

# Intertemporal Risk Pooling in Village Economies\*

Marcel Fafchamps<sup>†</sup> and Aditya Shrinivas<sup>‡</sup>

June 2024

## Abstract

We propose a new method to test for efficient risk pooling that allows for intertemporal smoothing, non-homothetic consumption, and heterogeneous risk and time preferences. The method is composed of three steps. The first one allows for precautionary savings by the aggregate risk pooling group. The second utilizes the inverse Engel curve to estimate good-specific tests for efficient risk pooling. In the third step, we obtain consistent estimates of households' risk and time preferences using a full risk sharing model, and incorporate heterogeneous preferences in testing for risk pooling. We apply this method to recent panel data from Indian villages to generate a number of new insights. We find that food expenditures are better protected from aggregate shocks than non-food consumption, after accounting for non-homotheticity. Village-level consumption tracks aggregate village cash-in-hand, suggesting some form of coordinated precautionary savings. We also find a large unexplained gap between the variation in measured consumption expenditures and cash-in-hand at the aggregate village level. Contrary to earlier findings, risk pooling in Indian villages no longer appears to take place more at the sub-caste level than at the village level.

---

\*We thank Pierre Dubois, Cynthia Kinnan, and Pau Milan for their comments, as well as the participants to the WEFIDEV webinar of November 2021 and to the IZA Workshop on Gender and Family Economics of June 2022. We are grateful to ICRISAT for sharing the data with us from their website <http://vdsa.icrisat.ac.in>. The only input into this research is the time of the two researchers involved. No outside funding was secured for this project. We have no conflict of interest to report.

<sup>†</sup>Email: [fafchamp@stanford.edu](mailto:fafchamp@stanford.edu). Stanford University, 616 Jane Stanford Way, Stanford CA 94305.

<sup>‡</sup>Email: [aditya.shrinivas@iimb.ac.in](mailto:aditya.shrinivas@iimb.ac.in). Indian Institute of Management Bangalore, Bannerghatta road, Bengaluru, India 560076.

# 1 Introduction

People in village economies face income shocks due to drought, flood, unemployment, illness, and crop or business failure. Households that are uninsured against these shocks experience consumption fluctuations detrimental to their welfare (e.g., Gertler and Gruber, 2002; De Weerd and Dercon, 2006). Protection from such income shocks depends on the availability and effectiveness of institutions that distribute and share risk. Asset accumulation *ex ante* can help smooth consumption through precautionary saving (e.g., Zeldes, 1989a; Deaton, 1991). Risk can also be pooled *ex post* through various informal or formal agreements (e.g., Udry, 1990; Fafchamps and Lund, 2003; Dercon et al., 2006). Risk pooling is best at addressing idiosyncratic shortfalls in income that only affect some households. But it offers little or no protection against aggregate shocks that affect the whole village. To self-protect against such shocks, some form of precautionary saving is required – either at the aggregate level (e.g., cereal bank; provident fund) or at the individual level (Fafchamps et al., 1998). It follows that effective protection of village economies against income shocks requires a combination of risk sharing and precautionary saving.

In spite of this, the literature on consumption smoothing has tended to focus on either of these two mechanisms in isolation. Early studies of consumption smoothing across individuals have focused on risk pooling within periods (e.g., Mace, 1991; Altonji et al., 1992; Townsend, 1994; Schulhofer-Wohl, 2011; Mazzocco and Saini, 2012; Chiappori et al., 2014) and those that have examined risk pooling across periods have ignored or assumed away assets (e.g., Kocherlakota, 1996; Ligon, 1998; Ligon et al., 2002; Kinnan, 2022). Similarly, studies of intertemporal risk coping via precautionary savings have typically assumed away contemporaneous risk pooling across households (e.g., Zeldes, 1989a; Deaton, 1991; Rosenzweig and Wolpin, 1993; Lim and Townsend, 1998; Fafchamps et al., 1998; Kazianga and Udry, 2006). While past empirical tests of risk pooling are nested in our model and, as such, are unbiased, studies of precautionary saving that ignore risk pooling across individuals may yield misleading results. This is because, in the presence of risk pooling, individually held assets can be used to smooth the consumption of others. Hence precautionary saving can only be meaningfully studied at the level of the risk pooling group aggregate.

In this paper, we propose a novel approach to testing for efficient risk pooling that combines both strands of the literature into a single theoretically-based framework. Its purpose is to test whether village economies optimally pool risk without unrealistically assuming away wealth accumulation, as tests of optimal risk pooling have done to date. Perfect risk pooling at the village level is calculated by solving a social planner optimization that maximizes the weighted sum of discounted expected utilities of all households subject to an intertemporal budget constraint. In this framework, the decision of the social planner can be divided into two steps: in each period, the social planner first optimally chooses aggregate village consumption as in a precautionary saving model (e.g., Zeldes, 1989a; Deaton, 1991); in a second step, this aggregate consumption is optimally divided among households (e.g., Townsend, 1994). This framework

allows for individual asset accumulation but, in an efficient equilibrium, all assets are de facto held in trust for the entire risk pooling group and all village savings and income are pooled to smooth consumption over time and within periods. In such equilibrium, the optimal division of wealth across villagers only depends on whether returns to individual assets are decreasing, increasing, or constant.<sup>1</sup> This is because individually held assets can be used to smooth the consumption of others. This framework bears some similarity with how lending by saving and credit associations to their members can help non-members and enhance risk pooling (e.g., Attanasio et al. (2023) ). New important insights are gained from this approach.

As second main innovation, we extend the above model to multiple consumption categories and use it to develop a new risk sharing test that allows for non-homothetic consumption preferences. There has been a recent revival of interest in Engel curves (e.g., Dunbar et al., 2013; Atkin et al., 2020; Ligon, 2020; Almås et al., 2018; Escanciano et al., 2021). Drawing inspiration from this literature, we propose a new way of using partial consumption data to test for efficient risk pooling when income elasticities are non-unitary. This approach allows us to test whether certain components of consumption are better insured than others (e.g., Mace, 1991) – as could arise, for instance, in the presence of paternalism or imperfect altruism.<sup>2</sup> All this is done while allowing for heterogeneity in risk and time preferences among households, as has successfully been done in the recent literature (e.g., Schulhofer-Wohl, 2011; Chiappori et al., 2014; Mazzocco and Saini, 2018). This approach is based on the intuition, dating back to Wilson (1968), that efficient risk pooling allocates more risk to more risk-tolerant households, implying that a household whose consumption strongly co-moves with village consumption must be relatively more risk tolerant. Finally, we replicate the analysis assuming aggregation of risk either at the village-level or the sub-caste level.

We start by replicating results obtained with other datasets. In contrast to the several papers testing risk pooling in Indian villages that use the ICRISAT data collected in six villages between 1975 and 1985 (e.g., Townsend, 1994; Ravallion and Chaudhuri, 1997; Mazzocco and Saini, 2018), we rely on a new wave of ICRISAT panel data from a much larger sample of 30 Indian villages. For total expenditures, we find evidence of perfect risk pooling at the village level - the correlation with earned income is not statistically different from zero. This finding is different to earlier studies that strongly reject full risk pooling using other datasets. We find that household consumption expenditures move nearly one-for-one with the village average. Even though year-to-year variation in household-level cash-in-hand has a statistically significant coefficient, the magnitude of the correlation with household expenditures remains very small. From this we conclude that the study villages engage in a considerable pooling of idiosyncratic within-year risk.<sup>3</sup> Our results also confirm previous findings indicating that there is substantial

---

<sup>1</sup>If returns are convex, all precautionary assets should be held by a single actor (e.g., an insurance corporation or a large money lender). If they are concave, wealth should be equally distributed so as to maximize aggregate returns. If the saving technology has linear returns (e.g., financial savings), the distribution of assets does not matter.

<sup>2</sup>The long-lasting US Food Stamp program is a well-known example where one dimension of consumption – i.e., food – is better insured than others – e.g., utilities, heating, transport (e.g., Hastings and Shapiro, 2018).

<sup>3</sup>We are not saying, however, that risk pooling is the result of the explicit sharing of risk among villagers. Indeed, risk pooling tests do not identify the mechanism by which risk is being pooled. This reality, which is

and significant heterogeneity in estimated risk preferences across households – and we show that this conclusion extends to time preferences. As previous authors have argued theoretically, we find the failure to correct for this heterogeneity biases coefficient estimates in any risk pooling test. In our data, however, this bias is empirically negligible and it does not affect the qualitative conclusions from the analysis.

Our first novel contribution is to show that food and non-food expenditures do not vary with average village consumption in a way consistent with a within-year utility-maximizing allocation of total expenditures: based on estimated non-linear Euler curves, households increase their food consumption less – and their non-food consumption more – than what the increase in average village expenditures would predict. We also do not find that non-food expenditures respond more to household income and wealth than food expenditures, ruling out the idea that there is less risk pooling in non-food consumption than in food consumption. Instead, the results suggest that, in a good year for the village as a whole, people increase their non-food expenditures more than proportionally, thereby sheltering their average food consumption from aggregate shocks. In the context of risk pooling, this is consistent with the same kind of altruistic paternalism that often lead high-income countries to favor welfare interventions that favor the consumption of specific goods (e.g., food stamps; subsidized health care and education; affordable housing initiatives).

Our second novel contribution relates to consumption smoothing across years. Here, we find a small but significant response of village consumption to aggregate village cash-in-hand. On its own, this finding could be consistent with a high level of consumption smoothing achieved through precautionary saving. Indeed, individuals with high wealth are predicted to consume only a small fraction of temporary variation in cash-in-hand (e.g. Zeldes (1989b); Deaton (1991)). But we do not find that, as a rule, villages hold large stocks of liquid wealth. Furthermore, the small and insignificant coefficient on village earned income indicates perfect asset integration: villages optimally draw on their liquid assets to smooth aggregate income shocks. These findings are (consistent?) with efficient precautionary saving at the village level. We also find that household consumption does vary significantly (albeit weakly) with their cash-in-hand, which contradicts the asset pooling prediction from the efficient intertemporal risk pooling model. Results further indicate that, after correcting for non-homothetic consumption preferences using Engel curves, village food consumption responds less than anticipated to variations in village cash-in-hand while non-food expenditures respond significantly to cash-in-hand.

Taken all together, these findings suggest that while our study villages manage to achieve a considerable smoothing of consumption in response to idiosyncratic within-year income shocks, they do not appear to use aggregate village savings as a buffer against aggregate income fluctuations across years. At the end of the paper we discuss whether the villages' success at dealing with the first problem – protection against within-year idiosyncratic shocks – may explain their difficulty at using liquid wealth as a protection against across-year collective shocks.

---

shared by all risk pooling test papers in the literature, is the reason why, throughout the paper, we have refrained from using the expression 'risk sharing' since it implicitly suggests some deliberate sharing intent.

We largely replicate these findings at the sub-caste (*Jati*) level, instead of the village level. We find that, contrary to the papers using earlier Indian data (e.g., Townsend, 1994; Mazzocco and Saini, 2012; Shrinivas and Fafchamps, 2018), risk pooling within sub-castes in each village is less strong – i.e., less responsive to aggregate consumption and more responsive to individual income and liquid assets – than pooling across all households in the village. Similar conclusions are reached for precautionary saving: aggregate sensitivity to cash-in-hand is very small – probably smaller than it should be given the relatively low level of liquid wealth in the data: in a poor population such as the one we study, we would expect a stronger dependence of consumption to cash-in-hand in a precautionary savings optimum. The sub-caste evidence also suggests that the excess sensitivity to income found in the linear (CARA) model may be due to misspecification; a log model provides a better fit for the theory.

This paper makes several contributions to the literature. Previous studies of consumption smoothing have either focused solely on inter-household risk pooling or inter-temporal precautionary savings. We show how the two empirical approaches can be combined within a single framework, and we go well beyond existing studies in the scope of our findings. Our results suggest that consumption smoothing is achieved more from within-period risk pooling, as compared to precautionary savings. In particular, villages achieve very little protection from year-to-year shocks via precautionary savings: group consumption hardly responds to aggregate liquid assets or cash-in-hand. Yet aggregate village consumption also responds little to aggregate village income we find a large variable gap between village income and village consumption. Furthermore, we find no evidence that this gap is bridged by net flows of funds to the village in loans or gifts and transfers. Taken together, these findings suggest that average village consumption is largely insulated from collective income shocks. But the process by which this is achieved is unclear, apart from the fact that it is not due to precautionary savings.

Second, we integrate non-homothetic Engel curves into our testing strategy. If households have homothetic consumption preferences, perfect risk pooling implies that all consumption categories respond equally to shocks: as the total expenditures of the household increase or contract, all consumption categories should rise and fall in the same proportion (e.g., Mace, 1991). When preferences are not homothetic, the consumption of goods with an income elasticity larger than one should respond more than proportionally to a rise in total expenditures, while the opposite holds for goods with low income elasticity. By estimating Engel curves on cross-section data, we are able to predict by how much expenditures of different expenditure categories should vary with total expenditures to remain on the household's indifference curve. If some expenditure categories contract more than predicted while others contract less, it means that smoothed consumption deviates from what would be optimal. Such deviations are common in social protection programs of high income countries and they have been associated with altruistic paternalism (reference). We document a similar pattern in a low-income informal village setting where risk pooling depends on people's goodwill and is thus likely to be affected by social norms about redistribution (e.g., Fafchamps, 1992). This could arise, for instance, because consumption categories are not all equally observable. As a result, individuals may claim a negative shock but spend the insurance payout on luxuries. To deter such behavior,

recipients of an insurance payouts may be forced to spend it on necessities such as food. We do indeed find evidence that food consumption is over-insured in the sense that it co-moves less with village consumption than is predicted by Engel curves, while non-food consumption co-moves more. This suggests that if households were able to spend their consumption budget freely, they would opt to spend less on food and more on non-food in bad years.

Third, this paper adds to the new strand of literature on risk sharing with heterogeneous preferences (e.g., Schulhofer-Wohl, 2011; Mazzocco and Saini, 2012; Chiappori et al., 2014; Dubois, 2001). This literature has brought to light the existence of an omitted variable bias in standard tests of risk pooling under homogeneous preferences, and has shown that this bias drives the income coefficient upwards, leading to spurious rejections of full insurance. Different parametric and semi-parametric tests have been proposed to account for heterogeneous risk preferences. We follow in the footsteps of this literature by building on the testing approach proposed by Chiappori et al. (2014). We start by estimating relative risk preferences between households under the assumption of perfect risk sharing and then use the resulting estimates to correct the test of full risk pooling.<sup>4</sup>Our approach, however, differs from Chiappori et al. (2014) in important ways: it generalizes the approach by simultaneously estimating relative welfare weights and time preferences between households under perfect risk pooling; and it is easier to implement since it relies on a simple linear regression.<sup>5</sup>

Lastly, we make an empirical contribution to the literature on efficiency in risk sharing. Earlier results have emphasized that full risk sharing is not achieved in village economies. Using the new wave of ICRISAT's panel data from 2010-2015, our results are qualitatively consistent with previous estimates from the 1975-85 ICRISAT data: household income is statistically significant. We nonetheless find that the marginal propensity to consume (MPC) out of household's own income is a very low 0.02, which is substantially lower than the MPC of 0.14 estimated by Townsend (1994) using the old data. In absolute terms, Rosenzweig and Binswanger (1993) find that a 100 rupee decline in income reduces food consumption by 7 rupees in the old sample. In the new sample, an equivalent decline in earned income reduces food consumption by half of that: 3.4 rupees. This suggests that consumption smoothing has improved in these villages over time. In addition, our findings corroborate recent papers that find substantial heterogeneity in

---

<sup>4</sup>Since we test efficient risk pooling, it is natural that preference heterogeneity be estimated under that assumption as well, so as to ensure internal consistency of the test. This being said, Chiappori et al. (2014) show that the bias caused by a violation of this assumption is minimal, and the estimates predominantly identify the true preference parameters. We revisit this point in Section XX.

<sup>5</sup>Our approach employs a simple linear regression to estimate the correlation between household consumption with that of aggregate village consumption for each household. Our method differs from previous studies in several ways. Chiappori et al. (2014) and Mazzocco and Saini (2012) consider household-pairs and estimate correlation of consumption between each pair of households in the village. Chiappori et al. (2014) use moment conditions generated by consumption of each pair of distinct households in a village to impute preference parameters of each household by GMM. Mazzocco and Saini (2012) test for efficient risk sharing for household pairs by examining whether a households consumption is monotonically increasing with the sum of consumption for the pair. Their non-parametric test uses a risk-sharing function that allows for a general class of utility functions with heterogeneous risk preferences. On the other hand, Schulhofer-Wohl (2011) treats risk preferences as nuisance parameters that must be eliminated from the full risk sharing equation. To this effect, the author uses quasi-fixed effects that controls for household specific trends and household specific effects of aggregate shocks, thereby removing any heterogeneity in preferences.

MPC across households (e.g., Lewis et al., 2019; Aguiar et al., 2020). Earlier work based on the 1976-85 data indicated a higher level of risk pooling within sub-castes. Our findings from the more recent data show little difference between village and sub-caste results: sensitivity to cash-in-hand is very small in both cases – and possibly smaller within sub-castes.

Previous studies have explored several mechanism that limit full risk pooling.<sup>6</sup> Although our empirical analysis largely abstracts away from studying these mechanisms, it is the first to document a distorsion in consumption patterns without assuming homothetic preferences. In addition, at the end of the paper, we present a theoretical argument based on limited commitment for why villagers hold little liquid wealth.

## 2 Risk pooling with Assets: A conceptual framework

### 2.1 Adding liquid assets to the standard risk pooling model

Since we are interested in testing risk pooling within villages, we follow Townsend (1994), and many others and assume a closed village economy over time. In each period  $t$  each individual  $i$  in the village receives an earned income  $y_{its}$  that varies with the state of nature  $s$ . We assume that the joint income distribution of all individuals in the village is stationary over time with known mean, variance, and covariance vectors. This allows correlation in outcomes across individuals within period but, for simplicity, we abstract from autocorrelation of incomes across time.<sup>7</sup> The probability of state of the world  $s$  is denoted  $\pi_s$ . Each individual starts the period with liquid wealth  $(1 + r_s)w_{it}$  where  $w_{it}$  is the stock of liquid assets of individual  $i$  saved in the previous period and available at the beginning of period  $t$ , and  $r_s$  is the return to those assets, which is allowed to vary with the state of the world  $s$ . Each individual's cash-in-hand at the beginning of the period is thus  $x_{its} \equiv y_{its} + (1 + r_s)w_{it}$ . We restrict our attention to cases where the total liquid wealth of the village must be non-negative. But the individual net liquid wealth of individuals can be negative.

---

<sup>6</sup>Several explanations have been proposed for the failure of full risk pooling. One is moral hazard: when risk taking behavior is not observable to others, being insured may lead to excessive risk taking; this in turn may induce caps on insurance coverage (e.g., Rogerson, 1985; Golosov et al., 2003). Another is limited commitment: households receiving a high income draw may leave the insurance arrangement when asked to make a large contribution to the insurance pool; this limits the contributions households can be expected to make (e.g., Kimball, 1988; Coate and Ravallion, 1993; Ligon et al., 2002; Laczó, 2015). A third possibility is hidden income, opening the possibility of insurance fraud; this in turn results in some forms of income risk being better insured than others (e.g., Townsend, 1982). Kinnan (2022) examines several of these possibilities in a joint nested model: using panel data from rural Thailand, the author provides evidence that supports the hidden income hypothesis but rejects limited commitment and moral hazard. There is also an extensive literature on the role of network structures on risk-sharing (e.g., Ambrus et al., 2014; Ambrus and Elliott, 2020; Ambrus et al., 2021). For instance, Ambrus et al. (2014) show that the degree of risk-sharing is governed by the expansiveness of the network and that household consumption comoves more with that of socially proximate households. Ambrus et al. (2021) study local information constraints in risk-sharing networks and predict that network centrality is positively correlated with consumption volatility. Similarly, Chandrasekhar et al. (2018) show that the well-connected central agents engage more in risk-sharing when income risk is high, when income shocks are positively correlated, and when attitudes towards risk are more sensitive in the aggregate.

<sup>7</sup>Differences in the mean of income across individuals get subsumed in the welfare weights.

The utility that an individual derives from consumption expenditures  $c_{its}$  is given by a standard instantaneous utility function  $U_i(c_{its})$  specific to individual  $i$ . This allows for heterogeneous risk preferences. Each individual discounts the future with constant discount factor  $\rho_i$ , which similarly allows for heterogeneous time preferences.

We identify the Pareto efficient allocation of consumption across individuals within and across periods by solving a social planner problem of the form:

$$\begin{aligned}
& \text{Max}_{\{c_{its}, w_{its}\}} && \sum_{t=1}^T \sum_{i=1}^N \eta_i \rho_i^t \sum_{s=1}^S U_i(c_{its}) \pi_s \\
\text{s.t.} &&& \sum_i^N c_{its} = \sum_i^N ((1 + r_s)w_{it} + y_{its} - w_{it+1,s}) \quad \forall t, s \\
&&& \sum_i^N w_{it+1,s} \geq 0 \quad \forall t, s
\end{aligned} \tag{1}$$

where  $\eta_i$  is a particular set of (time-invariant) welfare weights with  $\sum_{i=1}^N \eta_i = 1$ . The middle equation (1) denotes the aggregate feasibility constraint that must hold in each time period  $t$  and state of the world  $s$ . To each particular set of welfare weights  $\{\eta_i\}$  corresponds a different Pareto efficient solution.

We now characterize the solution to the social planner's problem. We begin by noting that any income vector that has the same aggregate income  $y_{ts} = \sum_i^N y_{its}$  produces the same optimal solution. The same can be said for  $w_{ist}$ : any distribution of assets across individuals that generates the same total wealth  $w_{ts} \equiv \sum_i^N w_{its}$  generates the same optimal solution. It follows that the allocation of consumption across individuals does not depend on individual incomes and wealth: only village aggregates  $y_{ts}$  and  $w_{ts}$  matter. Put differently, the social planner's problem satisfies income and asset pooling: within each period, individual welfare does not depend on individual income or assets realizations; rather, it depends on welfare weights and individual preferences. Second, we note that since the return to wealth is linear and identical across individuals, the way assets are distributed across individuals is irrelevant and thus undetermined. This means that the solution to the social planner's problem does not stipulate the distribution of liquid assets across individuals – only its aggregate.

Next we note that the social planner's problem can be decoupled into an inner optimization problem – how to allocate consumption across individuals, conditional on a choice of future savings  $w_{t+1,s}$  for each  $s$  – and an outer optimization problem – how to allocate total consumption across periods by choosing the contingent path of  $\{w_{t+1,s}\}$ . The inner optimization problem takes the familiar risk sharing form:



$$\begin{aligned}
& \text{Max}_{\{c_{its}|w_{t+1,s}\}} \sum_{i=1}^N \eta_i \rho_i^t \sum_{s=1}^S U_i(c_{its}) \pi_s \\
& \text{s.t.} \quad \sum_i^N c_{its} = (1 + r_s)w_t + y_{ts} - w_{t+1,s} \equiv c_{ts} \quad \forall s
\end{aligned} \tag{2}$$

where  $w_{t+1,s}$  is taken as given. Since  $(1 + r_s)w_t + y_{ts}$  is predetermined by past savings and the state of the world  $s$ , and  $w_{t+1,s}$  is taken as given for the purpose of this optimization, the above optimization boils down to an allocation problem: how a given  $c_{ts}$  is divided among individuals. To characterize the properties of the solution, let us denote  $\lambda_{ts}\pi_s$  as the Lagrange multiplier associated with the feasibility constraint. The first order conditions for the consumption levels  $c_{its}$  and  $c_{jts}$  of two arbitrary individuals in the same village are:

$$\eta_i \rho_i^t U'_i(c_{its}) = \lambda_{ts} = \eta_j \rho_j^t U'_j(c_{jts}) \tag{3}$$

which implies the usual condition for optimal risk pooling: since all individuals face the same realization of the aggregate resource constraint  $\lambda_{ts}$ , weighted marginal utilities of consumption are equated across individuals in each state of the world  $s$ . Since  $\lambda_{ts}$  is a deterministic function of  $c_{ts}$ , this leads to the standard testable prediction: individual consumption  $c_{its}$  varies with aggregate village consumption  $c_{ts}$ , not with individual income  $y_{ist}$  or wealth  $w_{ist}$ . This theoretical result forms the basis for all tests of efficient risk pooling.

We now turn to the outer optimization problem that selects the contingent aggregate level of savings  $w_{t+1,s}$ . Let  $W_t(c_{ts})$  denote the value, to the social planner, of the *optimal* solution to the inner optimization problem for a total consumption level  $c_{ts}$ . Function  $W_t(\cdot)$  is indexed with  $t$  because, as we just discussed, when time preferences vary across individuals, the way the social planner divides the same amount of aggregate consumption  $c_{ts}$  across individuals varies over time. For clarity of exposition, let us define  $R(t) \equiv \sum_{i=1}^N \eta_i \rho_i^t$ . Since  $\sum_{i=1}^N \eta_i = 1$  by construction,  $R(t)$  is nothing but an average of individual discount factors weighed by the welfare weights.<sup>8</sup> Further, let us normalize individual discount factors as  $\hat{\rho}_i^t = \frac{\rho_i^t}{R(t)}$  such that  $\sum_{i=1}^N \eta_i \hat{\rho}_i^t = 1$ . With this normalization, the outer optimization can be written in the form of the following Belman equation:

$$V_t(x_{ts}) = \max_{w_{t+1,s}} W_t(x_{ts} - w_{t+1,s}) + R(t)EV_{t+1}((1 + r_{s'})w_{t+1,s'} + y_{ts'})$$

where  $s'$  denotes the (yet unrealized) state of nature in period  $t + 1$  and where we made use of the fact that  $c_{ts} = x_{ts} - w_{t+1,s}$ . This is a standard optimization problem (e.g., Stokey and Robert E. Lucas, 1989). It yields as solution a policy function of the form  $w_{t+1,s} = S_t(x_{ts})$ . The case with homogenous time preferences has been extensively studied in the precautionary savings literature (e.g., Zeldes, 1989a; Deaton, 1991). It is well known that  $c_{ts} = C_t(x_{ts})$  is a concave function of  $x_{ts}$ .

---

<sup>8</sup>Note that, as  $t \rightarrow \infty$ ,  $R(t)$  converges to the largest discount factor in the village.

## 2.2 Interpretation

In order to correctly interpret results from the above risk pooling model, it is essential to understand the logic and purpose of the testing methodology. First, it does not constitute a *causal* analysis of any kind. The theory makes predictions regarding specific correlations patterns that should or should not be present in the data, and we test for the presence of these patterns. These tests rest heavily on a simple accounting identity – consumption has to be financed somehow. They do not imply – or require – any causal process.

Second, we look at consumption smoothing after *income* smoothing has taken place. It does not matter for the validity of the tests whether or not individual households have already reallocated resources in response to an income shortfall – e.g., by increasing their supply of labor to the market (e.g., Kochar, 1999). The income shocks we consider here are those that the household did not absorb to keep disposable income unchanged. As long as the household is unable to completely smooth income on its own, the need for consumption smoothing remains and it must be achieved by other means, which are the focus of this paper.

Third, we have followed other papers in using the word ‘risk’ throughout. It is important to realize, however, that the tests we conduct do not require income shocks to be unanticipated or unpredictable. The need for consumption smoothing applies even for perfectly predictable income shortfalls, such as seasonal income variation or the birth of a child. While other papers have attempted to empirically distinguish between anticipated and unanticipated variation in income (e.g., Paxson, 1992, 1993; Kochar, 1999; Fafchamps and Lund, 2003), no such attempt is made here because our focus is elsewhere.

Finally, and perhaps most importantly, the methodology followed here does *not*, by itself, identify the mechanisms by which pooling/sharing is achieved. In particular, evidence of risk pooling at the village level does not imply the existence of a *deliberate* mechanism, whether formal or informal, for sharing risk among villagers. Indeed it has long been recognized that individual precautionary saving can closely approximate efficient risk pooling if individuals have sufficient liquid wealth: household with a shortfall cover their income deficit by dis-saving to smooth consumption while households with an income surplus save it for a rainy day (e.g., Deaton, 1992; Pan, 2009).<sup>9</sup> This being said, in poor village economies such as those studied here, the historical lack of liquid savings instruments has made it difficult if not impossible for households to accumulate wealth for consumption smoothing purposes. For this reason, it is commonly believed

---

<sup>9</sup>Sargent (1987), Chapter 3, provides a simple illustration of this fact. Imagine an isolated island in which people consume a non-storable food, e.g., ripening bananas from their orchard. Some days, more banana ripen, other days less. The banana harvest is also affected by common shocks, such as hurricanes. Each household is endowed with a stock of liquid wealth in the form of currency. With this currency, they purchase bananas when their daily harvest is insufficient, and sell bananas when they have more than they need. The price of bananas equilibrates the market such that, when there are few bananas in aggregate, the price is higher, and vice versa if there are bananas in abundance. In this economy, households with excess bananas de facto transfer them to households with a banana deficit by trading bananas against currency. As long as their stock of currency is sufficient, households can perfectly smooth their consumption against idiosyncratic shocks, and absorb common shocks by reducing consumption when the price of bananas is high. In this example, islanders pool risk while thinking they are self-insuring. Deaton (1992) provides an illustration of this model using data from Cote d’Ivoire.

that the explicit sharing of monetary and food resources among villagers plays an important role in how risk pooling takes place (e.g., Rosenzweig and Stark, 1989; Ligon et al., 2002; Pan, 2009; Kinnan, 2022). An important contribution of this paper is to revisit the role of precautionary savings with more contemporary data from rural India.

The validity of the tests nonetheless requires a number of conditions to be satisfied: the stock of assets  $w_{it}$  is measured before, not after, consumption during time interval  $t$ , so as to omit asset sales and loans that help the household smooth consumption at  $t$  (e.g., Zeldes, 1989a; Dubois, 2001); our definition of income does not include transfers and remittances received from – or given to – other households during period  $t$  since these often serve the purpose of smoothing consumption; we use first differences to control for household time-invariant characteristics that are correlated with its average consumption (e.g., Ravallion and Chaudhuri, 1997); and we observe a large enough proportion of households in a village to construct a sufficiently accurate measure of average village consumption (e.g., Suri, 2011). These issues are discussed in more detail below. As presented so far, the method also requires that all households share identical time and risk preferences (e.g., Mazzocco and Saini, 2012, 2018; Chiappori et al., 2014; Schulhofer-Wohl, 2011), a requirement that we dispense with in the next sub-section.

### 2.3 Accounting for heterogeneous time and risk preferences

The generalization of the methodology to heterogeneous time preferences does not change the main predictions. But the shape of  $C_t(x_{ts})$  changes over time. This is because the relative weights associated with ratios of marginal utility vary over time: if  $i$  is more patient than  $j$ , then  $\rho_i^t/\rho_j^t$  increases with  $t$ . This means that  $i$ 's expected share of aggregate consumption increases over time. This implies that early on, the social planner's discount factor  $R(t)$  puts more weight on impatient individuals. As time passes, their weight in the average  $\sum_{i=1}^N \eta_i \rho_i^t$  falls and  $R(t)$  gets dominated by the most patient individuals whose weight  $\rho_i^t$  falls less fast. This means that as time passes, the marginal propensity to consume  $\frac{\partial C(x_{ts})}{\partial x_{is}}$  out of village assets falls. With infinitely lived agents, in the long run all village cash-in-hand  $x_{is}$  is consumed by the most patient individual(s) since the welfare weight  $\eta_i \hat{\rho}_i^t$  of all the others converge more rapidly to 0. These are stark, unrealistic predictions that we do not expect to observe in practice, but they serve to outline the gradual unequalizing role that heterogeneous time preferences play in a risk pooling social optimum with assets.

Next we turn to the behavior of the model when individuals differ in their risk preferences. To this effect, we parameterize the utility function to have the constant-absolute-risk-aversion (CARA) form  $U_i(c) = -\frac{e^{-\gamma_i c}}{\gamma_i}$  where parameter  $\gamma_i$  is the coefficient of absolute risk aversion of individual  $i$ .<sup>10</sup> With this functional form, the first order condition (3) simplifies to:

$$\eta_i \rho_i^t e^{-\gamma_i c_{its}} = \lambda_{ts}$$

---

<sup>10</sup>Assuming constant relative risk aversion (CRRA) yields a similar result, except that estimating equations are expressed in logs rather than levels. See Appendix B for a formal derivation

Taking logs and rearranging yields, we get:

$$c_{its} = \frac{\log \eta_i}{\gamma_i} + \frac{\log \rho_i}{\gamma_i} t - \frac{1}{\gamma_i} \log \lambda_{ts} \quad (4)$$

Averaging over all  $N$  individuals in the village and solving for  $\log \lambda_{ts}$  yields an expression for average village consumption  $\bar{c}_{ts} \equiv \frac{1}{N} \sum_{i=1}^N c_{its}$ , which we use to replace the common Lagrange multiplier in equation (4). We obtain:

$$c_{its} = \frac{1}{\gamma_i} \left[ \log \eta_i - \frac{\frac{1}{N} \sum_{j=1}^N \frac{\log \eta_j}{\gamma_j}}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \right] + \frac{1}{\gamma_i} \left[ \log \rho_i - \frac{\frac{1}{N} \sum_{j=1}^N \frac{\log \rho_j}{\gamma_j}}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \right] t + \frac{1/\gamma_i}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \bar{c}_{ts} \quad (5)$$

which shows that the consumption of individual  $i$  is a linear function of the average individual consumption  $\bar{c}_{ts}$  and each parameter has been suitably normalized relative to its village average. Equation (5) shows that individual  $i$ 's consumption increases linearly in  $1/\gamma_i$ , which captures  $i$ 's willingness to bear risk. More risk averse individuals consume, other thing being equal, a smaller fraction of village consumption but, thanks to the intercept, their consumption is, as we would expect, more stable. This means that individuals who are less risk averse than the rest of the village consume less in bad years, but make up for it in good years, i.e., their consumption depends more on  $\bar{c}_{ts}$ . We also confirm that  $c_{its}$  increases in  $i$ 's relative welfare weight and relative discount factor, with the latter effect increasing over time as noted earlier.

## 2.4 Accounting for consumption categories

Next, we examine the predictions that the model makes regarding specific components of consumption, e.g., food and non-food non-durables. If individuals have homothetic preferences, income elasticities are unity for all goods and consumption shares are constant. This implies that consumption of good  $k$  is simply:

$$c_{itsk} = \alpha_k c_{its}$$

In this case, model (5) applies equally to all consumption goods, except that all coefficients are premultiplied by  $\alpha_k$ . This means that risk pooling can be tested with any component of consumption.

This is no longer the case when consumption preferences are not homothetic, i.e., when  $c_{itsk} = \alpha_k(c_{its})$  where  $\alpha_k(\cdot)$  now denotes an Engel curve. If the shape of this Engel curve can be estimated separately, e.g., from an analysis of the relationship between consumption shares and total consumption expenditures in a cross-section, model (5) can still be fitted to specific consumption categories provided the dependent variable is suitable transformed as:<sup>11</sup>

$$\hat{c}_{its}^k \equiv \hat{\alpha}_k^{-1}(c_{itsk}) \quad (6)$$

<sup>11</sup>For this transformation to yield a usable  $\hat{c}_{its}^k$  in our test, function  $\hat{\alpha}_k(\cdot)$  must be monotonic over the relevant range.

In an efficient risk pooling economy, applying model (5) to each  $\hat{c}_{its}^k$  should yield the same coefficient estimates. This would indicate that all consumption categories move with total expenditures in a way consistent with preferences across goods.

It is nonetheless conceivable that risk pooling focuses more on basic necessities such as food, but ignore luxuries. In this case, the consumption share of luxuries would fall faster with a fall of total expenditures than predicted by the Engel curve, i.e.,  $\hat{c}_{its}^k$  would vary more with  $\bar{c}_{ts}$  when  $k$  is a luxury than when  $k$  is food expenditures. This can be investigated by comparing coefficient estimates of model (5) applied to consumption categories  $\hat{c}_{its}^k$  with low and high income elasticities.

## 2.5 Summary of predictions

We now summarize the main model predictions regarding efficient risk pooling:

1. Individual consumption  $c_{its}$  is independent of individual cash-in-hand  $x_{its}$  and individual income  $y_{its}$ .
2. Individual consumption  $c_{its}$  is a function of village aggregate consumption expenditures  $c_{ts}$
3. Average village consumption  $\bar{c}_{ts}$  is a concave function of aggregate village cash-in-hand  $x_{ts}$  – i.e., the village smoothes consumption over time using all village liquid assets as pooled precautionary savings.
4. The share of village consumption that individuals receive falls over time if they are more impatient than the (suitably weighted) village average.
5. Individuals who are more risk averse than the (suitably weighted) village average receive, other things being equal, a smaller share of average village consumption. As a result their consumption is smoother than that of less risk averse individuals in the village.
6. The consumption of goods with a low income elasticity is smoothed more than the consumption of goods with a high income elasticity. Once transformed by the inverse of the Euler curve, expenditure shares on specific goods all respond identically to aggregate village expenditures  $c_{ts}$
7. Liquid assets are used to buffer income shocks at the village level.

## 3 Testing strategy

Before we present our testing strategy in detail, we must first recognize that, while the model presented in Section 2 applies at the individual level, in our data, as in most, consumption,

assets and income are all measured at the household level. As a result, we cannot estimate the extent to which risk is pooled within households (e.g., Dercon and Krishnan, 2000; Dunbar et al., 2013). We can only test whether it is pooled across households.

To do so in a way consistent with theory, we need to normalize the data in such a way that, if risk were perfectly pooled within and across households, our methodology would conclude that it is. This implies that, in order to obtain a correct village average  $\bar{c}_{ts}$ , we must weigh each household's per-capita consumption by the number of its members.<sup>12</sup> The same reasoning applies to income and assets, as well as to the risk pooling tests themselves. For this reason, all regressions presented in the paper are weighted by household size, so as to ensure that our tests aggregate individuals in a way that is consistent with theory.<sup>13</sup> In practice, we measure the size of each household by its number of adult-equivalents to reflect the fact that consumption needs vary by age and gender.

### 3.1 Standard risk pooling test

We test the risk pooling predictions, as listed in section 2.5, under both CARA (in levels) and CRRA (in logs) functional forms. We discuss CARA model test strategy in detail below and derive similar tests under the CRRA model in Appendix B. We start by testing the predictions of the CARA model under the assumption of homogeneous risk and time preferences. Since there is only one realized state of the world per time period, equation (5) simplifies to a perfect risk pooling relationship of the form:

$$c_{it} = \beta_i + \beta_1 t + \beta_2 \bar{c}_t \quad (7)$$

with  $\beta_1 = 0$  when all  $\rho_i$  are identical. It is important to note that assets are absent from this equation. This is because, thanks to predictions 2 and 3 above,  $\bar{c}_t$  is a sufficient statistic about the social planner's choice of future savings for the village. It follows that standard tests of risk pooling that use village average also work when the village saves.

As in the rest of the literature (e.g., Mace, 1991; Cochrane, 1991; Townsend, 1994; Ravallion and Chaudhuri, 1997), we first-difference equation (7) to eliminate the individual specific welfare weight term  $\beta_i$ . We also add two regressors: income  $y_{it}$ , which refers to the flow of income collected during period interval  $t$ ; and cash-in-hand  $x_{it}$  which refers to the sum of income and the stock of liquid assets of household  $i$  at the beginning of period  $t$  (measured in the month of July at the beginning of every panel year or at the end of previous panel year at  $t - 1$ ).

<sup>12</sup>This is best illustrated with a simple example. Imagine two households 1 and 2, respectively with 1 and 2 members. Total consumption in household 1 is 100, which is also the consumption per head. In household 2, total consumption is also 100, which means that consumption per head is 50. If we take the simple average of consumption per head across the two households we obtain average village consumption of  $75 = \frac{1}{2}100 + \frac{1}{2}50$ . If, however, we average across individuals, the average village consumption is  $66.67 = \frac{1}{3}100 + \frac{2}{3}50$ .

<sup>13</sup>To the best of our knowledge, however, this easy correction is not implemented by Mace (1991); Cochrane (1991); Townsend (1994), and those that followed in their footsteps (e.g., Schulhofer-Wohl, 2011; Mazzocco and Saini, 2012; Chiappori et al., 2014).

The estimated CARA model thus has the form:

$$\Delta c_{it} = \beta_1 + \beta_2 \Delta \bar{c}_t + \beta_3 \Delta y_{it} + \beta_4 \Delta x_{it} + \epsilon_{it} \quad (8)$$

We include cash-in-hand in addition to the income variable that traditionally appears in the efficient risk pooling test in order, so as to also allow for the possibility of precautionary saving at the household level. Under the null of efficient risk pooling at the village level, the coefficients of both household income and cash-in-hand should be 0 while the coefficient of average village consumption should be 1. In contrast, under a pure household precautionary savings model, household consumption should not vary with village consumption or household income, but it should respond to variations in household cash-in-hand (Zeldes, 1989a,b) – in which case its coefficient is the so-called marginal propensity to consume. If we find a significant coefficient for income *and* cash-in-hand (in which income is included) – this indicates excess sensitivity to income, e.g., because of asset liquidation costs or mental accounting. Since the household can make up for certain income shocks by engaging in other income earning activities – e.g., redirecting labor supply to the market – the regression model (8) examines the extent to which unsmoothed variation in income is reflected in household consumption. Since cash-in-hand is measured at the beginning of time interval  $t$ , it does not include any transfer of funds across households that may have served to smooth period  $t$  consumption – which is exactly what we want.

Given the relatively small number of households within each village, the mechanical correlation between  $c_{it}$  and  $\bar{c}_t$  generates a bias in  $\beta_i$  when the null of perfect risk pooling is false (see Appendix A for illustration).<sup>14</sup> To correct for this bias, we estimate (10) by replacing the village mean  $\bar{c}_t$  by the leave-out-mean  $\bar{c}_{-i,t} \equiv \frac{1}{N-1} \sum_{j \neq i} c_{jt}$ .<sup>15</sup>

Our main null hypothesis is that risk pooling is efficient, which implies that  $\beta_2 = 1$  and  $\beta_1 = \beta_3 = \beta_4 = 0$ . Equation (8) also enables us to consider the following alternative hypotheses:

1. Hand-to-mouth: Each individual consumes his or her income  $y_{its}$ , which implies  $\beta_3 = 1$  and  $\beta_1 = \beta_2 = \beta_4 = 0$
2. Individual precautionary saving: Each individual consumes a concave fraction of his or her cash-in-hand  $x_{its} \equiv y_{its} + w_{its}$ , which implies that  $\beta_3 = \beta_4 > 0$  and  $\beta_2 = 0$
3. Individual precautionary saving with excess sensitivity to income:  $\beta_3 > \beta_4 > 0$  and  $\beta_2 = 0$
4. Partial pooling of income but full pooling of assets:  $1 > \beta_2 > 0$  and  $1 > \beta_3 > 0$  and  $\beta_4 = 0$
5. Partial pooling of income and assets:  $1 > \beta_2 > 0$  and  $\beta_3 > 0$  and  $\beta_4 > 0$

This regression is complemented by a village level analysis to test whether the village collectively uses assets to smooth consumption. The estimated regression is the standard test of the

<sup>14</sup>For instance, when the true  $\beta_i = 0$  in equation (10), the OLS estimate has a bias equal to  $1/N$ .

<sup>15</sup>It is easy to show that, under the null of perfect risk pooling, estimating (10) with the leave-out-mean still yields the correct estimate of  $\hat{\beta}_i = 1$  but multiplies  $\alpha_i$  by  $\frac{N}{N-1}$ . See Appendix A for details.

precautionary saving developed by Zeldes (1989b). It takes the following form:

$$\Delta \bar{c}_t = \beta_1 + \beta_3 \Delta y_t + \beta_4 \Delta x_t + \epsilon_t \quad (9)$$

Efficient precautionary saving requires asset integration, which implies that consumption responds to cash-in-hand ( $1 > \beta_4 > 0$ ): when available cash falls, consumption is reduced, but less than one-for-one, and consumption should not respond to current income ( $\beta_3 = 0$ ), since it is already included in cash-in-hand. If the village does not use assets to smooth consumption across periods, then  $\beta_1 = \beta_4 = 0$  and  $\beta_3 = 1$ . It is also conceivable that the village achieves a modicum of intertemporal consumption smoothing from other sources that are not identified in the data (e.g., external transfers from migrants, government, or NGOs), in which case  $\beta_1 > 0$  and  $1 > \beta_3 \geq 0$ . We also estimate (9) in CRRA log form.

### 3.2 Consumption categories and Engel curves

Next we estimate inverse Engel curves (equation 6) for various consumption goods. This is achieved by non-parametrically regressing total expenditures  $c_{its}$  on expenditures  $c_{itsk}$  on good  $k$ . We do this using cross-section data, which means that the income elasticities embedded in these inverse Engel curves are estimated using variation in expenditure shares across households with different total levels of expenditures. We then use the fitted model  $\hat{\alpha}_k^{-1}(c_{itsk})$  to obtain a prediction of total expenditures  $\hat{c}_{its}^k$  for each household in each period. If households are unconstrained in the consumption choices they make after risk sharing, they should, on average, be on their Engel curve for each good. In contrast, if assistance from the village favors certain goods – e.g., food<sup>16</sup> – then households should spend a higher proportion of their total expenditures on food when they receive assistance. This observation forms the basis of our test.<sup>17</sup>

To implement this idea in the simplest way, model (8) is estimated separately for each  $\Delta \hat{c}_{its}^k$  dependent variable. We then test whether estimated coefficients  $\beta_2$ ,  $\beta_3$ , and  $\beta_4$  are identical across consumption goods. This constitutes an alternative test of the perfect risk pooling model. The alternative is that certain expenditures are better insured than others – e.g., luxuries are consumed when the village as a whole is enjoying a higher income and, in bad times for the village, individual consumption patterns are adjusted towards food consumption. Differentiated insurance is a common occurrence in all societies: social safety nets typically seek to guarantee individuals a minimum consumption level, with a focus on necessities such as food, shelter, and basic clothing – but typically excludes luxuries. To verify whether this pattern is also present in our data, we test whether  $\beta_2$  is smaller (i.e., less sensitive to aggregate shocks) for goods with a low income elasticity, and vice-versa for goods with a high income elasticity, such as luxuries.

<sup>16</sup>as the US Food Stamps welfare program used to do.

<sup>17</sup>To illustrate with an example, imagine that a household optimally spends 700 Rps on food and 300 on non-food when its total expenditures is 1000 Rps, and 800 on food and 400 on non-food when its total expenditures is 1200. Then if this household is unconstrained and we observe it to spend 800 Rps on food, its total expenditures should be 1200. If, at the same time, we observe it consuming 300 on non-food, we would predict that its total expenditures is 1000. Hence a systematic discrepancy between the two predicted values of total expenditures  $\hat{c}_{its}^{food}$  and  $\hat{c}_{its}^{non-food}$  indicates that consumption choices are constrained.



### 3.3 Heterogeneous risk and time preferences

We now introduce heterogeneity in risk and time preferences across households. It is well known that tests of risk pooling are biased in the presence of heterogeneous risk preferences (e.g., Schulhofer-Wohl, 2011; Mazzocco and Saini, 2012; Chiappori et al., 2014). Ignoring heterogeneous risk preferences leads to an upwards bias in  $\hat{\beta}_2$ , the coefficient of  $\bar{c}_t$  in equation (8). This is because  $\epsilon_{it}^{homog} = \left(\frac{1}{\gamma_i} - \frac{1}{\gamma}\right)\bar{c}_t + u_{it}$ , which introduces a positive correlation between  $\bar{c}_t$  and the error term.

To address this issue, we proceed in two steps. We do not have (reliable) information on monthly income and assets: this information is only available annually. But we do have reliable information on monthly consumption for each households over a period of five years (60 months). We therefore have enough observations to fit a perfect risk sharing model to each household separately, since this model does not require data on income and assets. This yields estimates of household-specific risk and time preference parameters that are consistent under the null of perfect risk sharing, and thus can then be used, in a second step, to re-estimate model (8) with household-specific risk and time preferences on annual data.

We follow in the footsteps of Chiappori et al. (2014) and Mazzocco and Saini (2012) and test for heterogeneity in preference parameters under the assumption of perfect risk sharing. The assumption of full risk sharing is important for our estimates. If full risk sharing does not hold, the method may not consistently estimate the true preference parameters. However, Chiappori et al. (2014) show that the bias caused by the violation of this assumption is minimal, and the estimates predominantly identify the true preference parameters.<sup>18</sup> Therefore, while we acknowledge that perfect risk sharing is required for our estimates of the preference parameters to be consistent, the small size of the likely bias means that we can interpret the estimates as largely reflecting the true preference parameters even if perfect risk sharing is rejected.

Formally, in the first step we estimate model (5) separately for each of the 1300 households in our data. To achieve this, we start by noting that, as pointed out by Wilson (1968), doubling every household's coefficient of risk aversion does not change the set of Pareto-efficient allocations. This means that absolute risk preferences cannot be identified – but relative risk preferences can. To reflect this, we follow Chiappori et al. (2014) and normalize risk preferences up to a

---

<sup>18</sup>Chiappori et al. (2014) show that, if full insurance does not hold, then the method rather identifies a weighted linear combination of true risk tolerance and cyclicity of income. The study further performs back of the envelope calculations and shows that, even if full insurance fails, the estimated risk preferences are a combination of at least 91 percent risk tolerance and at most 9 percent 'cyclicity of income'. The latter is defined by the authors (see page 9, equation 10) as the coefficient of a regression in which the log of household income is regressed on a household-specific intercept, a household-specific trend, household-specific monthly dummies, and an average village-specific income shock. Applying this approach to our annual data, monthly dummies drop out and we estimate a cyclicity of income coefficient of 0.83. Applying this bias estimate to equation (13) in Chiappori et al. (2014), we see that, if full insurance fails but is 'at least as good as in the United States', our risk tolerance estimates are a mixture of 91 percent true risk tolerance and  $0.09 \times 0.83 = 0.075$ . For instance, if the true risk tolerance is 2, our estimated risk tolerance would be  $0.91 \times 2 + 0.05 = 1.9$ . Similar calculations for time preference using equation (11) in Chiappori et al. (2014) and our estimate of the coefficient of the time trend of 0.011 imply that, if fully insurance fails, our estimates of  $\frac{\log \rho_j}{\gamma_j}$  would equal to 0.91 times the true value  $+0.09 \times 0.011 = 0.01$ . For instance, for a household with discount factor 0.95 and risk tolerance 2, the estimated  $\rho_j$  would be 0.956.

village-specific scale by setting  $\frac{1}{N} \sum_{i=1}^N \frac{1}{\gamma_i} = 1$ . With this normalization, equation (5) reduces to:

$$c_{it} = \frac{1}{\gamma_i} \left[ \log \eta_i - \frac{1}{N} \sum_{j=1}^N \frac{\log \eta_j}{\gamma_j} \right] + \frac{1}{\gamma_i} \left[ \log \rho_i - \frac{1}{N} \sum_{j=1}^N \frac{\log \rho_j}{\gamma_j} \right] t + \frac{1}{\gamma_i} \bar{c}_{-i,t}$$

We use 60 months of consumption data to estimate, for each household  $i$ , an OLS model of the form:

$$c_{it} = \alpha_i + \theta_i t + \beta_i \bar{c}_{-i,t} + \epsilon_{it} \quad (10)$$

The mapping between estimated coefficients and structural parameters is given by:

$$\beta_i = \frac{1}{\gamma_i} \quad (10A)$$

$$\alpha_i = \beta_i \left[ \log \eta_i - \frac{1}{N} \sum_{j=1}^N \beta_j \log \eta_j \right] \quad (10B)$$

$$\theta_i = \beta_i \left[ \log \rho_i - \frac{1}{N} \sum_{j=1}^N \beta_j \log \rho_j \right] \quad (10C)$$

Coefficient  $\beta_i$  represents the risk tolerance (i.e., the inverse of risk aversion  $\gamma_i$ ) of individual  $i$  relative to the village mean – e.g.,  $\beta_i > 1$  implies that  $i$  is more risk tolerant than others in the village, and as a result has a consumption level that varies more than others with the village average.<sup>19</sup>

Estimates of structural parameters  $\gamma_i$ ,  $\eta_i$ , and  $\rho_i$  can be recovered from OLS estimates of  $\hat{\beta}_i$ ,  $\hat{\alpha}_i$ , and  $\hat{\theta}_i$ , subject to suitable normalization. For time preferences, the same reasoning applies as for time preferences: only relative preferences can be recovered from (10). We therefore set  $\frac{1}{N} \sum_{j=1}^N \rho_j = 1$ . For welfare weights, we follow convention and normalize them to sum to 1 within each village. With these normalizations,  $\hat{\gamma}_i = 1/\hat{\beta}_i$  and estimates of relative welfare weights  $\eta_i$

<sup>19</sup>The reader may wonder whether the recovered risk coefficient  $\beta_i$  measures the curvature of the instantaneous utility function  $U_i(c_{its})$  or that of household  $i$ 's value function. In the original Townsend (1994) model, there are no assets and, consequently, the two coincide. In a precautionary savings model with liquid assets, the household's intertemporal value function inherits part of the curvature of  $U_i(c_{its})$  and the curvature of the value function increases monotonically with the curvature of  $U_i(c_{its})$  (e.g., Stokey and Robert E. Lucas, 1989; Deaton, 1991). Hence the estimated risk tolerance coefficient  $\beta_i$  can be said to be isomorphic to the risk aversion coefficient of  $U_i(c_{its})$ . This statement nonetheless requires the maintained assumption that all households share the same technology for accumulating liquid wealth. Indeed, it is well known that, the more flexible the liquid wealth technology is, the less curvature the value function has relative to its corresponding instantaneous utility function. For instance, if the liquid wealth technology allows a perfect smoothing of consumption, the value function becomes linear in cash-in-hand, irrespective of how much risk aversion is exhibited by  $U_i(c_{its})$ .

and relative time preference parameters  $\rho_i$  can be recovered using the following formulas:<sup>20</sup>

$$\log \eta_i = \frac{\alpha_i}{\beta_i} + \log \left[ \frac{1}{\sum_{j=1}^N e^{\frac{\alpha_j}{\beta_j}}} \right] \quad (10A)$$

$$\log \rho_i = \frac{\theta_i}{\beta_i} + \log \left[ \frac{1}{\frac{1}{N} \sum_{j=1}^N e^{\frac{\theta_j}{\beta_j}}} \right] \quad (10B)$$

This yields a set of household-specific estimates of  $\hat{\gamma}_i$ ,  $\log \hat{\eta}_i$ , and  $\log \hat{\rho}_i$ , all estimated under the maintained assumption of perfect within-village risk sharing. These estimates represent what the relative welfare weights and the relative risk and time preferences of households *would be* if income risk is perfectly shared among villagers. Therefore, our method of estimating preferences is simple to implement and relies on a simple linear regression, and differs from previous studies in important ways.<sup>21</sup>

The second step of our test is to use these inferred parameters to control for household-level risk and time preferences when estimating our risk pooling test with annual data on income and wealth. We extend model (8) to allow for household heterogeneity in risk and time preferences as follows:

$$c_{it} = \alpha_i + \theta_i t + \beta_i \overline{c_{-i,t}} + \xi y_{it} + \zeta w_{it} + \epsilon_{it} \quad (11)$$

As before, perfect risk pooling requires that  $\xi = 0$  and  $\zeta = 0$ . To test this prediction, we write (11) so as to eliminate all the household-specific coefficients. First,  $\alpha_i$  is eliminated by first-differencing the data. Second, we use the  $\hat{\beta}_i$  and  $\hat{\theta}_i$  estimates obtained in the first step to create an estimable model of the form:<sup>22</sup>

$$\Delta \left( c_{it} - \hat{\beta}_i \overline{c_t} \right) - \hat{\theta}_i = \xi \Delta y_{it} + \zeta \Delta w_{it} + \Delta \epsilon_{it} \quad (12)$$

---

<sup>20</sup>Let  $X = -\frac{1}{N} \sum_{j=1}^N \beta_j \log \eta_j$  and  $\phi_i = \frac{\alpha_i}{\beta_i}$ . From equation (10C) we get  $\log \eta_i = \phi_i + X$  (\*) and thus  $\eta_i = \exp^{\phi_i + X}$ . By the normalization of welfare weights  $\sum_{j=1}^N \eta_j = 1$  we get  $\sum_{j=1}^N e^{\phi_j + X} = 1$ , which implies  $X = \log \left[ \frac{1}{\sum_{j=1}^N e^{\phi_j}} \right]$ . Substituting  $X$  back into (\*) yields the reported formula. A similar approach yields the  $\rho_i$  formula, except for the division by  $N$  which comes from the different normalization rule  $\frac{1}{N} \sum_{j=1}^N \rho_j = 1$ .

<sup>21</sup>Chiappori et al. (2014) and Mazzocco and Saini (2012) consider household-pairs and estimate correlation of consumption between each pair of households in the village. Chiappori et al. (2014) use moment conditions generated by consumption of each pair of distinct households in a village to impute preference parameters of each household by GMM. Mazzocco and Saini (2012) test for efficient risk sharing for household pairs by examining whether a households consumption is monotonically increasing with the sum of consumption for the pair. Their non-parametric test uses a risk-sharing function that allows for a general class of utility functions with heterogeneous risk preferences. On the other hand, Schulhofer-Wohl (2011) treats risk preferences as nuisance parameters that must be eliminated from the full risk sharing equation. The author uses quasi-fixed effects that controls for household specific trends and household specific effects of aggregate shocks, thereby removing any heterogeneity in preferences.

<sup>22</sup>Although regression model (12) makes use of predicted variables  $\hat{\beta}_i \overline{c_t}$  and  $\hat{\theta}_i$ , estimates of  $\xi$  and  $\zeta$  are not subject to sampling error since the constructed variables only appear in the dependent variable (e.g., Murphy and Topel, 1985).

## 4 Data

We use the new wave of ICRISAT’s VDSA (Village Dynamics of South Asia) panel data of about 1400 households observed over 60 consecutive months from July 2010 to June 2015.<sup>23</sup> Households were randomly selected from 30 villages in eight Indian states, chosen to represent the agro-climatic conditions in India’s semi-arid and humid tropical regions.<sup>24</sup> Households in each village were randomly selected to represent households in four landholding classes: large, medium, small, and landless. The data collection timeline follows the agricultural cycle in India, beginning from July to June. Attrition in the VDSA data is minimal - only about 10% of households have an unbalanced panel of less than 60 months of data. For our analysis, we use a balanced panel of 1,296 households that reported 60 months of consecutive monthly data.

To construct the main consumption outcomes, we use data on food expenditures, non-food expenditures and total expenditures collected every month for each household. Food consumption includes all food items sourced from home production and purchases. Non-food consumption includes expenses on services and utilities such as travel, education, medical, and energy.

Our measure of earned income includes all net earnings from crops, livestock and off-farm labor. We also include value of food consumed from home production as part of earned income. Crop and livestock income is calculated as the revenue from sales of crop and livestock products, minus production costs that include the value of material inputs and the imputed cost of own labor. Off-farm labor income is the sum of earned wages for all household members and the net income earned from household businesses. The majority of individuals in the sample are at least partially employed in the casual labor market. A few individuals are employed in business or a salaried job in the formal sector.

In the analysis, we use cash-in-hand as a measure of household assets. Cash-in-hand is constructed as the sum of liquid wealth, earned income and unearned income. Liquid wealth is defined as the sum of the household’s net credit position (savings, minus borrowing plus lending) and the value of liquid assets such as livestock, consumer durables (gold and silver, cars and motorcycles, furniture and farm equipment). Unearned income is the sum of income from government transfers and rental income received from renting out land, household or farm equipment.

Although the VDSA has rich monthly data on consumption and income, household assets are only measured annually at the beginning of each panel year, which coincides with the onset of the

---

<sup>23</sup>ICRISAT’s new wave of VDSA panel data is a continuation of Village level studies (VLS) panel of household data collected between 1975 to 1985 in six villages in the semi-arid tropics of India. In the VDSA data, in addition to the 6 old VLS villages, 12 more villages in the semi-arid tropics and 10 more villages from East India were included, summing to a total of 30 villages across 8 states in India. The VDSA data collection started in 2009, however, the data for panel year 2009 has a gap between July to October 2009. To maintain consistency, this paper uses data beginning from panel year 2010 until 2014.

<sup>24</sup>The eight states are Andhra Pradesh, Bihar, Gujarat, Jharkhand, Karnataka, Madhya Pradesh, Maharashtra, and Orissa). Four villages were selected from each state, except in Madhya Pradesh where only two were selected. See figure A1 in Appendix for the precise location of the 30 villages.

main agricultural season in July. Consequently, all regressions that require asset information are estimated by aggregating monthly data on household consumption and income to the beginning of the agricultural cycle. All values are deflated and expressed in 2010 Indian rupees. Income, consumption, and assets are expressed per capita by dividing them by their adult-equivalent weight.<sup>25</sup> We winsorize the top and bottom 1% of the data to remove outliers and large measurement errors.

Table 1 presents descriptive statistics for the main variables used in the analysis. Annual consumption expenditures per adult equivalent are on average Rs. 15,937 in 2010 rupees. This is equivalent to 3 US\$ per day and per adult-equivalent, based on a purchasing parity rate of 14.6 Rs. per US\$ in 2010 (World Bank, 2014).

## 5 Preparatory analysis

### 5.1 Engel curves

Before launching the main part of our analysis, we complete the preparatory analysis on Engel curves and heterogeneous preferences. We begin by fitting Engel curves to annual consumption data in 2011, a good rainfall year when village cash-in-hand is the highest in our data. This is the year in which risk pooling would be least likely to impose constraints on consumption. Figure 1 uses a flexible polynomial to plot household budget shares against the log of total household expenditure per capita. The Engel curve for food is approximately log-linear and downward sloping, confirming that the food expenditure share falls with income, in accordance with the literature. The poorest quintile of the income distribution spend about 65% of their total budget on food, whereas the richest quintile spend about 40%. The inverse is true for non-food expenditures. These results constitute strong evidence against homotheticity in food and non-food preferences. Since the computed Engel curves are monotonic, they can be inverted to obtain the function  $\hat{\alpha}_k^{-1}(c_{itsk})$ .

### 5.2 Risk and time preferences

Next, we estimate individual risk and time preferences as described in Section 3.3. To recall, we estimate the CARA-based regression (10) for each household, only using monthly data on consumption expenditures. We then use the results to recover estimates of absolute risk aversion and time preference by using the formulas reported in Section 3.3. The risk tolerance measure  $\beta_i$  is normalized to a village-specific scale – i.e., mean risk tolerance of each village is set to one. To recall, risk tolerance is the inverse of the coefficient of risk aversion. The estimate of the

---

<sup>25</sup>Following Townsend (1994), we define the age-sex weights as : 1.0 and 0.9 for adult males and females; 0.94 and 0.83 for adolescent males and females aged 13-18, 0.67 for children aged 7-12 regardless of gender; 0.52 for toddlers 1-3 and 0.05 for infants

discount factor  $\rho_i$  is similarly normalized to average to one in each village. These normalizations arise from the fact that only relative values of risk aversion and time preference can be inferred from the coefficients of regression (10). We also estimate a CRRA version of these parameters using a model similar to (10), but in logs. Apart from the normalizations, it is important to remember that the estimated risk and time preference parameters are obtained under the *maintained assumption* of perfect risk pooling and that their main purpose is to test perfect risk pooling in the heterogeneous-corrected regression model (12). This being said, these estimated parameters contain valuable information that we summarize here.

In Figure 2 we plot the distributions of the household-specific estimates of risk tolerance  $\hat{\beta}_i$  under both the CARA and CRRA models.<sup>26</sup> These parameters are identified from whether household  $i$ 's consumption varies more than that of household  $j$ : if it does – and we are in a perfect risk pooling equilibrium – then  $i$  must be less risk averse than  $j$ . Both sets of estimates are normalized to have a mean equal to 1 within each village, which means that they capture relative risk tolerance rather than absolute values. Since the CARA and CRRA estimates are not measured in the same units, their magnitude is not directly comparable; but their frequency distribution is.

Overall, we find considerable heterogeneity in risk tolerance within villages, suggesting that, if we are in a perfect risk pooling equilibrium, large welfare benefits are achieved not only from pooling risk, but also from shifting risk from highly risk averse households to more risk neutral ones. Since more variation in  $\beta_i$  around its average  $\bar{\beta}$  translates into more correlation between  $\bar{c}_t$  and  $u_{it}$  in  $\epsilon_{it}^{homog} = (\beta_i - \bar{\beta}) \bar{c}_t + u_{it}$ , Figure 2 also constitutes prima facie evidence that ignoring heterogeneity in risk preferences may bias risk pooling tests that assume homogeneous risk preferences.

Figure 3 shows the distribution of household-specific estimates of discount factors  $\hat{\rho}_i$ . As for risk preferences, the  $\hat{\rho}_i$  are normalized to have unit mean within each village. Discount factors are identified from whether household  $i$ 's consumption increases more over time than that of household  $j$ : if it does – and we are in a perfect risk pooling equilibrium – then  $i$  must be more patient than  $j$ . We see that, visually, there is much less dispersion in discount factors than in risk tolerance. This suggests that, if the study area is in a perfect risk pooling equilibrium, few welfare gains are achieved by accommodating differences in impatience. This being said, even small differences in discount factors can, over time, translate into increasing differences in consumption levels across households, even if they share the same risk preferences and the same welfare weight.

---

<sup>26</sup>Because each  $\hat{\beta}_i$  is estimated from a regression with 60 observations, its distribution suffers from excess variance due to sampling error. To assess the magnitude of the excess variation that this error generates in Figures 2 and 3, we shrunk the sample distribution of  $\hat{\beta}_i$  in such a way that, when we add the sampling noise to the 'shrunk'  $\hat{\beta}_i$ , we obtain a frequency distribution with the same variance as that of the original  $\hat{\beta}_i$ . In practice, this procedure entails turning the original  $\hat{\beta}_i$  distribution to have zero mean, and using the standard error of  $\hat{\beta}_i$  in each regression as estimate of the sampling error in that sample. Using this approach shrinks the standard deviation of the estimated parameters by 17% on average, while respecting the general shape of the original distribution. The qualitative conclusions reported here are not affected, however.

### 5.3 Inequality

We also infer welfare weights  $\hat{\eta}_i$  from estimated coefficients from regression (10), using the formula presented in Section 3.3. The welfare weights are identified from the household-specific intercept in model (10). In the CARA model, this intercept measures the minimum level of consumption that is assigned at  $t = 0$  to household  $i$  in a perfect risk pooling equilibrium. Keeping risk and time preferences constant, and keeping  $\bar{c}_t$  the same, household  $i$  consumes more than  $j$  if  $i$  has a higher  $\hat{\eta}_i$  than  $j$ . In the CRRA version, the intercept is the base *share* of consumption that goes to  $i$ , but the reasoning is the same: *ceteris paribus*,  $i$  consumes more if  $i$  has a larger welfare weight. Following common practice, the welfare weights themselves are normalized to sum to 1 within each village. It follows that equal treatment of all households in a village of size  $N_v$  requires that they all have  $\hat{\eta}_i = 1/N_v$ . Since  $N_v$  varies across villages, it is useful to take  $1/N_v$  as yardstick to judge intra-village inequality.

Using this approach, we calculate, for each village, the proportion of households for whom  $\hat{\eta}_i < 1/N_v$ . The larger this proportion is, the more unequal the distribution of welfare is in the village. We present in Figure 4 a histogram of these proportion across all the villages in our study. While there is some variation between the histograms depending on whether the welfare weights were estimated using CARA or CRRA, it is nonetheless clear that welfare weights are quite unequal in most villages. Across all villages, the proportion of households with welfare weights less than the equitable share  $1/N_v$  is estimated to be about 88% from the CARA version of regression (10), and 89% for the CRRA version.

Figure 4 presents two histograms of village-averages of the 30 villages, one from the CARA  $\hat{\eta}_i$  estimates and the other from the CRRA estimates. In most of the villages, more than 95% of households have welfare weights less than  $1/N_v$ , and the overwhelming majority of them have 80% or more of households below the average welfare weight of  $1/N_v$ . This implies that these villages have 20% or less of their households enjoying above average welfare weights – and thus consistently above average consumption across time. We also find that the frequency distribution of  $\hat{\eta}_i$  has a fat upper tail, with some households receiving welfare weights close 1, indicative of very high consumption inequality. Keeping in mind that these estimates all assume perfect risk pooling at the village level, they remind us that risk pooling is not equivalent to income redistribution – and that it is quite compatible with a lot of consumption inequality in equilibrium. This inequality would be further reinforced in good years if richer households – i.e., those with a high welfare weight and, thus, a high average consumption – are also those who are more risk neutral, as is likely.

## 6 Main empirical analysis

We now turn to our main estimation results. We start by reporting the results of the risk pooling tests under the assumption of CARA utility. We then repeat the exercise for the CRRA model

to check the robustness of our findings. Next we estimate the extent of precautionary saving at the village level. In the last part of this Section, we re-estimate our main results for perfect risk pooling within castes (Jatis) in the same villages, instead of within villages.

## 6.1 CARA model

Table 2 summarizes all our test results for perfect risk pooling within villages under a CARA model. As explained in Section 3, these estimates are obtained from first-difference regressions in levels, using a pooled panel of all the sample households. In all regressions, standard errors are clustered at the village-level. Panel A in Table 2 reports the test results under homogeneous preferences. As shown in Column (1) Panel A, for total expenditures, we fail to reject full risk pooling of income. These results are in contrast to earlier studies that reject full risk pooling (e.g., Townsend, 1994; Ravallion and Chaudhuri, 1997). The small magnitude and statistical insignificance on coefficient on income earned suggests substantial pooling of idiosyncratic income risk. However, we reject full pooling of assets. The small magnitude of the coefficient on cash-in-hand nonetheless suggest substantial risk pooling. A Rs.100 change in annual cash-in-hand is associated with Rs. 2.1 change in annual consumption, all measured in real 2010 rupees per-adult equivalent. The coefficient on average village expenditure (0.769) is significantly different from 1 at the 1% significance level. This indicates extensive (even if not perfect) mutual risk pooling within villages, thereby rejecting the pure hand-to-mouth or household precautionary savings models discussed in Section 2.5.<sup>27</sup> Taken together, these findings are consistent with a pooling of income and assets that is partial but nonetheless achieves a considerable amount of co-movement in household consumption across years. We can also reject a model in which assets are pooled for risk purposes but incomes are not. We find no evidence of excess sensitivity of consumption to household income

Next, we examine the implications of risk pooling separately for food and non-food, under homothetic and non-homothetic preferences. Under homothetic preferences, consumption shares are assumed constant with income, which implies that, in the absence of constraints on consumption, each consumption expenditure category should, on average, respond equally to village average expenditures as well as income and asset shocks. If, in contrast, food consumption is better insured than non-food consumption, food expenditures should respond less to village average expenditures and less to variation in household assets and income. Results under the assumption of homothetic preferences are shown in columns (2) and (3) for food and non-food expenditures, respectively. To maintain direct comparability for column (1), food and non-food expenditures are divided by the *average* budget share for food and non-food, respectively. For instance, if the average food share is 50%, the food expenditure variable is multiplied by 2, mak-

---

<sup>27</sup>Because we do not observe all the households in each village, this coefficient may be affected by measurement error in  $\bar{c}_{-i,t}$  due to sampling. In Appendix C we offer a simple back-of-the-envelope correction for the downward bias due to sampling error in  $\bar{c}_{-i,t}$ . Applied to this coefficient, the correction yields a point estimate of  $0.769/0.976 = 0.787$ . In analysis not shown here to save space, we also estimate the contrast estimator proposed by Suri (2011) and we obtain similar results to those presented here. Our results are thus unlikely to be driven by measurement error in average village expenditures.



ing it comparable to the total expenditure variable used in column (1). In Table 2 column (2), we see that - food expenditures respond less to average village expenditures. We also note that non-food expenditures respond more strongly to variations in household cash-in-hand, suggesting less smoothing of non-food expenditures across individuals. Taken together, these findings suggest that year-to-year variation in household expenditures on food and non-food depart from what cross-section expenditure shares would suggest. This indicates that food expenditures are not only less responsive to assets shocks, but also that they fluctuate less with average village consumption than non-food expenditures – i.e., they are better protected from collective shocks. The contrast between food and non-food expenditures suggests that this is achieved by smoothing food consumption from collective shocks more than non-food expenditures.

This interpretation, however, can be misleading because it incorrectly assumes that consumption preferences are homothetic. Under non-homothetic preferences, the approach must be amended to account for the systematic variation of expenditure shares with total expenditures. To correct for this, we transform the dependent variable by the inverse of the Engel curve (see Section 3.2 for details). As a result, it becomes the level of total expenditures that is predicted from observed food expenditures and cross-section Engel curve estimates for year 2011.<sup>28</sup> Using this approach, we can test whether variation in food and non-food expenditures responds similarly to village average expenditures and income and asset fluctuation – as they should if total consumption is redistributed among households but households can spend optimally. A side benefit of this transformation is that the dependent variables in columns (4) and (5) are expressed in the same units as in columns (1) to (3), making coefficients comparable between them.

As anticipated, we find that the coefficient estimates on village average expenditure shown in columns (4) and (5) are less different from each other than the coefficient estimates obtained by assuming homothetic preferences. This confirms that assuming homothetic preferences overestimates the excess sensitivity of non-food expenditures to village shocks relative to food expenditures: part of this higher sensitivity is due to the fact that non-foods have a higher income elasticity. With this correction, we nonetheless continue to observe that non-foods respond more to aggregate village shocks than food expenditures.

Taken together, these findings indicate that food expenditures (which have low income elasticity) are more insulated from from village-level shocks than non-food expenditures (which have high income elasticity) *over and beyond* the smoothing in food consumption that would arise purely as a result of optimal reorganization of consumption towards food when individual total expenditures fall. The converse is true for non-food expenditures since they are, on average, more volatile than would be optimal based on their cross-section income elasticity.

Lastly, we examine the implications of risk pooling while allowing for heterogeneous risk and time preferences. As described in Section 3.3, this involves two steps. In the first step, we

---

<sup>28</sup>To illustrate with an example, let the food expenditures per capita of household  $i$  be Rs.1000. Further suppose that, based on the Engel curve, a household with a total expenditures of Rs.2500 has an average food share of 0.4 and thus spends Rs.1000 on food. It follows that a household that spends Rs.1000 on food must, on average, have a total expenditure of Rs.2500 in order to be on its Engel curve.

estimate *individual* households' risk and time preferences using 60 months of consumption data and assuming perfect risk pooling. The results from this estimation were discussed in Section 5. The second step relies on equation (12) to estimate the coefficients of cash-in-hand and earned income. These results are presented in Panel B of Table 2. As shown in Column (1) of Panel B for total expenditures, correcting for the bias from heterogeneous preferences does, as anticipated, reduce the magnitude of the coefficients on income and cash-in-hand. But the difference is minimal, suggesting that the bias is small. A similar conclusion emerges for columns (2) to (5): coefficient estimates in Panel B are similar to those reported in Panel A. Overall, this confirms our earlier interpretation of the findings under homogeneous preferences.

## 6.2 CRRA model

Next, we re-estimate all the regressions presented in Table 2 under the CRRA assumption. The main change is that the dependent variable and the regressors are now expressed in (first differences of) logs instead of levels. Other changes relate to the way risk and time preferences are estimated – a point already discussed in Section 5. The interpretation of the reported coefficients nonetheless remains the same.

The results are presented in Table 3. A number of observations are lost when taking logs due to zero or negative values in cash-in-hand or earned income. Negative values arise, for instance, when a household is a net borrower or when the imputed value of inputs (including family labor) allocated to crops and livestock exceeds crop and livestock revenues –e.g., due to crop or livestock losses. This loss of observations means we suffer some loss of power compared to Table 2. Standard errors are clustered at the village-level in all regressions.

The first thing to notice is how similar results in Panel A are to those reported in Table 2. The coefficient of average village expenditures in column (1) is significantly below 1 – indicating less than perfect risk pooling – but the magnitude of the difference is not large. Similarly, we find that year-to-year variations in household-level cash-in-hand have a statistically significant effect on consumption expenditures – but the magnitude of the effect is small: a doubling of cash-in-hand, for instance, translates, on average, into a 7.6% increase in household consumption per capita. Furthermore, similar to the CARA model estimates in Table 2, the coefficient on earned income is very small in magnitude and statistically indistinguishable from zero, indicating no excess sensitivity to household income. Given this, the cash-in-hand coefficient of 0.076 can be interpreted as the marginal propensity to consume (MPC) out of cash-in-hand and is precisely estimated with a standard error of 0.011. It is slightly higher than the MPC of 0.05 estimated by Blundell et al. (2008) for households in the US using PSID data, and is on the lower end of the mean MPC of 0.21 reported in a recent meta-analysis of 246 MPC estimates by Havranek and Sokolova (2020). The low MPC found in our study is far lower from the original estimate of 0.5 by Campbell and Mankiw (1989). It implies that a very low proportion of households in the VDSA villages live 'hand-to-mouth' and that most households engage in consumption smoothing strategies of mutual risk pooling.

Turning to Table 3 columns (2) to (5), we also find that food and non-food expenditures do not move with total expenditures in the same way across years within households as they do across households within years. If we take variation across households within years to compute average consumption shares (columns 2 and 3) or Engel curves (columns 4 and 5), we again find that correcting for non-homogeneity in consumption reduces the difference in estimated sensitivity to village shocks in average expenditures. We nonetheless find that, even with the correction for non-homogeneity, food and non-food still do not vary with average village consumption in a way consistent with a within-year utility-maximizing allocation of total expenditures: households increase their food consumption less – and their non-food consumption more – than what the increase in average village expenditures would predict.

We also do not find that non-food expenditures respond more to household income and wealth than food expenditures, ruling out the idea that there is less risk pooling in non-food consumption than in food consumption. Instead, the results suggest that, in a good year for the village as a whole, people increase their non-food expenditures more than implied by their cross-section income elasticity, thereby sheltering their average food consumption from aggregate shocks. Why this is the case is not a question that our estimation is designed to answer, but it could be due to social pressure to avoid spending on luxuries when other villagers are hungry.

Lastly, Panel B in Table 3 presents the test results of full risk pooling after accounting for heterogeneous risk and time preferences. In Column (1) we find a small reduction in the coefficients of cash-in-hand in Panel B compared to Panel A. This finding extends to food and non-food expenditures, whether we assume homothetic consumption preferences or allow for non-unitary income elasticities of consumption. Overall, the results remain similar for the CRRA model even after controlling for heterogeneous preferences.

Overall, the results under CRRA model indicate that accounting for non-unitary income elasticities are important when testing for risk pooling. The data indicate that food consumption is smoothed across years to a greater degree than non-food consumption. This could arise naturally from the fact the income elasticity of non-food expenditures is larger than for food expenditures. Results in Panel A and B demonstrate that this is not the case: the difference in the variation of food and non-food expenditure exceeds what can be predicted from the variation of expenditure shares with total consumption. Taken together, these findings indicate that food consumption is better insured than non-food consumption against village-level risk. These results suggest that village-level risk pooling aims to smooth year-to-year food consumption more than non-food consumption. This consistent with our initial hypothesis that risk pooling institutional arrangements put more emphasis on necessities such as food – and that this is achieved, at least in part, by preventing or discouraging non-food expenditures in bad years, while allowing their explosion in good years.

### 6.3 Village precautionary savings

So far we have shown that variation in household consumption expenditures is strongly correlated with changes in village consumption expenditures, and only mildly correlated with variation in household cash-in-hand. This suggests that the sharing of idiosyncratic income and asset risk is close to optimal.

We now examine whether our study villages are also close to optimal in terms of smoothing aggregate shocks over time. To this effect, we investigate whether the village pools precautionary savings to partially or fully insulate average village consumption from income and asset shocks. This is done using model 9 to estimate the response of total village consumption to village cash-in-hand. Efficient precautionary saving predicts a positive and significant response of consumption to cash-in-hand: when available cash falls, consumption is reduced, but less than one-for-one. For a large enough liquid wealth, consumption approaches certainty equivalent consumption whereby the response of consumption to an income shock is equal to the discount rate (e.g., Zeldes, 1989a). To illustrate, if the discount rate is 5%, certainty equivalent consumption changes by approximately Rs.5 in response to a Rs.100 temporary increase in income – and the coefficient of cash-in-hand in regression 9 should be around 0.05. At the same time, we should also observe a stable aggregate consumption over time and a large stock of liquid wealth.

Model 9 also includes village income as regressor in order to estimate the excess sensitivity of (aggregate) consumption to current income. Under the null of efficient savings at the village level, there should be no excess sensitivity to income since it is already included in cash-in-hand, and the coefficient of income should be zero.

Estimated results are reported in Table 4. Panel A presents the coefficient estimates for a model in first differences (CARA) and Panel B does the same for a model in log differences (CRRA). The coefficient of cash-in-hand Column (1) of Panel A shows a small but significant response of village consumption to village cash-in-hand. Based on this co-efficient estimate, a Rs.100 change in village cash-in-hand is associated with Rs. 7 change in total village consumption – a ratio (7%) equal to (or possibly lower than) what we would expect the discount rate to be for the study population. This at prima facie indicates a high level of consumption smoothing through precautionary saving – i.e., close to certainty equivalence. This is further supported by the small and insignificant coefficient on village earned income. Overall, these results suggest that villages optimally draw on liquid assets to smooth aggregate income shocks.

Columns (2) to (5) repeat the same procedure on food and non-food expenditures separately. As before, columns (2) and (3) assume linear Engel curves, that is, constant expenditure shares, while columns (4) and (5) allow expenditure shares to vary systematically with total expenditures. Results indicate that village food consumption does not respond to village cash-in-hand while village non-food expenditures respond significantly to it – and that this difference is somewhat stronger when correcting for non-homotheticity in consumption. This is consistent with

earlier results. Neither forms of consumption display excess income sensitivity.

The CRRA results shown in Panel B are broadly similar to those reported in Panel A for a model in first differences. We find a large response of village cash-in-hand to village consumption. Based on the estimates in Panel B Column (1), a 10% increase in village cash-in-hand translates on average into 2.7% increase in village consumption. This co-efficient is even larger for non-food. Based on the estimates in Panel B Column (3), a 10% increase in village cash-in-hand translates on average into 4.6% increase in village non-food consumption. The coefficient on non-food reduces a bit after correcting for non-homotheticity, as shown in Column (5).

## 7 Robustness and mechanisms

We have tested an expanded model of perfect risk sharing that allows for precautionary savings. We find strong evidence that household consumption co-varies with average village consumption, but responds significantly to variation in household cash-in-hand. These findings remain, albeit in a weaker form, when we also allow for household heterogeneity in risk aversion and discount rate. We also find that, correcting for differences in income elasticity, food consumption responds less to the village average than non-food consumption, suggesting that food consumption is smoothed preferentially. Our results also indicate that the village does not appear to smooth aggregate year-to-year shocks through precautionary savings.

In this Section, we subject these findings to additional analysis to assess their robustness to possible confounds and to examine possible mechanisms behind our results.

### 7.1 Sub-caste analysis

So far we have focused on risk pooling within entire villages. This implicitly assumes that the village is the correct unit of risk pooling. A number of studies have however suggested that, in the context of India, this assumption need not be appropriate: endogamous marriage groups called *Jati* or sub-castes may be a more likely social unit within which income sharing takes place, if only because of the strong bonds they create through marriage and family-based social events (e.g., Townsend, 1994; Mazzocco and Saini, 2012; Shrinivas and Fafchamps, 2018).

In this sub-section, we replicate our analysis at the level of sub-caste units within study villages.<sup>29</sup> The testing strategy is identical to that used in Section 3, except that village units are replaced by sub-castes. In particular, for the regressions with homogeneous preferences, we use average sub-caste expenditures instead of village expenditures. For the regressions with heterogeneous preferences, we first estimate risk and time preferences at the individual household

---

<sup>29</sup>To identify caste groups, we use the available VDSA data on *Jati*/sub-caste of each household. There are a total 251 unique village-specific sub-caste groups in the data. For the analysis, we drop the sub-castes counting only one household and focus on the 168 sub-caste groups with at least 2 households.

level, using monthly data as before. We then normalize household preference estimates to the level of their sub-caste.

Table 5 contains the results for both the CARA and the CRRA model, and is thus the equivalent of Tables 2 and 3. Contrary to expectations based on the literature, the sub-caste results are, if anything, less compelling than those at the village level. In particular, all estimated coefficients for the average sub-caste expenditure variable are smaller – and thus more different from one – than in the regressions using village expenditures. A similar pattern can be seen for cash-in-hand and earned income coefficients which tend to be slightly larger. Taken together, this suggests that risk pooling within sub-castes with each village is less strong than pooling across all households in the village.<sup>30</sup> This being said, the two sets of results are qualitatively similar in terms of patterns across regression models. This provides additional support for our earlier conclusions regarding the difference between food and non-food consumption smoothing and the stronger results using the model with a correction for non-homothetic preferences. It also confirms that the conclusions drawn in Panel A are not an artifact of ignoring the heterogeneity of risk and time preferences across households: conclusions are similar in sign and significance with those in Panel B. The results do, nonetheless, confirm that imposing homogeneity of preferences generates a bias: the coefficients on household income and liquid wealth get slightly smaller in Panel B – but not so much smaller as to change the take-away message of the analysis.

Table 6 replicates the precautionary savings regression analysis presented in Table 4, at the level of sub-castes. The main change is a drastic increase in the number observations, a change that increases power across the board. The pattern of results is similar to that discussed in Table 4, except that coefficient estimates are slightly smaller in magnitude. Conclusions are similar: sensitivity to cash-in-hand is very small – probably smaller than it should be given the relatively low level of liquid wealth in the data. Indeed, in a poor population such as the one we study, we would expect a stronger dependence of consumption to cash-in-hand in a precautionary savings optimum. This concern is somewhat confirmed by noting that, if anything, the unexplained variation in group-level consumption is larger in Table 6 than in Table 4.<sup>31</sup> We again note there is no excess sensitivity to income, as in Table 4.

Taken together, these results suggest that, at the time our data were collected, sub-castes were no longer the social and economic unit at which risk pooling was taking place.

## 7.2 Consumption shares and relative prices

We have found that different components of consumption are not smoothed in the same way. Could this be an artifact of changes in relative prices? Our concern is that covariate shocks may induce a correlation between idiosyncratic income and consumption that could lead to false

---

<sup>30</sup>See Appendix B where we show that sampling error in  $\bar{c}_{-i,t}$  is not large enough to explain the difference of  $\beta_2$  from 1.

<sup>31</sup>The  $R^2$  in column 1 of Table 4 is 0.084 in Panel A and 0.090 in Panel B, compared to 0.065 and 0.046 in Table 6, respectively.

rejection of full risk pooling, as shown by Ligon (2023). For instance, a negative co-variate shock such as a drought may lead to a fall in the price of food relative to non-food while a positive co-variate shock may lead to a rise in relative food price. Such co-variate shocks would result in a reorganization of consumption towards food in bad years and away from food in good years. This in turn would result in a smoother food consumption across time, and a more variable non-food consumption.

To account for changes in relative prices, we would ideally control for a village-specific index of the price of food relative to that of non-food. Unfortunately, the VSDA survey data does not include any information about non-food prices. To overcome this difficulty, we construct two separate variables: a village-specific food price index, which we can construct from the data; and a time-varying variable capturing the evolution of the relative price of food versus non-food at the national level. Given that non-food products include manufacturing goods and services that are traded at the national level, the combination of the two variables should proxy for relative price changes, one component of which varies across villages over years, and one component of which is largely determined at the national level.

The village food price index is constructed using the village price data and consumption quantity data, which are reported in the VSDA Monthly price schedule and Transaction schedule, respectively. This produces one food price index per village. For each village, this index is set to 100 for the year 2011 in which the Euler curves are estimated. We then include in the estimating equations this Village Food CPI as an additional regressor, the coefficient of which could be interpreted in terms of price elasticities. We expect food consumption to increase when this village-specific food price index rises, since food is widely believed to be relatively inelastic. The ratio of national non-food price index to the national aggregate price index is obtained from national CPI data. It measures the price change of non-food relative to food. We expect food consumption to fall with an increase in the relative price of non-food.

Estimation results are reported in Appendix Table A1 and A2 for the CARA and CRRA models respectively. Consistent with our prediction, they show that food consumption decreases with an increase in relative price of non-food. However, the standard errors on these coefficients are large, probably due to the lack of variation in price indices, as reported in the summary statistics Table 1. Nonetheless, these results validate that our main results on full risk pooling from the CARA and CRRA model in Tables 2 and 3 remain unchanged, suggesting that our findings are unlikely to be driven by changes in relative prices of food vs non-food.

If changes in relative prices cannot account for our findings, what can? One possibility is suggested by the observation that public or private social assistance to the poor often seeks to restrict the use that is made of given funds. Examples include food aid and food stamps, food and housing subsidies, shelter for the homeless, and subsidized education and health care. It has long been recognized that these forms of aid can constrain individual choices (Cunha, 2014) and that they reflect paternalistic preferences on the part of those who give – i.e., voters in the case of government programs and donors in the case of non-governmental organizations (Currie

and Gahvari, 2008). But evidence of such paternalistic altruism had not yet been documented in the case of informal risk sharing.

### 7.3 Wealth and assets

Arguably the most puzzling result from our analysis is the lack of responsiveness of village average consumption to aggregate village cash-in-hand. What could be behind this finding?

To throw some light on the issue, we start by considering the different assets that compose village wealth. As noted earlier, in an optimal village risk pooling equilibrium, village wealth  $w_t$  can compensate a shortfall in village income  $y_t$  by dissaving  $d_t \equiv w_{t+1} - w_t$ . Given the village's accounting identity  $c_t = y_t + w_t - w_{t+1}$  estimating a regression of the form:

$$d_t = \alpha_0 + \alpha_1 y_t + \alpha_2 c_t + u_t \quad (13)$$

should yield an estimated coefficient  $\hat{\alpha}_1$  identically equal to  $-1$ . In practice, there is measurement error in the data, especially in income, which biases  $\hat{\alpha}_1$  towards 0.

In our study, household wealth has five main components: land  $a_t$ ; financial assets  $f_t$ ; jewelry  $g_t$ ; consumer durables  $b_t$ ; and livestock  $l_t$ . In practice, there are hardly any land sales and purchases in our data, which rules out the use of land as buffer stock against income shocks. For this reason land is omitted from village liquid wealth and we define  $w_t \equiv f_t + g_t + b_t + l_t$ . Of the four components of  $w_t$ , we expect financial assets and jewelry to be the most liquid and consumer durables to be the least liquid due to the gap between buying and resale price. It is also conceivable that livestock is an asset that farmers are reluctant to part with (e.g., Fafchamps et al., 1998). From these observations, we expect more adjustment in  $f_t$  and  $g_t$  – and less in  $b_t$  and  $l_t$  – when the village dissaves. This can be tested by estimating regression:

$$d_t^j = \alpha_0^j + \alpha_1^j y_t + \alpha_2^j c_t + u_t^j \quad (14)$$

for  $j \in \{f, g, b, l\}$ . If all four assets are equally liquid, we should find  $\hat{\alpha}_1^j = s_j \hat{\alpha}_1$  where  $s_j$  is the share of asset  $j$  in  $w_t$ . But if, as we expect, the most liquid assets are used more to smooth income shocks, their estimated  $\hat{\alpha}_1^j$  should be larger. We present in Appendix D a more detailed treatment of this issue.

We estimate equation (14) for aggregate village wealth and for different components of liquid wealth. Results are reported in Appendix Table A4. From the estimated coefficient of earned income, we find that most of the variation in village assets associated with shortfalls in income is borne by financial assets and by jewelry (i.e., gold and silver). Livestock hardly responds and consumer durables have the wrong sign. This confirms that, as predicted, it is the most liquid assets that adjust to income shocks. This being said, the earned income coefficients are far from summing to  $-1$ , indicating that changes in assets do not fully bridge the gap between village income and consumption.



These results suggest the existence of large changes in average village consumption that do not closely follow reported changes in village income and liquid wealth. To investigate the magnitude of this issue, we calculate the unexplained gap  $G_{vt}$  between village-level consumption and cash-in-hand for village  $v$  in year  $t$ , which we calculate as:

$$G_{vt} \equiv C_{vt} - Y_{vt} - L_{v,t-1} + L_{vt}$$

where  $C_{vt}$  is consumption in year  $t$ ,  $Y_{vt}$  is income in year  $t$ , and  $L_{vt}$  is the liquid wealth at the end of year  $t$ . In principle,  $G_{vt}$  should be zero or close to zero in a precautionary savings model where consumption is financed out of income and year-to-year changes in liquid wealth. We indeed find that the mean of  $G_{vt}$  is close to zero – see Appendix Figure A2. But its variance is about four times the variance of  $C_{vt}$ .<sup>32</sup>

This indicates the presence of variation in village consumption that is driven by sources other than variation in earned income  $Y_{vt}$  and liquid wealth and indebtedness  $L_{vt} - L_{v,t-1}$ . What are these other sources of fluctuations is unclear. Measurement error undoubtedly plays a part – income is notoriously difficult to measure in rural economies and consumption expenditure data is subject to recall bias. But village averaging should, in principle, smooth out some of these measurement errors and diminish their impact on the estimation. Other possibilities include fluctuations in remittance flows from migrant family members, variation in transfers to and from relatives in other villages, and unreported financial transactions – such as unrecorded debt forgiveness. This calls for further research.

## 7.4 Unanswered questions

What have we learned from this exercise that we did not already know? We already knew that village economies in India and elsewhere enjoy a great deal of consumption smoothing with respect to idiosyncratic variation in income (e.g., Townsend, 1994). We also knew that this consumption smoothing is not consistent with perfect within-year risk sharing, even after allowing for heterogeneity in risk preferences (e.g., Schulhofer-Wohl, 2011; Chiappori et al., 2014; Mazzocco and Saini, 2018). On both counts, our results replicate these earlier findings, albeit with more recent data and after also allowing for heterogeneity in time preferences.

It has been suggested that such findings may be explained by limited commitment in informal risk sharing arrangements: although villagers may agree to sharing risk *ex ante*, they may renege *ex post* if they are asked to share more than the expected utility of what they expect to receive in return later (e.g., Coate and Ravallion, 1993). Some partial evidence has been provided for and against this interpretation (e.g., Ligon et al., 2002; Fafchamps and Lund, 2003; Weerdt and Fafchamps, 2011), without fully closing the debate.

---

<sup>32</sup>The variance of village consumption  $C_{vt}$  is 116,672 while the variance of the gap  $G_{vt}$  is 502,025. A test that the ratio of the two variances is equal to 1 is rejected with a  $p$ -value= 0.000.

What has largely been lacking in this debate, however, is a bridge between risk sharing and asset accumulation. Indeed households can achieve a great deal of consumption smoothing via precautionary saving (e.g., Zeldes, 1989b; Deaton, 1991). It is also well known that consumption smoothing through individual asset accumulation can, in practice, achieve a degree of consumption smoothing that is extraordinarily similar to risk pooling (e.g., Deaton, 1992).<sup>33</sup> This is where the main contribution of our paper comes in. If consumption smoothing is achieved via individual savings, we should observe a sensitivity of consumption with respect to cash-in-hand, that is, to the sum of current income and liquid wealth, especially among a poor population such as the one we study. This is not what we find: while household consumption does show some sensitivity to cash-in-hand, this sensitivity is low – much lower than what we would expect at very low levels of savings (e.g., Zeldes, 1989b). Plus we see strong dependence on village average consumption, which should not arise if people are saving individually and are subject to idiosyncratic variation in income.

The most novel aspect of our approach is that we consider precautionary saving at the village level. Indeed, in a perfect village risk pooling equilibrium, the village would accumulate liquid assets to smooth consumption from fluctuation in village income over time – e.g., in a rural economy, against weather- and pest-induced variation in crop production. Perfect risk pooling would thus make individual consumption vary with average village consumption, as the standard test assumes – and average village consumption would vary with aggregate village cash-in-hand. We test this prediction and find that average village consumption responds only weakly to village cash-in-hand. We also find that average village consumption is less variable than average village income or cash-in-hand. This raises two main unanswered questions: how do villages smooth their average consumption if it is not through precautionary savings; and why don't villagers accumulate savings to protect against village-wide shocks?

To address the first question, we look at two possible sources of insurance from outside the village on which we have data: gifts and transfers from outside the village – through family networks (e.g., Rosenzweig and Stark, 1989) or via remittances from migrants (e.g., Gubert, 2002; Azam and Gubert, 2006; Munshi and Rosenzweig, 2016) and loans from formal and informal sources (e.g., Udry, 1990; Fafchamps and Lund, 2003). By a simple aggregate accounting reasoning, if these flows of fund finance the deviation of village consumption from village cash-in-hand, we should observe a positive coefficient when regressing average village consumption on inflows of funds from gifts and loans over the same time period. Results presented in Table A3 reject this hypothesis: in both cases, the estimated coefficient is small in magnitude and not statistically significant. From this evidence we conclude that, based on the extensively detailed data at our disposal, we are unable to account for the discrepancy between village consumption and village cash-in-hand from all the sources we observe in the data. While this phenomenon has already been observed at the individual household level (e.g., Lim and Townsend, 1998), it has never, to our knowledge, been discussed at the village level. What accounts for this discrepancy thus remains a mystery. Possible explanations include measurement error in income and assets (e.g.,

---

<sup>33</sup>See Sargent (1987), Chapter 3, for a model of risk pooling through individual accumulation of liquid wealth.

Carletto et al., 2022); and hidden income and assets (e.g., Kinnan, 2022). More research is needed on this topic.

The second unanswered question relates to the lack of precautionary saving, either at the individual or collective level, to deal with village-level income shocks that, by definition, cannot be smoothed via within-village redistribution. Since risk pooling among villagers is unable to protect households against such shocks, one would expect villagers to individually self-insure by accumulating liquid assets at the individual or household level. Yet they do not. Why?

One possible explanation is offered by the limited commitment model of risk sharing. In this model, individuals who can self-insure with their accumulated wealth are, as a result, less in need of mutual insurance. This, in turn, implies that to be convinced to assist others financially, they must receive a net positive financial return (e.g., Fafchamps, 1999). This process, in turn, creates a strong divergence in wealth levels over time (e.g., Fafchamps, 2003): the rich become richer by insuring others in difficulty in exchange for a stream of transfers and favors, a situation that has been discussed in the literature on patron-client relationships (e.g., Platteau, 1995b,a). The danger, however, is that individuals with sufficient wealth leave the mutual insurance arrangement (e.g., Scott, 1977). To prevent this from happening, informal risk pooling arrangements tend to nurture redistributive norms that tax the accumulation of wealth (e.g., Platteau et al., 1998; Platteau, 2014; Fromell et al., 2021). This could explain the limited accumulation of liquid wealth by the villagers in our sample and, thus, the scant evidence of the use of savings as a buffer against shocks, either at the individual or village level, in our study area as in others (e.g., Fafchamps et al., 1998; Kazianga and Udry, 2006).

Another possible explanation relates to the empirical observation that many people are unable to accumulate savings out of windfall gains, either because of overconfidence in the future (Pfäuti et al., 2024) or because of social pressure (Goldberg, 2017). These two mechanisms are related to each other through the often-made observation that, in poor communities, there is always someone whose marginal utility of cash is high. Given this, it may be ex ante optimal for a community to redistribute any excess cash to more needy households before that cash gets wasted in futile expenditures and fraught business investment. Redistribution is then a way to avoid wasteful expenditures, an interpretation that is consistent with our finding regarding the contrast between the smoothing of necessities (i.e., food) and non-necessities (i.e., non-food) depending on the average consumption of the village: when the village is poor, non-food consumption deviates from what would be individual optimal, while food consumption is smoothed more than total village consumption.

## 8 Conclusion

In this paper we have revisited the seminal and influential literature on risk pooling in village economies. We make a number of methodological contributions. First, we show how precaution-

ary savings and non-homothetic preferences can be incorporated in standard tests of risk pooling in a simple and easy-to-implement way. Second, we expand the recent work of Schulhofer-Wohl (2011), Mazzocco and Saini (2012) and Chiappori et al. (2014) on heterogeneity in risk preferences to also allow for individual heterogeneity in time preferences. Third, we integrate all these improvements into a single, elegant testing strategy using more recent data.

The usefulness of our proposed approach is illustrated in a geographical and cultural context similar to that studied by Townsend (1994). While our results on aggregate risk pooling mirror those of the existing literature, our approach generates a number of new insights. First, we show that risk pooling creates a distortion in consumption such that food consumption is better protected from aggregate village shocks than non-food consumption, even after accounting for non-unitary income elasticities. This finding echoes many social welfare policies in developed economies, which similarly prioritize specific types of consumption by the poor (e.g., food stamps, housing and shelter, health care, primary and secondary education, public transport) while taxing luxuries.

Second, we find that, contrary to findings based on earlier data, risk pooling in Indian villages no longer appears to occur primarily at the sub-caste level rather than at the village level; household consumption better tracks average village expenditures than average sub-caste expenditures within villages.

Third, we find evidence that household consumption tracks aggregate village cash-in-hand, suggesting some form of precautionary savings at the village level. But there is considerable excess sensitivity to aggregate income, indicating a lack of full asset integration. We also find a large unexplained gap between the variation in measured consumption expenditures and cash-in-hand at the aggregate village level. Possible explanations include fluctuations in remittance flows from migrant family members, variation in transfers to and from relatives in other villages, and borrowing and lending to deal with shocks. We examine these possibilities and find little or no evidence that the consumption-income gap at the village level is filled by these flows of fund. Other possible explanations for this unexplained gap include mis-measurement in household expenditures, income, and wealth. But village averaging should, in principle, smooth out some of these measurement errors. Another unexplored possibility is unrecorded debt forgiveness, i.e., loans that are never paid. More research is needed on these possible sources of insurance against village-level collective shocks.

As final observation, we should make clear that the presence of risk pooling in a village does not, by itself, constitute evidence that villagers explicitly share risk in the form of mutual insurance arrangements or self-help groups, or via contingent transfers and gift exchange. Within village risk pooling can indeed be achieved by other means, such as informal peer-to-peer loans and credit from MFIs and shopkeepers. But we find no evidence that consumption smoothing is mostly achieved via individual precautionary saving.

# Figures and Tables

## Engel curves

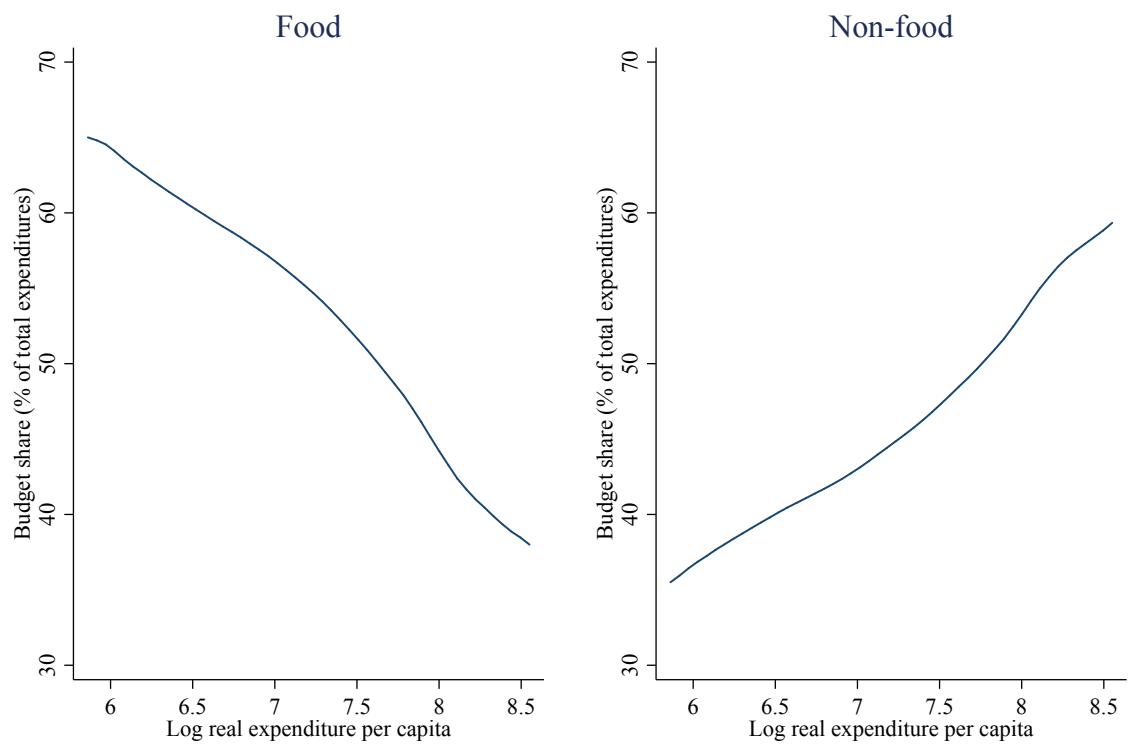


Figure 1: Engel Curves for food and non-food

## Frequency distribution of relative risk tolerance estimates




Figure 2: Normalized risk tolerance  $\beta$

Each panel of this Figure depicts the frequency distribution of  $\hat{\beta}_i$ , the estimate of household-specific risk tolerance obtained from regression (10). In each case, the sample mean of  $\hat{\beta}_i$  is set to 1 within each village. This means that  $\hat{\beta}_i$  measures the risk tolerance of household  $i$  relative to other households in the same village.

### Frequency distribution of relative time preference estimates




Figure 3: Normalized time preference parameter  $\rho$

Each panel of this Figure depicts the frequency distribution of  $\hat{\rho}_i$ , the estimate of household-specific discount factor obtained from regression (10). As for risk tolerance, the sample mean of  $\hat{\rho}_i$  is set to 1 within each village – which that  $\hat{\rho}_i$  measures the discount factor of household  $i$  relative to other households in the same village.

## Histogram of the frequency count of villages by their proportion of households below the equitable welfare weight




Figure 4: Consumption Inequality

Each panel of this Figure depicts the histogram of village-average proportion of households with welfare weight  $\hat{\eta}_i$  less than the average welfare weight  $1/N_v$  for the 30 villages, one from CARA and the other from CRRA. The household-specific welfare weight estimate  $\hat{\eta}_i$  is obtained from regression (10). In each case, the sum of  $\hat{\eta}_i$  is set to 1 within each village. Based on CARA estimates, the overall mean is 0.88, implying that 88% of households have welfare weights less than equitable share. Similarly, the overall mean for CRRA estimates is 0.89.



Table 1: Summary Statistics

	Mean	SD
<b><i>Consumption</i></b>		
Total expenditure	15937	10884
Food expenditure	8321	3776
Non-food expenditure	7615	8072
<b><i>Income</i></b>		
Crop and Livestock income	2238	15204
Wages income	15598	20053
Earned income	19774	25100
Unearned Income	2860	7235
<b><i>Assets</i></b>		
Wealth	88637	117235
Liquid wealth	25384	44424
Cash-in-Hand	48018	58492
<b><i>Household characteristics</i></b>		
Household size	4.8	2.3
Adult-eq weight	4.1	1.8
Education of head (in years)	5.1	4.7
Age of head	50	13
CPI food village	112	14
Non-food CPI / Agg CPI	0.98	0.02
Households	1296	
Villages	30	
Observations	6589	

*Notes:* This table reports descriptive statistics - mean and standard deviation - for household consumption, income, assets and demographic characteristics. Consumption, Income and Asset variables represent annual values, adjusted to 2010 rupees per adult-equivalent. Total expenditures is the sum of food and non-food expenditures, and earned income is the sum of crop and livestock income and wage income. Wealth is sum of net credit position (saving-borrowing+lending) and total assets (liquid assets+capital assets); Liquid wealth is the sum of net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Similarly, cash-in-hand is constructed as the sum of liquid wealth and eared income. For all values, top and bottom 1% are winsorized to account for measurement error.

Table 2: First differences in levels - CARA model

	Total expenditure (1)	Homothetic preferences		Non-homothetic	
		Food (2)	Non-food (3)	Food (4)	Non-food (5)
<b>Panel A: Homogeneous preferences</b>					
Average village expenditures	0.769*** (0.038)	0.544*** (0.091)	1.014*** (0.059)	0.625*** (0.114)	0.758*** (0.049)
Cash-in-hand	0.021*** (0.007)	0.012*** (0.004)	0.032** (0.012)	0.016*** (0.005)	0.020** (0.007)
Earned Income	0.013 (0.011)	0.015* (0.008)	0.010 (0.016)	0.015 (0.011)	0.005 (0.010)
Observations	5256	5256	5256	5256	5256
<b>Panel B: Heterogeneous preferences</b>					
Cash-in-hand	0.019** (0.007)	0.007* (0.004)	0.031** (0.012)	0.012** (0.005)	0.017** (0.007)
Earned Income	0.010 (0.010)	0.012* (0.007)	0.007 (0.015)	0.013 (0.009)	0.002 (0.009)
Observations	5256	5256	5256	5256	5256

*Notes:* This table reports the results from household panel-pooled estimation of first differences in levels, based on CARA model. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. The unit of observation is a household-year. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. Panel A reports the test results of full risk pooling under homogeneous preferences and Panel B reports the test results of full risk pooling after accounting for heterogeneity in risk and time preferences. Village expenditures represents the village leave-out-mean. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures that correspond to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. In Panel B, the dependent variable is transformed to account for heterogeneity in risk and time preferences – see text for details. The top and bottom 1% of all values are trimmed to eliminate outliers and large measurement errors. All standard errors are clustered at the village level. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.

Table 3: First differences in logs - CRRA model

	Total expenditure (1)	Homothetic preferences		Non-homothetic	
		Food (2)	Non-food (3)	Food (4)	Non-food (5)
<b>Panel A: Homogeneous preferences</b>					
Average village expenditures	0.895*** (0.020)	0.608*** (0.064)	1.394*** (0.134)	0.652*** (0.063)	0.793*** (0.074)
Cash in hand	0.076*** (0.011)	0.056*** (0.012)	0.090*** (0.023)	0.063*** (0.013)	0.057*** (0.013)
Earned Income	-0.001 (0.007)	0.011 (0.009)	-0.018 (0.012)	0.013 (0.009)	-0.009 (0.008)
Observations	4491	4491	4491	4491	4491
<b>Panel B: Heterogeneous preferences</b>					
Cash in hand	0.069*** (0.011)	0.043*** (0.012)	0.096*** (0.023)	0.050*** (0.013)	0.048*** (0.014)
Earned Income	0.000 (0.006)	0.012 (0.009)	-0.017 (0.012)	0.014 (0.009)	-0.009 (0.008)
Observations	4491	4491	4491	4491	4491

*Notes:* This Table reports the results from a household panel pooled estimation of first differences in logs, based on the CRRA model. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. The unit of observation is a household-year. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. Panel A reports the test results of full risk pooling under homogeneous preferences and Panel B reports the test results of full risk pooling after accounting for heterogeneity in risk and time preferences. Village expenditures represents the village leave-out-mean. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures that correspond to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. In Panel B, the dependent variable is transformed to account for heterogeneity in risk and time preferences – see text for details. The top and bottom 1% of all values are trimmed to eliminate outliers and large measurement errors. All standard errors are clustered at the village level. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.

Table 4: Village precautionary savings

	Total expenditure (1)	Homothetic pref		Non-homothetic pref	
		Food (2)	Non-food (3)	Food (4)	Non-food (5)
<b>Panel A: First differences</b>					
Cash-in-Hand	0.070** (0.029)	0.008 (0.013)	0.062*** (0.021)	0.015 (0.029)	0.093*** (0.034)
Earned Income	-0.006 (0.047)	0.018 (0.022)	-0.024 (0.035)	0.043 (0.047)	-0.038 (0.056)
R-squared	0.084	0.028	0.096	0.028	0.081
Observations	120	120	120	120	120
<b>Panel B: Log differences</b>					
Cash-in-Hand	0.270*** (0.085)	0.117 (0.073)	0.467*** (0.155)	0.122 (0.080)	0.313** (0.128)
Earned Income	-0.070 (0.052)	-0.007 (0.045)	-0.162 (0.095)	-0.003 (0.049)	-0.116 (0.079)
R-squared	0.090	0.035	0.074	0.035	0.049
Observations	120	120	120	120	120

*Notes:* This table reports the test of precautionary saving at the village level. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. The unit of observation is a village-year and all variables represent village averages. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. Panel A reports the test results for first differences in levels and Panel B reports the results for first difference in logs. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures corresponding to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. Standard errors are reported in parenthesis. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.

Table 5: Risk pooling within sub-castes

	CARA Model : First difference in levels					CRRA model : First difference in log units				
	Homothetic preferences		Non-homothetic			Homothetic preferences		Non-homothetic		
	Total exp (1)	Food (2)	Non-food (3)	Food (4)	Non-food (5)	Total exp (6)	Food (7)	Non-food (8)	Food (9)	Non-food (10)
<b>Panel A: Homogeneous preferences</b>										
Average Caste Expenditure	0.545*** (0.047)	0.377*** (0.051)	0.729*** (0.067)	0.450*** (0.064)	0.527*** (0.046)	0.755*** (0.038)	0.498*** (0.053)	1.169*** (0.083)	0.536*** (0.052)	0.676*** (0.047)
Cash-in-hand	0.022*** (0.007)	0.011*** (0.004)	0.033*** (0.011)	0.016*** (0.005)	0.021*** (0.007)	0.081*** (0.010)	0.057*** (0.010)	0.103*** (0.020)	0.064*** (0.011)	0.064*** (0.011)
Earned Income	0.012 (0.011)	0.015** (0.007)	0.009 (0.017)	0.014 (0.009)	0.004 (0.010)	-0.001 (0.007)	0.011 (0.007)	-0.020* (0.011)	0.013 (0.008)	-0.011 (0.007)
Observations	4543	4543	4543	4543	4543	3862	3862	3862	3862	3862
<b>Panel B: Heterogeneous preferences</b>										
Cash-in-hand	0.021*** (0.007)	0.010*** (0.004)	0.034*** (0.011)	0.015*** (0.005)	0.020*** (0.007)	0.076*** (0.010)	0.047*** (0.010)	0.108*** (0.020)	0.053*** (0.011)	0.055*** (0.025)
Earned Income	0.010 (0.010)	0.012* (0.006)	0.009 (0.017)	0.012 (0.009)	0.002 (0.010)	-0.004 (0.006)	0.008 (0.007)	-0.021* (0.014)	0.009 (0.008)	-0.009 (0.009)
Observations	4543	4543	4543	4543	4543	3862	3862	3862	3862	3862

This table reports the results from household panel-pooled estimation of first differences in levels and logs, based on the CARA and CRRA models, respectively. The left-hand panel is similar to Table 2, but estimated at the Jati or sub-caste level. The right-hand panel is similar to Table 3, but estimated at the Jati or sub-caste level. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. The unit of observation is a household-year. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. Panel A reports the test results of full risk pooling under homogeneous preferences and Panel B reports the test results of full risk pooling after accounting for heterogeneity in risk and time preferences. Village expenditures represent the village leave-out-mean. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures that correspond to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. In Panel B, the dependent variable is transformed to account for heterogeneity in risk and time preferences – see text for details. The top and bottom 1% of all values are trimmed to eliminate outliers and large measurement errors. All standard errors are clustered at the village level. \* p<0.10 \*\* p<0.05 \*\*\* p<0.01.

Table 6: Caste-level precautionary savings

	Total expenditure (1)	Homothetic pref		Non-homothetic pref	
		Food (2)	Non-food (3)	Food (4)	Non-food (5)
<b>Panel A: First differences</b>					
Cash-in-Hand	0.039*** (0.013)	0.009* (0.005)	0.031*** (0.010)	0.019* (0.012)	0.044*** (0.014)
Earned Income	0.022 (0.021)	0.007 (0.009)	0.015 (0.016)	0.018 (0.019)	0.016 (0.024)
R-squared	0.065	0.023	0.063	0.025	0.054
Observations	360	360	360	360	360
<b>Panel B: Log differences</b>					
Cash-in-Hand	0.122*** (0.033)	0.099*** (0.026)	0.150** (0.061)	0.109*** (0.028)	0.094** (0.047)
Earned Income	-0.008 (0.019)	-0.017 (0.015)	-0.007 (0.035)	-0.018 (0.016)	-0.005 (0.027)
R-squared	0.046	0.044	0.023	0.045	0.015
Observations	354	354	354	354	354

*Notes:* This table reports the test of precautionary saving at the Jati or sub-caste level. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. The unit of observation is a sub-caste-year and all variables represent sub-caste averages. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. Panel A reports the test results for first differences in levels and Panel B reports the results for first difference in logs. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures corresponding to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. Standard errors are reported in parenthesis. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.

## References

- Aguiar, M. A., M. Bils, and C. Boar (2020, January). Who are the hand-to-mouth? Working Paper 26643, National Bureau of Economic Research.
- Almås, I., T. K. Beatty, and T. F. Crossley (2018). Lost in Translation: What do Engel Curves Tell us about the Cost of Living? CESifo Working Paper Series 6886, CESifo.
- Altonji, J. G., F. Hayashi, and L. J. Kotlikoff (1992). Is the extended family altruistically linked? direct tests using micro data. *American Economic Review* 82(5), 1177 – 1198.
- Ambrus, A. and M. Elliott (2020, 11). Investments in social ties, risk sharing, and inequality. *The Review of Economic Studies* 88(4), 1624–1664.
- Ambrus, A., W. Gao, and P. MilÅ;n (2021, 12). Informal Risk Sharing with Local Information. *The Review of Economic Studies* 89(5), 2329–2380.
- Ambrus, A., M. Mobius, and A. Szeidl (2014, January). Consumption risk-sharing in social networks. *American Economic Review* 104(1), 149–82.
- Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2020, March). Measuring welfare and inequality with incomplete price information. Working Paper 26890, National Bureau of Economic Research.
- Attanasio, O., A. Kochar, A. Mahajan, and V. Surendra (2023, May). Risk sharing, commitment constraints and self help groups. Working Paper 31245, National Bureau of Economic Research.
- Azam, J.-P. and F. Gubert (2006, 12). Migrants’ Remittances and the Household in Africa: A Review of Evidence. *Journal of African Economies* 15, 426–462.
- Blundell, R., L. Pistaferri, and I. Preston (2008, December). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Campbell, J. Y. and N. G. Mankiw (1989, April). Consumption, income, and interest rates: Reinterpreting the time series evidence. Working Paper 2924, National Bureau of Economic Research.
- Carletto, G., M. Tiberti, and A. Zezza (2022). Measure for measure: Comparing survey based estimates of income and consumption for rural households. *World Bank Research Observer* 37(1), 1–38.
- Chandrasekhar, A. G., R. Townsend, and J. P. Xandri (2018, March). Financial centrality and liquidity provision. Working Paper 24406, National Bureau of Economic Research.
- Chiappori, P.-A., K. Samphantharak, S. Schulhofer-Wohl, and R. M. Townsend (2014). Heterogeneity and risk sharing in village economies. *Quantitative Economics* 5(1), 1–27.
- Coate, S. and M. Ravallion (1993). Reciprocity without commitment: Characterization and performance of informal insurance arrangements. *Journal of Development Economics* 40(1), 1 – 24.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99(5), 957–976.
- Cunha, J. M. (2014, April). Testing paternalism: Cash versus in-kind transfers. *American Economic Journal: Applied Economics* 6(2), 195–230.

- Currie, J. and F. Gahvari (2008, June). Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature* 46(2), 333–83.
- De Weerd, J. and S. Dercon (2006). Risk-sharing networks and insurance against illness. *Journal of Development Economics* 81(2), 337 – 356.
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica* 59(5), 1221–1248.
- Deaton, A. (1992, 03). Saving and Income Smoothing in Cote d’Ivoire. *Journal of African Economies* 1(1), 1–24.
- Dercon, S., J. De Weerd, T. Bold, and A. Pankhurst (2006). Group-based funeral insurance in ethiopia and tanzania. *World Development* 34(4), 685–703.
- Dercon, S. and P. Krishnan (2000). In sickness and in health: Risk sharing within households in rural ethiopia. *Journal of Political Economy* 108(4), 688 – 727.
- Dubois, P. (2001). Consumption insurance with heterogeneous preferences. Can sharecropping help complete markets ? Economics Working Paper Archive (Toulouse) 25, French Institute for Agronomy Research (INRA), Economics Laboratory in Toulouse (ESR Toulouse).
- Dunbar, G. R., A. Lewbel, and K. Pendakur (2013, February). Children’s resources in collective households: Identification, estimation, and an application to child poverty in malawi. *American Economic Review* 103(1), 438–71.
- Escanciano, J. C., S. Hoderlein, A. Lewbel, O. Linton, and S. Srisuma (2021, October). Nonparametric Euler Equation Identification And Estimation. *Econometric Theory* 37(5), 851–891.
- Fafchamps, M. (1992). Solidarity networks in preindustrial societies: Rational peasants with a moral economy. *Economic Development and Cultural Change* 41(1), 147–174.
- Fafchamps, M. (1999). Risk sharing and quasi-credit. *The Journal of International Trade & Economic Development* 8(3), 257–278.
- Fafchamps, M. (2003). *Rural Poverty, Risk, and Development*. Elgar Press.
- Fafchamps, M. and S. Lund (2003). Risk-sharing networks in rural philippines. *Journal of Development Economics* 71(2), 261–287.
- Fafchamps, M., C. Udry, and K. Czukas (1998). Drought and saving in west africa: are livestock a buffer stock? *Journal of Development Economics* 55(2), 273–305.
- Fromell, H., D. Nosenzo, T. Owens, and F. Tufano (2021). One size does not fit all: Plurality of social norms and saving behavior in kenya. *Journal of Economic Behavior & Organization* 192, 73–91.
- Gertler, P. and J. Gruber (2002, March). Insuring consumption against illness. *American Economic Review* 92(1), 51–70.
- Goldberg, J. (2017). The effect of social pressure on expenditures in malawi. *Journal of Economic Behavior and Organization* 143, 173–185.
- Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003). Optimal indirect and capital taxation. *The Review of Economic Studies* 70(3), 569–587.
- Gubert, F. (2002). Do migrants insure those who stay behind? evidence from the kayes area (western mali). *Oxford Development Studies* 30(3), 267–287.
- Hastings, J. and J. M. Shapiro (2018, December). How are snap benefits spent? evidence from a retail panel. *American Economic Review* 108(12), 3493–3540.



- Havranek, T. and A. Sokolova (2020). Do consumers really follow a rule of thumb? three thousand estimates from 144 studies say "probably not". *Review of Economic Dynamics* 35, 97–122.
- Kazianga, H. and C. Udry (2006). Consumption smoothing? livestock, insurance and drought in rural burkina faso. *Journal of Development Economics* 79(2), 413–446. Special Issue in honor of Pranab Bardhan.
- Kimball, M. S. (1988). Farmers' cooperatives as behavior toward risk. *The American Economic Review* 78(1), 224–232.
- Kinnan, C. (2022). Distinguishing barriers to insurance in thai villages. *Journal of Human Resources* 57(1), 44–78.
- Kinnan, C., K. Samphantharak, R. Townsend, and D. Vera-Cossio (2024, January). Propagation and insurance in village networks. *American Economic Review* 114(1), 252–84.
- Kochar, A. (1999). Smoothing consumption by smoothing income: Hours-of-work responses to idiosyncratic agricultural shocks in rural India. *Review of Economics and Statistics* 81(1), 50–61.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies* 63(4), 595–609.
- Laczó, S. (2015). Risk sharing with limited commitment and preferences heterogeneity: Structural estimation and testing. *Journal of the European Economic Association* 13(2), 265–292.
- Lewis, D. J., D. Melcangi, and L. Pilossoph (2019, November). Latent Heterogeneity in the Marginal Propensity to Consume. Staff Reports 902, Federal Reserve Bank of New York.
- Ligon, E. (1998). Risk sharing and information in village economies. *The Review of Economic Studies* 65(4), 847–864.
- Ligon, E. (2020, February). Estimating household welfare from disaggregate expenditures. Working Paper 109, Center for Effective Global Action.
- Ligon, E. (2023, July). Risk sharing tests and covariate shocks. Working paper, UC Berkeley: Department of Agricultural and Resource Economics.
- Ligon, E., J. P. Thomas, and T. Worrall (2002, 01). Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *The Review of Economic Studies* 69(1), 209–244.
- Lim, Y. and R. M. Townsend (1998). General equilibrium models of financial systems: Theory and measurement in village economies. *Review of Economic Dynamics* 1(1), 59–118.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99(5), 928 – 956.
- Mazzocco, M. and S. Saini (2012). Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review* 102(1), 428–68.
- Mazzocco, M. and S. Saini (2018, October). Testing efficient risk sharing with heterogeneous risk preferences: Reply. *American Economic Review* 108(10), 3114–15.
- Munshi, K. and M. Rosenzweig (2016, January). Networks and misallocation: Insurance, migration, and the rural-urban wage gap. *American Economic Review* 106(1), 46–98.
- Murphy, K. M. and R. H. Topel (1985). Estimation and inference in two-step econometric models. *Journal of Business and Economic Statistics* 3(4), 370–379.

- Pan, L. (2009). Risk Pooling through Transfers in Rural Ethiopia. *Economic Development and Cultural Change* 57(4), 809–35.
- Paxson, C. (1993, February). Consumption and Income Seasonality in Thailand. *Journal of Political Economy* 101(1), 39–72.
- Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in thailand. *The American Economic Review* 82(1), 15–33.
- Pfäuti, O., F. Seyrich, and J. Zinman (2024, April). Bad luck or bad decisions? macroeconomic implications of persistent heterogeneity in cognitive skills and overconfidence. Working Paper 32305, National Bureau of Economic Research.
- Platteau, J.-P. (1995a). A framework for the analysis of evolving patron-client ties in agrarian economies. *World Development* 23(5), 767–786.
- Platteau, J.-P. (1995b). An indian model of aristocratic patronage. *Oxford Economic Papers* 47(4), 636–662.
- Platteau, J.-P. (2014). *Redistributive Pressures in Sub-Saharan Africa: Causes, Consequences, and Coping Strategies*, pp. 153–207. Cambridge University Press.
- Platteau, J.-P., Y. Hayami, and P. Dasgupta (1998). *Resource Endowments and Agricultural Development: Africa versus Asia*, pp. 357–412. London: Palgrave Macmillan UK.
- Ravallion, M. and S. Chaudhuri (1997). Risk and insurance in village india: Comment. *Econometrica* 65(1), 171 – 184.
- Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica* 53(1), 69–76.
- Rosenzweig, M. R. and H. P. Binswanger (1993). Wealth, weather risk and the composition and profitability of agricultural investments. *Economic Journal* 103(416), 56 – 78.
- Rosenzweig, M. R. and O. Stark (1989). Consumption smoothing, migration, and marriage: Evidence from rural india. *Journal of Political Economy* 97(4), 905 – 926.
- Rosenzweig, M. R. and K. I. Wolpin (1993). Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investment in bullocks in india. *Journal of Political Economy* 101(2), 223 – 244.
- Sargent, T. J. (1987). *Dynamic Macroeconomic Theory*. Harvard University Press.
- Schulhofer-Wohl, S. (2011). Heterogeneity and tests of risk sharing. *Journal of Political Economy* 119(5), 925–958.
- Scott, J. C. (1977). *The moral economy of the peasant: Rebellion and subsistence in Southeast Asia*. Yale University Press.
- Shrinivas, A. and M. Fafchamps (2018, October). Testing efficient risk sharing with heterogeneous risk preferences: Comment. *American Economic Review* 108(10), 3104–13.
- Stokey, N. L. and J. Robert E. Lucas (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Suri, T. (2011). Estimating the extent of local risk sharing between households. Unpublished.
- 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.
- Udry, C. (1990, 09). Credit Markets in Northern Nigeria: Credit as Insurance in a Rural Economy. *The World Bank Economic Review* 4(3), 251–269.

- Weerdt, J. D. and M. Fafchamps (2011). Social identity and the formation of health insurance networks. *The Journal of Development Studies* 47(8), 1152–1177.
- Wilson, R. (1968). The theory of syndicates. *Econometrica* 36, 119 – 132.
- World Bank (2014). World development indicators database. data retrieved from <https://data.worldbank.org>.
- Zeldes, S. P. (1989a). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy* 97(2), 305–346.
- Zeldes, S. P. (1989b, 05). Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. *The Quarterly Journal of Economics* 104(2), 275–298.

## A Additional Figures and Tables



Figure A1: Location of ICRISAT VDSA villages - 30 villages across 8 states in 2010

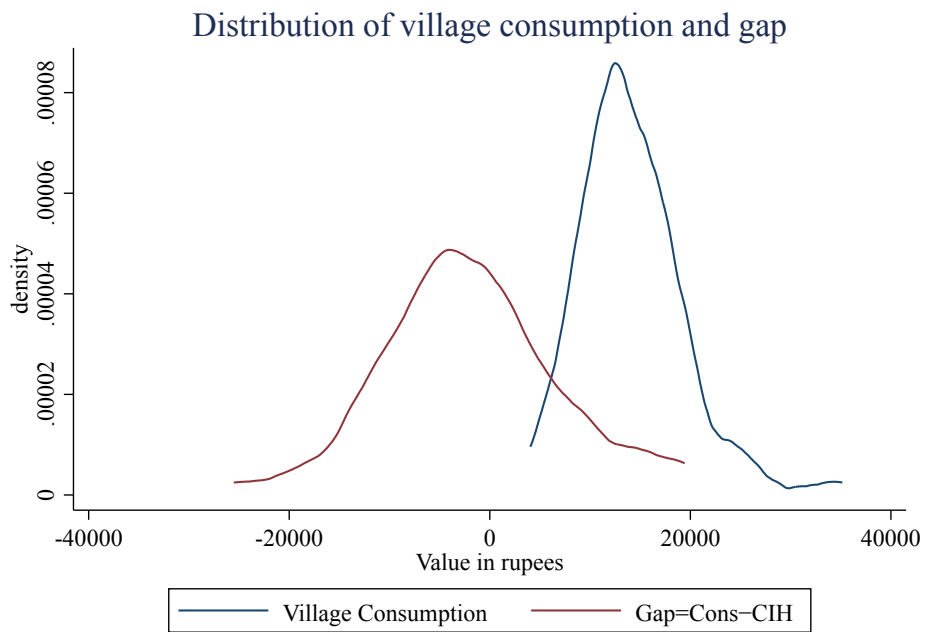


Figure A2: Distribution of Village Consumption and Unexplained Gap

Table A1: CARA model (difference in levels) controlling for Price index

	Total expenditure (1)	Homothetic preferences		Non-homothetic	
		Food (2)	Non-food (3)	Food (4)	Non-food (5)
<b>Panel A: Homogeneous preferences</b>					
Average village expenditures	0.770*** (0.038)	0.552*** (0.090)	1.007*** (0.058)	0.625*** (0.114)	0.718*** (0.041)
Cash-in-hand	0.021*** (0.007)	0.012*** (0.004)	0.032** (0.012)	0.016*** (0.005)	0.019** (0.007)
Earned Income	0.013 (0.011)	0.015* (0.008)	0.010 (0.016)	0.016 (0.011)	0.005 (0.010)
Village Food CPI	1.665 (6.814)	12.555 (16.868)	-10.234 (24.754)	16.544 (18.910)	-4.502 (17.381)
Non-food CPI / Aggregate CPI	-9.311 (32.018)	-107.460 (100.579)	97.940 (81.795)	-106.811 (119.368)	68.470 (56.706)
Observations	5256	5256	5256	5256	5256
<b>Panel B: Heterogeneous preferences</b>					
Cash-in-hand	0.019** (0.007)	0.008** (0.004)	0.030** (0.012)	0.012** (0.005)	0.016** (0.007)
Earned Income	0.010 (0.010)	0.012* (0.007)	0.007 (0.015)	0.012 (0.009)	0.002 (0.009)
Village Food CPI	-3.239 (6.951)	11.076 (19.358)	-18.880 (29.823)	13.924 (19.821)	-8.586 (16.758)
Non-food CPI / Aggregate CPI	-32.651 (21.385)	-155.947* (84.285)	102.078 (85.335)	-146.922 (105.973)	39.109 (66.526)
Observations	5256	5256	5256	5256	5256

*Notes:* This table reports the results from household panel-pooled estimation of first differences in levels, based on CARA model. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. The unit of observation is a household-year. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. Panel A reports the test results of full risk pooling under homogeneous preferences and Panel B reports the test results of full risk pooling after accounting for heterogeneity in risk and time preferences. Village expenditures represents the village leave-out-mean. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures that correspond to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. In Panel B, the dependent variable is transformed to account for heterogeneity in risk and time preferences – see text for details. The top and bottom 1% of all values are trimmed to eliminate outliers and large measurement errors. All standard errors are clustered at the village level. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.

Table A2: CRRA model(difference in logs) controlling for Price index

	Total expenditure (1)	Homothetic preferences		Non-homothetic	
		Food (2)	Non-food (3)	Food (4)	Non-food (5)
<b>Panel A: Homogeneous preferences</b>					
Average village expenditures	0.895*** (0.020)	0.630*** (0.065)	1.350*** (0.135)	0.676*** (0.063)	0.771*** (0.076)
Cash in hand	0.076*** (0.011)	0.056*** (0.011)	0.091*** (0.021)	0.063*** (0.012)	0.057*** (0.013)
Earned Income	-0.000 (0.007)	0.012 (0.008)	-0.020 (0.012)	0.013 (0.009)	-0.010 (0.008)
Village Food CPI	-0.022 (0.050)	-0.063 (0.212)	0.324 (0.467)	-0.050 (0.231)	0.138 (0.247)
Non-food CPI / Aggregate CPI	-0.011 (0.031)	-0.241* (0.131)	0.531* (0.281)	-0.261* (0.146)	0.257* (0.145)
Observations	4491	4491	4491	4491	4491
<b>Panel B: Heterogeneous preferences</b>					
Cash in hand	0.069*** (0.011)	0.043*** (0.011)	0.095*** (0.021)	0.051*** (0.012)	0.048*** (0.013)
Earned Income	0.001 (0.006)	0.013 (0.008)	-0.019 (0.012)	0.014* (0.008)	-0.009 (0.007)
Village Food CPI	-0.052 (0.044)	-0.096 (0.244)	0.300 (0.529)	-0.082 (0.253)	0.107 (0.229)
Non-food CPI / Aggregate CPI	-0.029 (0.029)	-0.323** (0.137)	0.622** (0.294)	-0.332** (0.148)	0.209 (0.148)
Observations	4491	4491	4491	4491	4491

*Notes:* This table reports the results from household panel-pooled estimation of first differences in levels, based on CARRA model. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. The unit of observation is a household-year. Each column presents the results from a separate regression on a different dependent variable. Column and row headings correspond to the dependent and regressor variables, respectively. Panel A reports the test results of full risk pooling under homogeneous preferences and Panel B reports the test results of full risk pooling after accounting for heterogeneity in risk and time preferences. Village expenditures represents the village leave-out-mean. Liquid wealth is the sum of the net credit position (saving-borrowing+lending) and liquid assets (livestock, consumer durables and inventory value of crops, inputs and fuel). Earned income is the sum of income from crop, livestock and wages. The results under the two columns 'Homothetic preferences' assumes linear Engel curves: food and non-food expenditures are divided by the sample average budget share. In the two columns 'Non-homothetic preferences', food and non-food expenditures represent the predicted total expenditures that correspond to a particular food or non-food level of expenditures, based on sample-estimated Engel curves. In Panel B, the dependent variable is transformed to account for heterogeneity in risk and time preferences – see text for details. The top and bottom 1% of all values are trimmed to eliminate outliers and large measurement errors. All standard errors are clustered at the village level. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.

Table A3: Village precautionary savings and transfers

	Total expenditure	
	(1)	(2)
Cash-in-Hand	0.070** (0.029)	0.066** (0.029)
Earned Income	-0.006 (0.047)	0.002 (0.047)
Net Gifts and remittances		-0.297 (0.258)
Net loans		0.014 (0.010)
R-squared	0.069	0.109
Observations	120	120

*Notes:* This table reports the test of precautionary saving at the village level, including outside transfers. Each column presents the results from a separate regression on total village expenditures. The unit of observation is a village-year and all variables represent village averages. All variables are in 2010 rupees per adult equivalent, aggregated over the entire year. Standard errors are reported in parenthesis. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.



## B Testing strategy with CRRA preferences

Building on the work of Mace (1991) and Townsend (1994), our testing strategy can easily be extended to the case where individuals have CRRA preferences of the form  $U_i(c) = \frac{1}{1-\gamma_i} c^{1-\gamma_i}$  where parameter  $\gamma_i$  is the coefficient of relative risk aversion of individual  $i$ . Under CRRA, the FOC for perfect risk sharing by the social planner simplifies to:

$$\eta_i \rho_i^t c^{-\gamma_i} = \lambda_{ts}$$

Taking logs and rearranging, we get:

$$\log c_{its} = \frac{\log \eta_i}{\gamma_i} + \frac{\log \rho_i}{\gamma_i} t - \frac{1}{\gamma_i} \log \lambda_{ts} \quad (15)$$

Averaging over all  $N$  individuals in the village and solving for  $\log \lambda_{ts}$  yields an expression for average village log consumption  $\overline{\log c_{ts}} \equiv \frac{1}{N} \sum_{i=1}^N \log c_{its}$ , which we use to replace the common Lagrange multiplier in equation (15). We then obtain:

$$\log c_{its} = \frac{1}{\gamma_i} \left[ \log \eta_i - \frac{\frac{1}{N} \sum_{j=1}^N \frac{\log \eta_j}{\gamma_j}}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \right] + \frac{1}{\gamma_i} \left[ \log \rho_i - \frac{\frac{1}{N} \sum_{j=1}^N \frac{\log \rho_j}{\gamma_j}}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \right] t + \frac{1/\gamma_i}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \overline{\log c_{ts}} \quad (16)$$

Under homogeneous risk preferences, the regression used to test efficient risk sharing becomes:

$$\Delta \log c_{it} = \beta_1 + \beta_2 \Delta \overline{\log c_t} + \beta_3 \Delta \log y_{it} + \beta_4 \Delta \log w_{it} + \epsilon_{it} \quad (17)$$

where the two exclusion restrictions in levels present in equation (8) have been suitably replaced by their equivalent in logs.

Under heterogeneous preferences, we similarly start by normalizing risk preferences relative to their mean by imposing that  $\frac{1}{N} \sum_{i=1}^N \frac{1}{\gamma_i} = 1$ . With this normalization, we obtain:

$$\log c_{it} = \frac{1}{\gamma_i} \left[ \log \eta_i - \frac{\frac{1}{N} \sum_{j=1}^N \frac{\log \eta_j}{\gamma_j}}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \right] + \frac{1}{\gamma_i} \left[ \log \rho_i - \frac{\frac{1}{N} \sum_{j=1}^N \frac{\log \rho_j}{\gamma_j}}{\frac{1}{N} \sum_{j=1}^N \frac{1}{\gamma_j}} \right] t + \frac{1}{\gamma_i} \overline{\log c_t}$$

To estimate this model, we first need to obtain estimates of all individual  $\gamma_i$  and  $\rho_i$  by running a model of the form:

$$\log c_{it} = \alpha_i + \theta_i t + \beta_i \overline{\log c_t} + \epsilon_{it} \quad (18)$$

using, as before, monthly consumption data on household  $i$ . We then recover structural param-

eters using the following equalities:

$$\beta_i = \frac{1}{\gamma_i} \tag{18A}$$

$$\alpha_i = \frac{1}{\gamma_i} \left[ \log \eta_i - \frac{1}{N} \sum_{j=1}^N \frac{\log \eta_j}{\gamma_j} \right] \tag{18B}$$

$$\theta_i = \frac{1}{\gamma_i} \left[ \log \rho_i - \frac{1}{N} \sum_{j=1}^N \frac{\log \rho_j}{\gamma_j} \right] \tag{18C}$$

In this case, the estimated  $\gamma_i$  can be interpreted as capturing the extent to which the coefficient of relative risk aversion of individual  $i$  differs from the average degree of relative risk aversion in the sample. The formulas for recovering welfare weights and time preference parameters are unchanged.

It follows that, apart from a slight difference in the normalization of risk preferences, the estimation of the CRRA and CARA model household by household is very similar, except that, in the CARA case, individual and village consumption are expressed in levels while they appear in logs in the CRRA case. The estimated regression model (10) for the CRRA case is thus:

$$\log c_{it} - \hat{\beta}_i \overline{\log c_t} - \hat{\theta}_i t = \xi \log y_{it} + \zeta \log w_{it} + \epsilon_{it}$$

or, expressed in first difference:

$$\Delta \log c_{it} - \hat{\beta}_i \Delta \overline{\log c_t} - \hat{\theta}_i = \xi \Delta \log y_{it} + \zeta \Delta \log w_{it} + \epsilon_{it}$$

## C Accounting for Measurement Error

Because we do not observe all the households in a village, but only a sample, there is a measurement error in the Townsend test that causes a downward bias in  $\beta_2$ . Let the true data generating process be:

$$\begin{aligned}c_i &= \beta_0 + \beta_2 \bar{x}_i + u_i \\x_i &= \bar{x}_i + e_i\end{aligned}$$

where  $\bar{x}_i$  is the true village mean for individual  $i$ ,  $x_i$  is a sample mean of  $\bar{x}_i$  based on a sample of size  $N$ , and  $e_i$  is the measurement error. Under the null of perfect risk sharing, the estimated model is:

$$c_i = \beta_0 + \beta_2^* x_i + u_i$$

and the magnitude of the bias is given by:

$$E[\beta_2^*] = \beta_2 \left(1 - \frac{\sigma_e^2}{\sigma_x^2}\right)$$

Since the standard error of a sample mean is  $\sigma_e = \frac{\sigma_x}{\sqrt{N}}$ , the downward bias is approximately:

$$\begin{aligned}E[\beta_2^*] &= \beta_2 \left(1 - \frac{\sigma_x^2/N}{\sigma_x^2}\right) \\&= \beta_2 \left(1 - \frac{1}{N}\right) \\&\simeq \beta_2 0.976\end{aligned}$$

when using the full sample of 1296 households divided into 30 villages, i.e.,  $N = 42.16$  on average across villages. We can use this result to perform an approximate correction of the  $\hat{\beta}_2$  coefficients estimated for total consumption, i.e., by dividing  $\hat{\beta}_2$  by 0.976. We can check whether, as the result of this correction, the revised  $\hat{\beta}_2$  is close enough to 1 to fail to reject full risk pooling. A similar calculation for the caste regressions yields a correction factor of 0.946. In both cases, we see that the correction for sampling error is not large enough to qualitatively change our reported findings.

## D Village precautionary saving, consumption durables and optimal portfolio

In this Appendix, we examine the optimal durable consumption and portfolio choice that corresponds to our perfect risk pooling hypothesis. To this purpose, we present a stylized model of a hypothetical social planner solving a standard precautionary saving problem. To recall,  $y_t$  represents the aggregate earned income during year  $t$ ,  $c_t$  is the aggregate village consumption in that year,  $w_t$  is the total value of village wealth at the beginning of year  $t$  and  $w_{t+1}$  is the value of wealth at the end of that year. In the absence of measurement error, the following accounting identity must hold:

$$c_t = y_t + w_t - w_{t+1} = y_t + d_t \quad (19)$$

where dissaving  $d_t \equiv w_t - w_{t+1}$  and where, for notational simplicity, we have included the return to productive assets into income  $y_t$ . Put in words, consumption must be financed either from income or from dissaving. In the presence of classical measurement error in income, consumption, or wealth, this identity should still be true on average. It follows that, for a given level of consumption  $c_t$ , we should observe a negative correlation between income  $y_t$  and dissaving  $d_t$ :

$$d_t = \alpha_0 + \alpha_1 y_t + \alpha_2 c_t + u_t \quad (20)$$

where  $\alpha_1$  should be converging to  $-1$  when measurement error tends to 0. This observation forms the starting point of our testing strategy.

We now turn to the breakdown of wealth across different assets. To derive the efficient risk pooling strategy, we imagine a village social planner who, at time  $t$ , optimally allocates village savings  $w_t$  across the four categories of assets present in our data: consumer durables  $b_t$  (e.g., bicycles, appliances); livestock  $l_t$ ; gold and silver  $g_t$  (typically in the form of jewelry); and the net stock of financial assets  $f_t$  (i.e., savings minus outstanding loans due to the rest of the world). By construction, the data always satisfies the following identity:

$$w_t = b_t + l_t + g_t + f_t \quad (21)$$

We are interested in the correlation between  $y_t$  and time changes in each of the components of wealth  $w_t$  that can theoretically be expected:

$$d_t^j = \alpha_0^j + \alpha_1^j y_t + \alpha_2^j c_t + u_t^j \text{ for } j \in \{b, l, g, f\} \quad (22)$$

where  $d_t^j$  is the change in asset  $j$  between the beginning and the end of the year. This is equivalent to asking whether some components of wealth adjust more to income shortfalls than others. Since all four types of assets are held on average in each village, it cannot be the case

that some assets dominate others. Hence, to derive a meaningful demand predictions for all four assets, we make simple but realistic assumptions about what differentiates them.

We start by assuming that all assets can be bought and sold, but that there is a price difference  $\beta_j$  (which we dub 'brokerage fee') between the buying and selling price of each asset type  $j \in \{b, l, g, f\}$ . For gold and silver and for financial assets,  $\beta_j$  is assumed to be 0 while it is assumed to be moderate for livestock and large for durables. This implies that jewelry and financial assets are the most liquid, while durables are the least liquid. Other things being equal, this predicts that villages should smooth income shocks using jewelry and financial assets first, and durables last (i.e., only when they have exhausted their stock of liquid assets, e.g., Deaton 1991). Based on these assumptions alone, we expect dissaving in jewelry and financial assets to be the most negatively correlated with income, conditional on a given consumption level  $c_t$  – i.e., we expect that  $\alpha_1^j$  is more negative for jewelry and financial assets than for livestock, and least negative for durables.

Next we refine the model to include a flow of consumption utility generated by some assets such as durables and jewelry. For jewelry and the kind of durable goods in our study, the demand for this consumption utility increases more than proportionally with income. It follows that, when income falls, it is optimal for households to shift consumption away from durables and towards necessities such as food products and non-durables. This implies a redistribution of assets away from jewelry and consumption durables. Hence, based on this income effect alone, we would expect that, for a given consumption, income is more negatively correlated with jewelry and durables than with financial assets. This, however ignores brokerage costs which, in the context of our study area, are larger for durables, which lose value over time, than for jewelry, which typically does not – i.e.,  $\beta_b > \beta_g$ . Combining the income and brokerage effects, we therefore expect the negative correlation between income and assets to be strongest for jewelry, weakest for consumer durables, and intermediate for financial assets.

Predictions for livestock are similar to those for financial assets given that, in our data, livestock consumables such as dairy products are imputed a monetary value that is counted as part of income and added to consumption if they are consumed by the household. It is, however, conceivable that livestock production benefits from increasing returns to scale – e.g., due to a minimal herd size to be economically profitable (e.g., Lybbert et al; Barrett et al.) – which creates a hurdle in the acquisition of livestock (cow paper showing flypaper effect). There is indeed evidence that livestock-raising households are reluctant to liquidate their livestock in times of duress (e.g., Fafchamps et al. 1998). Based on this, we expect little readjustment of the village asset portfolio away from livestock, and thus a negligible negative correlation between income and livestock. This logic applies, even more strongly, to land – which we have excluded from analysis given the very low frequency of land sales and purchases in our data, and the fact that these transactions take place almost exclusively within the village itself and thus cannot help smooth village consumption.

The predictions for the optimal village risk pooling model can be summarized as follows:

1. In the absence of measurement error, the level of village wealth should be perfectly inversely correlated with income, conditioning on village consumption – i.e., the coefficient  $\alpha_1$  in equation (20) should be  $-1$ . Measurement error in income biases the estimated  $\hat{\alpha}_1$  coefficient towards 0.
2. In the absence of brokerage costs, increasing returns, and flow of consumption utility from assets, the portfolio composition of village assets should remain constant, conditional on village consumption. Consequently, coefficients  $\alpha_1^j$  in equation (22) should be equal to the share  $s_j$  of asset  $j$  in the village portfolio. Given identity (21), measurement error in income biases  $\hat{\alpha}_1^j$  coefficients towards 0 in proportion to their portfolio share, i.e.,  $\hat{\alpha}_1^j = s_j \hat{\alpha}_1$ .
3. In the presence of brokerage costs  $\beta_j$  (i.e., a difference between the buying and selling cost of assets), it is optimal for liquid assets (those with lower brokerage cost) to vary more than illiquid assets. Since  $\beta_j$  is lowest for financial assets and jewelry and highest for durables, we expect  $\hat{\alpha}_1^f \simeq \hat{\alpha}_1^g < \hat{\alpha}_1^l < \hat{\alpha}_1^b$ .
4. In the presence of consumption utility for some assets, we expect portfolio composition to move away from those assets whose consumption is a luxury, such as jewelry and most consumer durables in our data. Combined with brokerage cost effects, this implies  $\hat{\alpha}_1^g < \hat{\alpha}_1^f < \hat{\alpha}_1^l$  with the relative position of  $\hat{\alpha}_1^b$  depending on the relative strengths of the brokerage cost effect (that makes durables less liquid) and the income elasticity effect (that makes them less consumed when income falls). In practice, we expect the brokerage effect to dominate.
5. While livestock is a relative liquid asset in the context of our study, the presence of increasing returns in livestock production creates threshold effects that militate towards using livestock as buffer against temporary income shortfalls. In this case, we expect  $\hat{\alpha}_1^l \simeq 0$ .

The optimal asset portfolio can also vary with asset prices. We do not observe asset price variations in our data, except for jewelry, which we expect to closely follow the international gold price, converted in Indian rupees. <sup>34</sup>Given that gold prices follow more or less a random walk over our period of study, the optimal jewelry share in village wealth should remain constant, which implies selling jewelry when the price of gold goes up.

---

<sup>34</sup>While study households also own silver jewelry, gold accounts for most of the jewelry value.

Table A4: Decomposition of wealth into its components

	Total Dissavings (1)	Components of Dissavings			
		Net Credit [Savings -Dues] (2)	Gold and Silver (3)	Livestock (4)	Other Consumer Durables (5)
Earned income	-0.093 (0.092)	-0.057 (0.055)	-0.049 (0.053)	-0.002 (0.019)	0.022 (0.040)
Consumption	0.143 (0.130)	0.295*** (0.078)	-0.118 (0.075)	0.010 (0.027)	-0.056 (0.056)
Observations	120	120	120	120	120
R-squared	0.013	0.117	0.051	0.001	0.009

*Notes:* Standard errors are reported in parenthesis. \*p<0.10 \*\* p<0.05 \*\*\* p<0.01.