004 Tanzanian Shillings. To adjust for family size and the ages of household members, consumption is expressed in per adult equivalent terms. Notice that there are considerable fluctuations in income and consumption both within and across households. This suggests that insurance against consumption risk is far from complete and that the welfare loss from incomplete risk sharing may be substantial. In the next section we formally test the model of full insurance and alternative models of constrained risk sharing.

---

[10] KHDS also collects data in two more rounds, in 2004 and 2010, after the first round of 1991–1994, but because the consumption modules were modified between rounds, the later periods were excluded from the analysis. Moreover, family compositions may have changed drastically over time, posing challenges in using the later rounds of surveys.

5 Tests of Risk-Sharing Models

5.1 Testing the Full Insurance Model

We start by using a standard regression specification to test the full insurance model (Townsend, 1994; Suri, 2005). This is essentially just a rewriting of equation (1) from the model, estimating the terms $\log \lambda_i^{1/\eta}$ as a set of household fixed effects; estimating the terms $\log C_t$ as a set of time effects; including an error process; and finally adding the log of contemporaneous income to test the exclusion restriction:

$$\log c_{it} = \alpha \log y_{it} + \delta_i + \phi_t + \varepsilon_{it}, \tag{11}$$

where c_{it} denotes household consumption, y_{it} household income, δ_i household fixed effects, ϕ_t time effects, and ε_{it} a disturbance term (e.g., due to measurement error or unobserved preference heterogeneity). The time effects control for economy-wide shocks to aggregate resources. Full insurance implies that $\alpha = 0$.

The logic behind this specification follows from equation (1): Once household fixed effects (the first term on the right-hand side) and aggregate resources (the second term) are controlled for, household consumption should not depend on idiosyncratic shocks such as the income of the household.

As an alternative specification, we add community-year fixed effects to capture the aggregate resource in each community (indexed by v) by estimating

$$\log c_{it} = \alpha \log y_{it} + \delta_i + \phi_{vt} + \varepsilon_{it}, \tag{12}$$

where ϕ_{vt} denotes community-year fixed effects. To check robustness, we also include time-varying demographic variables to capture households' idiosyncratic preference shifters that may affect consumption. Full insurance implies again that $\alpha = 0$.

Note the conceptual difference between the underlying assumptions about risk-sharing

arrangements embodied in the two testing equations above. In equation (11) the risk-sharing network consists of all communities. In equation (12), by contrast, risk sharing happens within each community, but there need not be any inter-community insurance. A comparison of the coefficients on household income across the two specifications may suggest whether inter-community insurance is quantitatively important in Kagera.

Table 3 shows the test results. The first two columns do not control for community-year fixed effects and thus presume the risk-sharing network to consist of households across all communities in the region, while the third and fourth columns control for community- or district-year fixed effects, respectively, which assume the mutual insurance network to be at the community or district level. A few patterns emerge from observing the regression results across specifications. First, full insurance has been soundly rejected. The elasticity of household consumption to household income is between 0.40 to 0.42—with all estimates statistically different from zero and fairly stable across specifications. In other words, a 10% rise in household income is correlated with about 4% increase in consumption. Second, comparing across specification with and without demographic controls, we find that time-varying household characteristics (e.g., the age of the household head or whether the household head is female) do not appear to affect the elasticity estimates. Third, based on the similar magnitude in the coefficient estimates across specifications (11) and (12), whether the risk-sharing entity is specified to be a community, a district, or the whole region does not seem to matter, either. The rejection of the full insurance model is not surprising given the vast amount of evidence in the literature (Townsend, 1994; Gertler and Gruber, 2002; Kinnan, 2019).

5.2   Testing Self-Insurance and Private Information Models

The strong evidence against full insurance leads us to explore mechanisms constraining risk sharing among rural households in Kagera. We do so by exploiting the Euler equations implied by alternative models with the empirical strategy outlined in section 3. The sample

analog of the moment condition in equation (10) is

$$\frac{1}{\bar{T}}\frac{1}{N}\sum_t\sum_i[g_{it}(\varphi)] = 0,$$

with $g_{it}(\varphi) \equiv \zeta_{it+1}(\varphi) \cdot x_{it}$, where N is the number of households, $\bar{T}$ one less than the number of time periods for which we have data, and $x_{it} \in I_{it}$ denotes variables pertaining to household i that are in the time t information set.

We apply a continuously-updating GMM estimator (Hansen, Heaton, and Yaron, 1996; Imbens, Spady, and Johnson, 1998) and estimate φ as[11]

$$\widehat{\varphi} = \arg\min_{\widetilde{\varphi}} \left(\frac{1}{\bar{T}}\frac{1}{N}\sum_t\sum_i g_{it}(\widetilde{\varphi})\right)' W(\widetilde{\varphi}) \left(\frac{1}{\bar{T}}\frac{1}{N}\sum_t\sum_i g_{it}(\widetilde{\varphi})\right), \qquad (13)$$

where $W(\widetilde{\varphi}) = \left(\frac{1}{\bar{T}}\frac{1}{N}\sum_t\sum_i g_{it}(\widetilde{\varphi})g_{it}(\widetilde{\varphi})'\right)^{-1}$ is the optimal weighting matrix. For the estimation procedure, we need to consider an important issue: the choice of variables to enter the information set. In principle, all lagged variables are in households' information set at time t and can be used to form the moment condition. To achieve the highest power of model testing, however, it is desirable to choose variables that inform households' consumption decisions. It is not advisable to choose variables that are "random," because even at false parameter values the correlation between the Euler forecast error and random variables will be close to zero—leading to weak identification (Stock, Wright, and Yogo, 2002).[12] In the empirical analysis we use lagged consumption, household size, and rainfall as instrumental variables to form households' information set.

Table 4 shows the estimation results. We test the models both for the whole sample and by district. The first row pools households across all districts; the second to the seventh

---

[11] The continuously-updating GMM estimator appears to have better properties than the traditional two-step or iterated GMM estimator, although they are asymptotically equivalent (see Hansen et al. (1996) and Imbens et al. (1998) for discussions). This estimator is robust to the issue of degeneracy of the moment condition at $\varphi = 0$.

[12] The continuously-updating GMM estimator is partially robust to weak identification. Another estimator that is also partially robust is the Generalized Empirical Likelihood estimator (Stock et al., 2002).

rows use samples from individual districts. The bottom row lists variables that enter the information set (which include consumption, household size, and rainfall, all lagged by one period). The first to third columns show the test results when we add the instruments one by one.

The magnitude of risk aversion parameters varies considerably, with a range in absolute value between 0.19 and 1.04 across districts. It is reassuring, however, that estimates for both the pooled sample and individual districts are stable across information sets, which range from 0.21 to 0.22 for the pooled sample, from 0.72 to 0.78 in Karagwe, from 0.19 to 0.20 in Bukoba Rural, from 0.64 to 0.65 in Muleba, from 0.61 to 0.75 in Biharamulo, from 1.00 to 1.04 in Ngara, and from 0.51 to 0.53 in Bukoba Urban. The range is moderately wider in Biharamulo probably due to its smaller sample size.

While the point estimate (in absolute value) appears low, at around 0.22, for the pooled sample, testing the model by district uncovers vast heterogeneity. This suggests that it is crucial to allow households living in different districts to face potentially different interest rates. The estimated range of risk aversion coefficients across districts—from 0.20 to 1.04 (based on results in column (3), where the information vector includes consumption, household size, and rainfall)—is broadly consistent with those estimated in the literature for developing countries. For example, using a pooled sample of three villages in South Indian, Ligon (1998) estimates risk aversion coefficients to range from 0.30 to 1.58.

Most important in the test results in table 4 are the signs of the point estimates. The estimated risk aversion coefficients are positive and statistically significant (with p-values below 0.01) across all specifications and samples, which suggests that the private information model is rejected and that households across all districts appear to engage in risk-sharing contracts consistent with the self-insurance regime.

To illustrate the importance of allowing for aggregate shocks, we also test the models by imposing a strict restriction on interest rates such that $r_{vt} = 1/\beta_v - 1$ for all v and t, as assumed in Ligon (1998), and compare the results with those obtained using the proposed

method in this study, which relaxes that assumption. Table 5 shows the comparison in test results both for the pooled sample and by district. For the pooled sample and for all districts, the signs of the point estimates are consistent across the two methods, so the inference of the risk-sharing regime is the same in our data. Both strategies find evidence consistent with the self-insurance model.

However, the magnitude of risk aversion parameters differs considerably across the two methods. The strict version estimates a lower risk aversion for all but one district, with a potential bias (in magnitude) ranging from 1% for Ngara to 49% for Karagwe. The bias is even worse for the pooled sample (69%): The strict version estimates the risk aversion coefficient to be merely 0.07, while the estimate is 0.22 once the interest rate restriction is relaxed. The reason for the different risk aversion estimates is that our tests permit aggregate shocks to interest rates, whereas the strict version does not. As we will show, implied interest rates do indeed vary significantly. These aggregate shocks lead to substantial variation in consumption over time despite households' high risk aversion; however, the more restrictive model cannot capture these shocks and misattributes the large intertemporal variation in consumption to low risk aversion. In line with this intuition, table 5 does indeed show that the restricted version tends to artificially suppress risk aversion in order to be able to explain observed consumption fluctuations. The vast differences in risk aversion estimates across methods suggest that welfare calculations may be biased were one to apply the strict method to analyze development policies.

Given that the tests reveal that the data are consistent with the self-insurance regime, we can compute the implied, theoretically-consistent interest rates according to the Euler equation and show their temporal variation. The calculation is based on the sample analogue of equation (8); that is,

$$r_{vt}^{\text{implied}} = \frac{1}{\beta_v} \frac{\sum_{j \in v} c_{jt}^{-\widehat{\eta}_v}}{\sum_{j \in v} c_{jt+1}^{-\widehat{\eta}_v}} - 1.$$

Assuming a typical value for the discount factor $\beta_v = 0.95$, we calculate the implied interest

rates for the pooled sample and by district, and the results are shown in table 6.[13] The implied interest rates vary vastly across time and districts. For the pooled sample, the implied rate increased slightly from 4.59% to 4.72% during 1992–1994, but this belies the heterogeneity across districts. During that period, the rates in Bukoba Rural and Bukoba Urban rose modestly from 4.56% to 5.15% and from 2.55% to 2.99%, respectively, while those in Muleba decreased from 3.67% to 2.90%. Other districts experienced more dramatic changes: The implied rates in Karagwe dropped from 20.33% to 4.06%, and those in Biharamulo from 2.61% to -1.01%; those in Ngara climbed from -0.38% to 5.13%. Clearly, the strict version of the tests, which imposes a constant interest rate, would have missed such information about interest rate heterogeneity across time and regions (the strict model imposes a homogeneous rate of $r_{vt} = 1/\beta_v - 1 = 5.26\% \ \forall v, t$).

5.3 Discussion

The evidence of self-insurance from model tests in this paper suggests that inter-household transfers play a limited role in mitigating idiosyncratic income risks in rural Tanzania. Another test employs an idea from Rosenzweig (1988): For households to share risk, one would expect income shocks to be negatively correlated with inter-household transfers; that is, a household that experiences a drop in income should receive a positive transfer from the risk-sharing network.

Table 7 tests this hypothesis. The regression specifications are similar to those in table 3, except that the dependent variable now indicates whether a household receives incoming or net transfers. In columns (1) and (2), the dependent variable is an indicator for whether a household receives positive incoming transfers (1 for yes, 0 for no); in (3) and (4), it is an indicator for whether a household receives positive net transfers (net of the outgoing amount). The first and third columns presume the risk-sharing network to consist of households belonging to the same community, while the second and fourth columns assume the network

---

[13] While we have assumed a common discount factor across regions in the calculation of interest rates, the tests we propose and implement do not impose that assumption.

to be at the district level. Across all specifications, the coefficient on log household income is economically and statistically insignificant, which suggests that inter-household transfers do not appear to respond to household income shocks and thus do not serve to effectively reduce those risks.

A note of caution is that throughout this paper we have restricted attention to distinguishing between risk-sharing mechanisms that deliver similar forms of the Euler equation, and tested them against one another in a unified framework. One might be concerned, however, that there may be other models we have not considered that can also generate the observed patterns in Tanzanian households' consumption behavior. Alternative explanations might include limited enforcement (Ligon, Thomas, and Worrall, 2002; Laczó, 2015; Ábrahám and Laczó, 2018), endogenous network formation (Genicot and Ray, 2003), transaction costs (Jack and Suri, 2014), interpersonal variation in time preferences (Dean and Sautmann, 2020), or a combination of multiple mechanisms. For example, Broer, Kapička, and Klein (2017) suggests that limited enforcement combined with ex post private information may have similar implications to self-insurance. Unfortunately, these models do not imply a form of the Euler equation that permits their testing using a unified approach. The researcher typically has to assume a specific friction a priori and estimate/simulate a fully specified parametric model, or needs to impose restrictive i.i.d. income processes to enable tests of more frictions (as in Kinnan, 2019).

In light of alternative mechanisms, it is useful to compare our findings to other researchers' in rural Tanzania or a similar setting. Using a dataset from a household survey in Nyakatoke, a small Haya community in the Bukoba Rural District of the Kagera region of Tanzania, De Weerdt and Dercon (2006) examine how insurance networks help insure idiosyncratic shocks such as health shocks and show evidence consistent with partial risk sharing for non-food consumption. To investigate the formation of insurance networks, De Weerdt and Fafchamps (2011) find no signs that transfers in the event of health shocks are reciprocal or that binding self-enforcement constraints are at play in risk-sharing arrangements in Nyaka-

toke. Observing that transfers systematically flow from the rich to poor households, they reject limited commitment models with wealth asymmetry (Fafchamps, 1999) and present evidence of risk sharing between kin based on altruism or cultural norms. Theoretically, social norms can interact with individual savings behavior in mutual insurance across households (Wahhaj, 2010). In contexts other than Tanzania, research has shown limited ability of informal networks to insure rural households against income risks using data from Côte dIvoire (Deaton, 1997), the Philippines (Fafchamps and Lund, 2003), and Thailand (Townsend, 1995). These findings are consistent with evidence in this paper that risk sharing between households in the same community is constrained.

6 Conclusion

Disentangling sources of partial risk sharing remains an essential area of research in development economics. In this paper we propose a method to distinguish among risk-sharing regimes that relaxes a strict restriction on interest rates without requiring data on them—a useful feature for studying informal insurance in developing economies—and develop tests of models including full insurance, self-insurance, private information (including hidden income), all allowing for arbitrary forms of temporal dependence in income processes. Applying the method to data from rural Tanzania, we find that households' intertemporal consumption behavior is most consistent with the self-insurance regime. Inter-household transfers appear to play a limited part in sharing risk between rural households in our data. Policies that facilitate these transfers may therefore help insure against idiosyncratic shocks for the poor and enhance welfare.

References

Ábrahám, Á. and S. Laczó (2018). Efficient risk sharing with limited commitment and storage. *Review of Economic Studies 85*(3), 1389–1424.

Ainsworth, M. (2004). User's guide to the Kagera health and development survey datasets. Development Research Group. The World Bank.

Alan, S., O. Attanasio, and M. Browning (2009). Estimating Euler equations with noisy data: Two exact GMM estimators. *Journal of Applied Econometrics 24*(2), 309–324.

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

Banerjee, A. V. and E. Duflo (2007). The economic lives of the poor. *Journal of Economic Perspectives 21*(1), 141–168.

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory 16*(2), 252–292.

Broer, T., M. Kapička, and P. Klein (2017). Consumption risk sharing with private information and limited enforcement. *Review of Economic Dynamics 23*, 170–190.

Chetty, R. and A. Looney (2006). Consumption smoothing and the welfare consequences of social insurance in developing economies. *Journal of Public Economics 90*(12), 2351–2356.

Chioda, L. (2004). Estimating Euler equations with measurement error: A nonparametric approach. Working Paper, Department of Economics, University of California at Berkeley.

De Weerdt, J. and S. Dercon (2006). Risk-sharing networks and insurance against illness. *Journal of Development Economics 81*(2), 337–356.

De Weerdt, J. and M. Fafchamps (2011). Social identity and the formation of health insurance networks. *Journal of Development Studies 47*(8), 1152–1177.

Dean, M. and A. Sautmann (2020). Credit constraints and the measurement of time preferences. *Review of Economics and Statistics*, forthcoming.

Deaton, A. (1997). *The analysis of household surveys: a microeconometric approach to development policy*. The World Bank.

Diamond, P. and J. Mirrlees (1978). A model of social insurance with variable retirement. *Journal of Public Economics 10*(3), 295–336.

Dubois, P., B. Jullien, and T. Magnac (2008). Formal and informal risk sharing in LDCs: Theory and empirical evidence. *Econometrica 76*(4), 679–725.

Fafchamps, M. (1999). Risk sharing and quasi-credit. *Journal of International Trade & Economic Development 8*(3), 257–278.

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

Friedman, M. (1957). *A Theory of the Consumption Function.* Princeton: Princeton University Press.

Genicot, G. and D. Ray (2003). Group formation in risk-sharing arrangements. *The Review of Economic Studies 70*(1), 87–113.

Gertler, P. and J. Gruber (2002). Insuring consumption against illness. *American Economic Review 92*(1), 51–70.

Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003). Optimal indirect and capital taxation. *The Review of Economic Studies 70*(3), 569–587.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy 86*(6), 971–87.

Hansen, L. P., J. Heaton, and A. Yaron (1996). Finite-sample properties of some alternative GMM estimators. *Journal of Business & Economic Statistics 14*(3), 262–280.

Hansen, L. P. and K. J. Singleton (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica 50*(5), 1269–1286.

Hong, H. and E. Tamer (2003). A simple estimator for nonlinear error in variable models. *Journal of Econometrics 117*(1), 1–19.

Imbens, G., R. H. Spady, and P. Johnson (1998). Information theoretic approaches to inference in moment condition models. *Econometrica 66*(2), 333–358.

Jack, W. and T. Suri (2014). Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review 104*(1), 183–223.

Karaivanov, A. and R. M. Townsend (2014). Dynamic financial constraints: Distinguishing mechanism design from exogenously incomplete regimes. *Econometrica 82*(3), 887–959.

Kinnan, C. (2019). Distinguishing barriers to insurance in Thai villages. Working Paper, Tufts University, https://econpapers.repec.org/paper/tuftuftec/0831.htm.

Kocherlakota, N. and L. Pistaferri (2009). Asset pricing implications of Pareto optimality with private information. *Journal of Political Economy 117*(3), 555–590.

Kocherlakota, N. R. (2005). Zero expected wealth taxes: A Mirrlees approach to dynamic optimal taxation. *Econometrica 73*(5), 1587–1621.

Laczó, S. (2015). Risk sharing with limited commitment and preference heterogeneity: Structural estimation and testing. *Journal of the European Economic Association 13*(2), 265–292.

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

Ligon, E. (2010). Measuring risk by looking at changes in inequality: Vulnerability in Ecuador. CUDARE Working Paper No. 1095.

Ligon, E. and L. Schechter (2003). Measuring vulnerability. *The Economic Journal 113*(486), C95–C102.

Ligon, E. and L. Schechter (2019). Structural experimentation to distinguish between models of risk sharing with frictions in rural Paraguay. *Economic Development and Cultural Change*, forthcoming.

Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *The Review of Economic Studies 69*(1), 209–244.

Mazzocco, M. and S. Saini (2012). Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review 102*(1), 428–68.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in Thailand. *The American Economic Review*, 15–33.

Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica 53*(1), 69–76.

Rosenzweig, M. R. (1988). Risk, implicit contracts and the family in rural areas of low-income countries. *The Economic Journal 98*(393), 1148–1170.

Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *The Review of Economic Studies 54*(4), 599–617.

Stock, J. H., J. H. Wright, and M. Yogo (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics 20*(4).

Suri, T. (2005). Spillovers in village consumption: Testing the extent of partial insurance. Working Paper, Yale University.

Thomas, J. and T. Worrall (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. (1982). Optimal multiperiod contracts and the gain from enduring relationships under private information. *Journal of Political Economy 90*(6), 1166–1186.

Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica 62*(3), 539–591.

Townsend, R. M. (1995). Financial systems in Northern Thai villages. *The Quarterly Journal of Economics 110*(4), 1011–1046.

Ventura, E. (1994). A note on measurement error and Euler equations: An alternative to log-linear approximations. *Economics Letters 45*(3), 305–308.

Wahhaj, Z. (2010). Social norms and individual savings in the context of informal insurance. *Journal of Economic Behavior & Organization 76*(3), 511–530.

Table 1: Model Testing Strategy

	Self-Insurance	Private Information
Moment condition	$\mathbb{E}[\zeta_{it+1}(\varphi) \cdot x_{it}] = 0$	$\mathbb{E}[\zeta_{it+1}(\varphi) \cdot x_{it}] = 0$
Sign of $\widehat{\varphi}$	> 0	< 0

Notes: $\zeta_{it+1}(\varphi) \equiv \left(\frac{c_{it+1}}{c_{it}}\right)^{-\varphi} - \frac{\sum_j c_{jt+1}^{-\varphi}}{\sum_j c_{jt}^{-\varphi}}$ and $x_{it} \in I_{it}$ are variables pertaining to household i's information set as of time t. In the empirical analysis, these variables include lagged consumption, household size, and rainfall.

Table 2: Summary Statistics

Variable	Mean	Std. Dev.
Household income	474,349	798,173
Income fluctuation over time	237,785	580,430
Consumption	135,434	136,590
Consumption fluctuation over time	50,373	86,831
Incoming transfer	53,625	538,228
HH head is female	0.277	0.448
HH head received schooling	0.816	0.388
HH head age	50	17
Rainfall	187.32	46.52
N	2,256	

Notes: Consumption is measured in per adult equivalent terms. The calculation is the same as that in Ligon and Schechter (2003): It assigns adult males a weight of 1 and adult females a weight of 0.9 (adults are at the ages of sixteen or older); children of ages 0 to 4 receive a weight of 0.32, ages 5 to 9 a weight of 0.52, and ages 10 to 15 a weight of 0.67. Consumption, income, and transfers are expressed in 2004 Tanzanian Shillings (TZS). 1 USD = 1,029 TZS as of January 1, 2004. Income and consumption fluctuation variables refer to the standard deviation of each household over different rounds of interviews. The standard deviations are pooled (i.e., across both time periods and households), except for income and consumption fluctuation over time, which is only across households (as those variables are income and consumption standard deviations already taken across time periods). The rainfall variable (measured in millimeters) is time series (1980-2004) averages of the two rainy seasons (March-May and October-December). It is matched to villages based on the nearest weather station.

Table 3: Testing the Full Insurance Model

Variable	(1)	(2)	(3)	(4)
Log household income	0.4177*	0.4184*	0.4028*	0.4166*
	(0.0225)	(0.0228)	(0.0251)	(0.0223)
Demographic controls	No	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Community-year FE	No	No	Yes	No
District-year FE	No	No	No	Yes
N	2,253	2,253	2,253	2,253

Notes: This table shows the regression results of testing the full insurance model. The dependent variable is log household consumption per capita (per adult equivalent). All variables are in per capita terms when appropriate. Demographic controls include the age of the household head and two indicators for whether the head is female and has received any formal schooling. The balanced panel has a total 2,256 observations, three of which have negative household income and are dropped from the regression analysis. Standard errors (in parenthesis) are clustered at the community level. $*p < 0.01$.

Table 4: Test Results by District

District	(1)	(2)	(3)	N
All	0.2092*	0.2187*	0.2192*	2,256
	(0.0266)	(0.0264)	(0.0261)	
Karagwe	0.7200*	0.7755*	0.7760*	237
	(0.1159)	(0.1198)	(0.1177)	
Bukoba Rural	0.1942*	0.1965*	0.1965*	768
	(0.0439)	(0.0429)	(0.0429)	
Muleba	0.6391*	0.6490*	0.6507*	372
	(0.0874)	(0.0876)	(0.0877)	
Biharamulo	0.6108*	0.6329*	0.7540*	156
	(0.1316)	(0.1239)	(0.1353)	
Ngara	1.0041*	1.0213*	1.0374*	258
	(0.0843)	(0.0841)	(0.0804)	
Bukoba Urban	0.5088*	0.5270*	0.5216*	465
	(0.0547)	(0.0557)	(0.0553)	
Instruments	c	c, hhsize	c, hhsize, rain	

Notes: This table shows test results for the pooled sample and by district. The values correspond to estimated CRRA coefficients. The variables that enter the information set are lagged by one period (they include consumption, household size, and rainfall). All information sets include a vector of constants. Standard errors are shown in parenthesis. $*p < 0.01$.

Table 5: Test Results by Method

District	Proposed Method	Strict Version	Potential Bias
All	0.2192*	0.0677*	-69%
	(0.0261)	(0.0096)	
Karagwe	0.7760*	1.1584*	49%
	(0.1177)	(0.1519)	
Bukoba Rural	0.1965*	0.1360*	-31%
	(0.0429)	(0.0328)	
Muleba	0.6507*	0.5745*	-12%
	(0.0877)	(0.0777)	
Biharamulo	0.7540*	0.5544*	-26%
	(0.1353)	(0.1000)	
Ngara	1.0374*	1.0311*	-1%
	(0.0804)	(0.1207)	
Bukoba Urban	0.5216*	0.3879*	-26%
	(0.0553)	(0.0569)	

Notes: This table shows test results based on two methods. The strict version assumes $r_t = 1/\beta - 1 \ \forall t$. The values correspond to estimated CRRA coefficients. The information vector includes the full list of variables (these include consumption, household size, and rainfall). Standard errors are shown in parenthesis. $*p < 0.01$.

Table 6: Implied Interest Rates

District	Proposed Method		Strict Version
	Year=1992	Year=1993	All Years
All	4.59%	4.72%	
Karagwe	20.33%	4.06%	
Bukoba Rural	4.56%	5.15%	
Muleba	3.67%	2.90%	5.26%
Biharamulo	2.61%	-1.01%	
Ngara	-0.38%	5.13%	
Bukoba Urban	2.55%	2.99%	

Notes: This table shows implied interest results based on two methods: the proposed method and a strict version. The strict version assumes $r_t = 1/\beta - 1 \; \forall t$. The information vector includes the full list of variables (these include consumption, household size, and rainfall). The calculation assumes $\beta = 0.95$.

Table 7: The Limited Role of Inter-household Transfers in Mitigating Income Risk

Variable	(1)	(2)	(3)	(4)
Log household income	-0.0151	-0.0115	-0.0249	-0.0233
	(0.0109)	(0.0110)	(0.0212)	(0.0193)
Demographic controls	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Community-year FE	Yes	No	Yes	No
District-year FE	No	Yes	No	Yes
N	2,253	2,253	2,253	2,253

Notes: This table shows the association between a change in income and whether a household receives transfers. In columns (1) and (2), the dependent variable is an indicator for whether a household receives positive incoming transfers (1 for yes, 0 for no); in (3) and (4), it is an indicator for whether a household receives positive net transfers (net of the outgoing amount). Demographic controls include the age of the household head and two indicators for whether the head is female and has received any formal schooling. The balanced panel has a total 2,256 observations, three of which have negative household income and are dropped from the regression analysis. Standard errors (in parenthesis) are clustered at the community level.

Appendix: The Inverse Euler Equation

Variants of the "inverse Euler equation" were established in a variety of special models quite early in the development of dynamic principal-agent models. For example, Diamond and Mirrlees (1978) present a version of the result in considering a model of public insurance in retirement. But early results relied on a first-order characterization of optimal arrangements, and thus were no more general than the first-order conditions used to establish them.

What is now the standard variational argument, originally due to Rogerson (1985), turns out to have much more general applicability. We develop this here for the sake of completeness, but without much claim to novelty.

We consider the case facing an intermediary and a single agent. Given independence across agents conditional on observables this problem generalizes to the case of multiple agents without complication. Note also that in the case of many agents this conditional independence assumption may have little bite. For example, if there is an unobserved time varying parameter that affects the mean of the distribution of output for all agents, then realizations of this mean become part of the observable history, and independence across agents is only required conditional on this mean (and other aspects of the observable history).

Consider an intermediary who seeks to minimize the net cost of inducing an agent to choose an effort profile e^*, while keeping future utility promises to the agent. The intermediary can influence the agent's utility by assigning consumption or effort in every period. Effort is not transferable, but consumption is. Interest rates may depend on the public history z^t, but not the private history, and after history z^t are denoted $r(z^t)$. The intermediary minimizes the expected present value of net transfers made to the agent, with a profile $\tau^* = \{\tau_t(z^t, \theta^t)\}_{t=0}^T$; thus the expected present value can be written as $\mathbb{E}_0 \sum_{t=0}^T \prod_{s=1}^t (1 + r(z^{s-1}))^{-1} \tau_t$.

Proposition 1. *Let $c^* = \{c_t(z^t, \theta^t)\}_{t=0}^T$ be an incentive-compatible consumption profile that minimizes the intermediary's cost of implementing the effort profile e^*. Then after any*

history (z^t, θ^t) the profiles c^* and e^* will satisfy

$$\mathbb{E}_t \left[\frac{u'(c_t(z^t, \theta^t))}{u'(c_{t+1}(z^{t+1}, \theta^{t+1}))} \right] = \beta(1 + r(z^t)).$$

Proof. The consumption profile c^* implements the effort profile e^* by assumption. Then given some history (z^t, θ^t) the intermediary assigns consumption $c_t(z^t, \theta^t)$, while for any subsequent realization of (z_{t+1}, θ_{t+1}) the intermediary plans to assign $c_{t+1}(z^{t+1}, \theta^{t+1})$.

Consider two adjacent periods $(t, t+1)$. The agent's utility from consumption at t after the history (z^t, θ^t) is denoted by $\omega_t(z^t, \theta^t) = u(c_t(z^t, \theta^t))$. Conditional on this history subsequent realizations just depend on (z_{t+1}, θ_{t+1}), delivering momentary utility from consumption $\omega_{t+1}(z^{t+1}, \theta^{t+1})$.

The intermediary can feasibly reschedule the agent's utility across these two periods by increasing utility at t by some ϵ, at a present cost to the intermediary of $u^{-1}(\omega_t(z^t, \theta^t) + \epsilon)$, but then decreasing utility by ϵ/β the next period, for every realization of (z_{t+1}, θ_{t+1}), which reduces discounted expected costs in the subsequent period:

$$\frac{1}{1 + r(z^t)} \mathbb{E}_t \left[u^{-1}(\omega_{t+1}(z^{t+1}, \theta^{t+1}) - \epsilon/\beta) \right],$$

where the expectation operator is conditional on the information set as of time t (this may include histories of shocks and the agent's efforts). Since utility of the agent is intertemporally separable, reschedulings of this sort will have no effect on incentive compatibility: Any such constraints are linear in utilities, and utility ϵ at t is a perfect substitute for utility ϵ/β at $t+1$. Then for any efficient consumption profile c^*, ϵ must be zero, and the derivative of the sum of the discounted costs across the two periods with respect to ϵ must also be zero, so that

$$u^{-1\prime}(\omega_t(z^t, \theta^t)) = \frac{1}{\beta(1 + r(z^t))} \mathbb{E}_t \left[u^{-1\prime}(\omega_{t+1}(z^{t+1}, \theta^{t+1})) \right],$$

for all (z^t, θ^t). Applying the inverse function theorem and rearranging the equation yields

the result. $\square$