Distinguishing barriers to insurance in Thai villages*

Cynthia Kinnan†

October, 2009

PRELIMINARY AND INCOMPLETE

Abstract

In both developing and developed countries, consumption insurance has been found to be incomplete: consumption co-moves with income. Models with limited commitment, moral hazard, and hidden income shocks have been proposed to explain this lack of full insurance. Limited commitment and moral hazard have generally been found to fit consumption and income data better than full insurance or permanent income models, but these models have not been tested against the alternative of hidden income. I show that the way history matters in forecasting consumption can be used to distinguish hidden income from other models of incomplete insurance. The inverse of a household's marginal utility last period is a sufficient statistic for "history" under limited commitment and moral hazard, but under hidden income, lagged income has additional predictive power in forecasting consumption. In a ten-year panel from rural Thailand, neither limited commitment or moral hazard can fully explain the relationship between income and consumption—the need to give households incentives to truthfully reveal their income appears to play a role.

JEL codes: D82, D91, O12

*I am grateful to Abhijit Banerjee, Esther Duflo, Jerry Hausman, Tavneet Suri, Robert Townsend and Iván Werning for extremely helpful comments, discussion and suggestions. John Felkner, Anan Pawasutipaisit and Archawa Paweenawat provided valuable advice on the Thai data. All remaining errors are my responsibility.

†ckinnan@mit.edu

1 Introduction

Risk to households' incomes is widespread in developing countries—crops and businesses fail, jobs are lost, livestock die, prices fluctuate, family members fall sick, expenses for marriages, births, and funerals arise. If perfect insurance were available, such income risk would not translate into fluctuations in household consumption. In fact, poor households in many developing countries are insured against short-term, idiosyncratic income shocks to a surprising degree, despite absent or imperfect markets for formal insurance, credit, and assets, as documented by Rosenzweig (1988), Townsend (1994, 1995), Udry (1994), Morduch (1995), Suri (2005) and others. However, households are generally not completely insured—income and consumption are typically found to be positively correlated[1], and serious income shocks like severe illness translate into reduced household consumption (Gertler and Gruber 2002). Households neither seem to live "hand to mouth," with shocks to income translating one-for-one to fluctuations in consumption, nor to be fully insured, with consumption completely buffered against shocks to income.

Furthermore, as discussed below, the hypothesis that households smooth consumption only with a borrowing-savings technology is generally rejected, and there is direct evidence that households make state-contingent transfers to others in their village ((Scott 1976), (Cashdan 1985), (Platteau and Abraham 1987), (Platteau 1991)). In other words, households do obtain insurance from others in their village. A natural question is then, why is this insurance not complete? Among the reasons proposed for the failure of full insurance are issues of information—one person's income realizations or actions may not be observable to others; and commitment—households with high incomes, who would be required by full insurance to make transfers to others, may leave the insurance arrangement instead. Knowing what barrier is most important in a given community is important for evaluation of policies that could affect the sustainability of informal insurance. These include access to saving (such as rural bank expansion; the impact of savings access under limited commitment is

[1] Hayashi et al. (1996) review the literature on full consumption smoothing in the US, and find that neither endogeneity of labor nor nonseparability between labor and consumption explains the rejection of full smoothing of food consumption in the PSID. Blundell et al. (2008) document that peristent income shocks are partially insured in the US, and even transitory shocks are not fully insured for low-wealth households.

discussed in Ligon, Thomas and Worrall 2000); technologies that make observing others' incomes easier (such as crop price information dissemination) or harder (taking individual deposits rather than collecting savings at a group meeting, or access to larger, more anonymous markets); weather, health or other formal insurance which makes leaving community insurance more palatable (Albarran and Attanasio 2003), (Attanasio and Rios-Rull 2000), (Rosenzweig and Wolpin 1993); and policies that expand communities' sanctioning ability (such as community-allocated aid; see Olken 2005), or restrict it (such as road access; see Townsend 1995). Changes in labor institutions may also affect problems of hidden income and moral hazard; these include replacing local labor-sharing arrangements with hired outside labor or augmenting family labor with local hired labor. Land-titling and conditional cash transfer programs may also affect–and be affected by–limited commitment, moral hazard and hidden income.

Several papers have examined whether limited commitment- or moral hazard-constrained insurance explains consumption data better than autarky or full insurance. Ligon (1998) tests moral hazard-constrained insurance against full insurance and borrowing-saving (i.e., the permanent income hypothesis) in India using ICRISAT village data, and finds that moral hazard best explains consumption data in 2 of 3 villages; in the third some households' consumptions are better explained by the PIH. Ligon, Thomas and Worrall (2002) test full insurance against limited commitment, also in the ICRISAT villages. They find that limited commitment explains consumption dynamics, but not why high-income households consume as little as they do relative to low-income households. Lim and Townsend (1998) incorporate capital assets and livestock into a moral hazard-constrained insurance model, and find that it fits the ICRISAT consumption data better than the PIH or full insurance. Cox et al. (1998) argue that features of lending in Peruvian villages are inconsistent with full insurance or the PIH, but consistent with limited commitment. Albarran and Attanasio (2003) show that the comparative statics of a limited commitment model are matched by data from Mexico following the introduction of Progresa. Dubois et al. (2008) develop a model with limited commitment and incomplete formal contracts and find that its predictions are matched in Pakistani data. Finally, a number of papers use data from developed countries to test for a particular insurance friction; Attanasio and Pavoni (2007) provides a review. Kocherlakota and Pistaferri (2009) review the literature which uses the asset-pricing implications of incomplete markets (borrowing-lending only) and private information (moral hazard/adverse selection) economies; Kocherlakota and Pistaferri find that the asset-pricing implications of the moral hazard/adverse selection model fit US, UK and Italian data with a "reasonable" coefficient of relative risk aversion, while the implications of the borrowing-lending model are rejected.

Also related is Paulson, Townsend and Karaivanov (2006). Using Thai data (from the initial 1997 wave of the survey used in this paper) on occupational choices, they consider whether limited liability or moral hazard better explains entry into entrepreneurship, and find that moral hazard limits access to business credit—a sufficient but not necessary condition for moral hazard to limit insurance. On the other hand, Karaivanov and Townsend (2008) use consumption, investment, and income data from Thailand to test between no-insurance regimes (PIH, savings only, and autarky) versus full insurance or moral hazard-constrained insurance regimes, and generally find that the PIH best fits the data.

To my knowledge, there have not been previous attempts to test a third class of models which predict incomplete insurance: the hidden income shock model of Townsend (1982) and Thomas and Worrall (1990). The closest papers may be Kaplan's 2006 working paper, which derives quantitative predictions about the amount of risk sharing that would arise, for a given wage distribution, under limited commitment versus a Mirrlees-style model with observed income but unobserved productivity; and the 2007 working paper of Ai and Yang which finds that a model with limited commitment and private information about productivity (but observed income) better fits quantitative features US data than a model with limited commitment alone. In contrast to a Mirrlees model with private information about productivity but observed income, I abstract away from unobserved productivity differences and consider a model in which the principal (i.e. the community) is unable to verify the income of the agent (household), and so truthful reporting must be incentive-compatible for the household.

If binding, such truth-telling incentive-compatibility constraints rule out full insurance, as do the participation constraints of the limited commitment model and the effort incentive-compatibility constraints of the moral hazard model[2], generating a positive correlation between changes in income and changes in consumption. Indeed, because all three models of incomplete insurance predict both a positive correlation between income and consumption changes, as well as some degree of "spillover," i.e. one household's income realizations affecting the consumption of other households in the village, it is possible that tests of a particular insurance friction versus borrowing-saving and/or full insurance will conclude in favor of the incomplete insurance model. The objective of this paper is to develop an empirically testable prediction which distinguishes this hidden income model from limited commitment- and moral hazard-constrained insurance.

To briefly preview my results, I show that, when insurance is constrained by limited commitment or moral hazard, a household's "history" matters in a specific way in predicting that household's current consumption: conditional on the village's shadow price of resources (a measure of the aggregate shock faced by the village), a household's lagged inverse marginal utility ("LIMU") is a sufficient statistic for forecasting the household's consumption: no other past information should improve the forecast of current consumption made using LIMU. Allowing the distribution of household income to depend on actions taken by the household in the past (investment, for instance) does not overturn the sufficiency result. On the other hand, when household income is unobserved, a household's LIMU is no longer a sufficient statistic in forecasting consumption. Because low-income households are optimally assigned the lowest consumption, hence the highest marginal utility, their temptation to claim *even lower* income (resulting in a higher transfer), is high. Incentive compatibility for these households is attained by reducing the expected future surplus promised to them, relative to their current consumption, with the result that using LIMU alone results in forecasts of current consumption that are too high for low-past-income households, and increasing in lagged income. The intuition for this difference between the hidden income model, on one hand, and limited commitment and moral hazard, on the other is that in the latter two, because income is observed–and hidden saving is ruled out, as explained below–the community can effectively control consumption, by controlling income-contingent

[2] The relationship between income and consumption need not be everywhere positive under a moral hazard model if the likelihood ratio is not monotone (Milgrom 1981). However, incentive compatibility requires that consumption be increasing in output on average. (Moreover, if agents can costlessly "burn output," this may also require monotonicity.)

transfers. As a result, there is no deviation from the optimal division of promised utility across periods–so that utility (via transfers) in the current period and (via promised utility) in future periods are equally valuable to the household. As a result, all past information relevant to forecasting future consumption When income is private information, in contrast, consumption is not effectively controlled by the community, and the constrained-optimal schedule of transfers and promised utilities distorts the tradeoff between current consumption and future expected utility, with households announcing low incomes being penalized more in terms of future utility than current consumption.

I examine the relationship between LIMU and current consumption in rural Thailand using 10 years (1997-2006) of the Townsend Thai Annual Survey, and find that sufficiency of LIMU is rejected: lagged income has predictive power in forecasting consumption. Moreover, the prediction errors generated with LIMU alone display a significant, positive correlation with lagged income, as predicted by the hidden income model. This suggests that the need to give households incentives to truthfully reveal their income plays a role in generating the observed comovement between income and consumption. While measurement error in lagged consumption is a concern, it does not appear to drive the results.

The rest of the paper is organized as follows: In section 2 I present evidence that households in this sample of rural Thais are insured to a large but incomplete extent: positive income changes predict positive changes in consumption, a fact consistent with no-insurance (PIH), limited commitment, moral hazard or hidden income. I also argue that villages do provide insurance beyond what households could achieve in autarky, and that villages do not face period-by-period budget constraints—villages appear to have access to a relatively frictionless intertemporal technology. This matters for the specification of the test of limited commitment, and for policy evaluation. In Section 3 I outline the approach for distinguishing hidden income from limited commitment and moral hazard, based on tests of limited commitment developed in Kocherlakota (1996). I derive a testable prediction which distinguishes hidden income from limited commitment, moral hazard and PIH regimes. In Section 4 I discuss the data and present results from these tests. Section 5 presents robustness checks, and Section 6 concludes. Proofs are contained in an Appendix.

2 Insurance and credit in Thai villages

2.1 Insurance is imperfect...

If households were perfectly insured, there would be no need to look for evidence of a particular insurance friction—if consumption did not move with own income, and all villagers' consumptions moved one-for-one with average village income, this would mean that none of hidden income, moral hazard, or limited commitment was a significant impediment to full insurance. This is not the case for rural Thailand. Townsend (1995) finds imperfect insurance in northern Thai villages in the years 1989-1991. In the 1997-2006 waves of the Townsend Thai Annual Survey, estimating

$$\Delta \ln c_{PCvt} = \alpha \Delta \ln y_{ivt} + \delta_{vt} + \Delta \varepsilon_{ivt} \tag{1}$$

where c_{PC} is per-capita consumption, Δy_{ivt} is change in household income, and δ_{vt} is a village dummy capturing a common change in villagers' consumption due to change in aggregate resources, yields $\hat{\alpha} = .042$ ($t = 4.3$). (See Table 1, col 2.) (Consumption is measured throughout this paper as expenditure and adjusted for household composition using the equivalence scale used by Townsend (1994). Using an equivalence scale which allows for economies of scale in household size does not change any of the results.)

Another telling feature of the data is a large amount of movement in the village per capita expenditure (PCE) distribution: the correlations between household PCE rankings in adjacent years range from .824 (1999-2000) to .539 (2000-2001). (See Table 2, Panel A.) Moreover, PCE rank changes are not at random, as they would be if driven by classical error in the measure of expenditure, but are predicted by income changes, with a 10% change in income associated with an increase in the PCE distribution of about one-half of a ranking; an ordered probit regression also yields a significant positive association. (See Table 2, Panel B.)

Absent taste shocks and assuming no heterogeneity in risk aversion, churn in the consumption distribution is incompatible with full insurance (see the proof of Proposition 1 in the appendix), as is $\alpha \neq 0$ in (1). However, insurance constrained by either limited commitment, hidden income, or moral hazard[3] would predict both $\alpha > 0$ and $corr(rank_{it}, rank_{it'}) < 1$. Thus, these findings cannot help distinguish among the models of imperfect insurance. Below I suggest how this can be done.

2.2 ...but villages do provide insurance

Finding $\alpha < 1$ in equation (1) does not establish that villages provide insurance: households could smooth consumption using borrowing and saving (Hall 1978),(Deaton 1991), or the relevant risk-sharing network might be a different group, such as kinship groups. (Finding $\alpha < 1$ is also compatible with households living "hand-to-mouth," consuming their income each period and using no smoothing devices at all, if income is measured with error. I discuss this possibility below.) As noted by Suri (2005), an implication of a set of households belonging to an insurance group is that household consumption is less correlated with household income, conditional on total group consumption, than group average consumption is correlated with group average income. This can be straightforwardly tested by testing the hypothesis that the village-year effects in (1) are jointly insignificant in explaining household income changes. This hypothesis is strongly rejected ($F(95, 2577) = 2.125, p = 0.000$), indicating that there is a significant tendency for household consumptions in a given village-year to move together.

Of course, in this same regression the coefficient on the change in household income, conditional on the average village change, is significantly different than zero: belonging to a village network does not remove all idiosyncratic risk, but village networks do manage to reduce dependence of household consumption on contemporaneous income. In Section 3 I consider three models that attempt to rationalize this finding of partial insurance: limited commitment, hidden income, and moral hazard.

2.3 Credit is available

As discussed below, the form of the contract that the hypothetical village social planner can offer to a household depends on whether the village's budget must balance each period.

[3] Grossman and Hart (1983) show that in a moral hazard regime, even if the PDF of output satisfies the monotone likelihood ratio property (MLRP), consumption will not necessarily be everywhere increasing in income, but incentive compatibility requires that consumption be increasing in income on average.

If so, a constraint on the planner's problem is that, at each date and state of the world, total consumption among the villagers ($j \in V$) cannot exceed their total income:

$$\sum_{j \in V} c_{jt} \leq \sum_{j \in V} y_{jt}, \forall t. \tag{2}$$

Alternatively, if borrowing and savings are possible, subject only to a terminal condition[4], village assets a_{vt} evolve according to

$$R^{-1} a_{v,t+1} = a_{vt} + \sum_{j \in V} (y_{jt} - c_{jt}) \tag{3}$$

where R is the gross interest rate and y_{jt}, k_{jt}, c_{jt} are respectively the income, investment, and consumption of villager $j \in V$.

Just as dependence of households' consumption on their contemporaneous income can be tested with (1), dependence of village consumption at time t on village income at t can be tested with

$$\Delta \overline{\ln c_{PCvt}} = \gamma \Delta \overline{\ln y_{vt}} + \varepsilon_{vt} \tag{4}$$

where $\overline{\ln c_{PCvt}}$ is village average log per capita consumption and $\overline{\ln y_{vt}}$ is average village log income. Estimating (4) yields $\hat{\gamma} = .035, t = 0.96$ (see Table 1, col 3); this is consistent with villages acting in accordance with the permanent income hypothesis for $R = \frac{1}{1-\gamma} = 1.036$, or with villages obtaining nearly-complete insurance—village average consumption is very insensitive to contemporaneous average income. This suggests that villagers and/or village institutions (banks, moneylenders, local government, etc.) have relatively frictionless access to a national-level credit market or a set of equivalent institutions, and perhaps additional state-contingent assets.[5]

Measurement error in income is a concern in interpreting the individual and village results. Classical measurement error in income (uncorrelated with the true values of income changes and with the error terms $\Delta\varepsilon_{it}$), will attenuate $\hat{\alpha}$ and $\hat{\gamma}$ toward zero. This would

[4] $A_{T+1} = 0$ if T is finite or, if T is infinite, $\sum_{t=1}^{\infty} R^{-t}(y_{jt} - k_{jt} - c_{jt}) \leq A_0$.

[5] This is consistent with the finding of Paxson (1992), using the 1975/76, 1981, and 1986 Socio-Economic Survey, that households save a high fraction of good transitory income shocks, and dissave when hit by negative income shocks.

make the extent to which income changes predict consumption changes in the data a lower bound on the true sensitivity of consumption to income—measurement error in income cannot explain the rejection of full insurance.

Having established that income changes predict consumption changes, that villages appear to provide insurance, and that village institutions do not face period-by-period budget constraints, we can turn to testing whether limited commitment, moral hazard or hidden income limit insurance.

3 Models of optimal insurance: full insurance, limited commitment, moral hazard, hidden income

3.1 Setting

As a simplified approximation to the environment in a Thai village, consider N risk-averse households ("agents") who interact over an infinite time horizon in a mutual insurance network. Each household evaluates consumption and effort plans according to:

$$U(\mathbf{c}_i, \mathbf{e}_i) = \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left[v(c_{it}) - z(e_{it}) \right]$$

All households have a common discount factor β, and the same utility of consumption and disutility of effort functions $v()$ and $c()$, with $v' > 0$, $v'' < 0$.

I assume that, while any agent is in the network, his savings decisions are contractible. As a result savings are determined as if chosen by a welfare-maximizing planner, not to maximize the agent's own expected payoff. There is a community-controlled credit market with return R, and an autarkic storage technology with return $r \leq R$. Contractible savings implies that any net borrowing or saving by agents in the network (such that (2) does not hold with equality) will take place via the community-controlled technology.

Output can take on S values, $\{y_1, ..., y_S\}$, where a higher index means more output: $r > q \Rightarrow y_r > y_q$. The distribution of household i's output is affected by i's choice of investment/effort. Effort can take on two values in each period,

$$\begin{aligned} e_t &\in \{0,1\} \\ z(0) &= 0 \\ z(1) &= 1 \end{aligned}$$

The distribution of income is affected by household's effort in the current and the previous period:

$$\Pr(y_t = y_r) = \Pr(y_r | e_t, e_{t-1})$$

Each income realization occurs with positive probability under either high or low effort

$$\Pr(y_r | e_t, e_{t-1}) \in (0,1), \forall e_t, e_{t-1}, r$$

Define $p_{ree'} \equiv \Pr(y_r | e_t = e, e_{t-1} = e')$, the probability of income realization y_r when an effort level e is exerted in the current period and e' was exerted in the last period. Finally, I assume that effort raises expected surplus, and that the return to effort in the current period is highest when effort was also exerted in the last period:

$$\sum_{r=1}^{S} [p_{r11} - p_{r01}]\, y_r > \sum_{r=1}^{S} [p_{r10} - p_{r00}]\, y_r \geq z(1) - z(0)$$

Effort may be contractible (or it may be constrained-optimal for the agent to exert no effort), in which case there will be no issue of moral hazard over effort choice[6]. However, if effort is non-contractible, the agent's choice must be incentive compatible. When discussing moral hazard I will assume that the principal wants to incentivize effort, that is, that expected surplus from the constrained-optimal contract induing effort is greater than expected surplus from the full-insurance contract (which induces low effort).

Having set out the environment, I will briefly characterize the full insurance benchmark before introducing the constraints which lead to incomplete insurance.

[6] Townsend and Mueller (1998) find evidence that inputs such as seed and fertilizer are controlled by lenders in the ICRISAT villages in India.

3.2 Full insurance

We can find the set of first-best allocations by considering the problem of a hypothetical risk-neutral planner who maximizes the utility of villager N such that each villager 1 to $N-1$ gets at least a value u_{it} in period t. Let $\mathbf{u}_t \equiv \{u_{it}\}_{i=1}^{N-1}$. The planner chooses effort e, transfers τ_{irt}, and future promises $u_{ir,t+1}$ for each villager. Transfers, which are equal to the difference between a household's income and its consumption, $\tau_{irt} \equiv c_{irt} - y_r$, and future promises, which summarize the utility the household can expect from next period onward (Spear and Srivastava 1987), are indexed by r because they are income-contingent (though this dependence will be degenerate in the case of full insurance). The planner's value function is

$$u_N(\mathbf{u}_t, a_t, \mathbf{e}') \equiv \max_{\mathbf{e}, \{\boldsymbol{\tau}_{rt}\}, \{\mathbf{u}_{r,t+1}\}} \tag{OBJ}$$
$$\sum_{r=1}^{S} p_{ree'} v(y_r + \tau_{Nrt}) - z(e_N) + \beta \mathbb{E}_{\{y\}} u_N(\mathbf{u}_{t+1}, a_{t+1}, \mathbf{e})$$

subject to the promise-keeping constraints that each household 1 to $N-1$ must get their promised utility (in expectation):

$$\sum_{r=1}^{S} p_{ree'} [v(y_r + \tau_{irt}) - z(e_i) + \beta u_{ir,t+1}] = u_{it}, \forall i < N \tag{λ_{it}}$$

and the law of motion for assets:

$$R^{-1} a_{t+1} = a_t - \sum_{i=1}^{N} \tau_{irt} \tag{η_t}$$

As is well known (but derived in the appendix for completeness), with perfect village-level credit access, and absent problems of commitment or information, every network member's consumption is independent of their own income realization, given aggregate network resources. Therefore we have

Proposition 1 *Under full insurance, (a) realized household income has no effect on household consumption, given aggregate resources, and (b) with no preference shocks and a com-*

mon discount factor, households never change place in the village consumption distribution.

Proof. *In appendix.* ■

In summary, full insurance predicts a complete decoupling of idiosyncratic income shocks and consumption changes. Since this implication fails to hold in virtually every dataset where it has been tested, the next question is how to distinguish among models that do predict a correlation between income shocks and consumption changes. I will first discuss the limited commitment model, which has become one of the workhorse models of village insurance.

3.3 Limited commitment

3.3.1 Sufficiency of lagged inverse marginal utility

If an agent can walk away from the insurance network at any time if he can do better in autarky, Proposition 1 no longer necessarily holds (Coate and Ravallion 1993). Limited commitment imposes further constraints on the planner's problem (OBJ), which is now subject to the promise-keeping constraints (whose multipliers are λ_{it}), the budget constraint (with multiplier η_t) and the participation constraints that the expected utility an agent gets in the insurance network be at least as great as the expected utility he could achieve in autarky, choosing his own savings and effort optimally. That is, a household will only remain in the network if

$$v(y_r + \tau_{irt}) + \beta u_{ir,t+1} \geq u_{aut}(y_r, e), \forall i, r \tag{5}$$

where

$$\begin{aligned} u_{aut}(y_r, e) \equiv & \max_{s_t, e_{t+1}} v(y_r - s_t) - \beta z(e_{t+1}) \\ & + \beta \mathbb{E}\left[u_{aut}(y_{t+1} + R s_t) | e_{t+1}, e\right] \end{aligned}$$

Kocherlakota (1996) showed that, under limited commitment, the vector of lagged marginal utility ratios for every member of the insurance group,

$$\left\{ \frac{v'(c_{N,t-1})}{v'(c_{i,t-1})} \right\}_{i=1}^{N-1}$$

is a sufficient statistic for history when forecasting any household's consumption. However, Kocherlakota's result is not testable if the econometrician does not have information on all the members of the insurance group. Since consumption and income data generally come from surveys, rather than censuses, this renders the test of very limited applicability. In Kocherlakota's setting, the need to keep track of every member of the insurance network arises due to the assumption that the village as a whole cannot borrow or save. If the village can borrow and save, the shadow price of resources at time t serves as a summary measure of how much consumption must be given to other households in the village. In this case, we have the following result, which is testable with time series data for only a sample of households in a network.

Proposition 2 *With village-level credit access, conditional on the time t shadow price of resources η_t, household i's LIMU $\left(\frac{1}{v'(c_{i,t-1})}\right)$ is a sufficient statistic for household i's time t inverse marginal utility under limited commitment. When i's participation constraint binds, i's current and expected future consumption are increasing in i's income.*

Proof. *In appendix.* ■

The intuition for this result is that, when the only barrier to full insurance is the fact that the agent can walk away when income is high, the principal can allocate consumption to an agent who is tempted to walk away without affecting the incentive of any other agent to stay in the network. So the constrained agent gets current consumption and a future promise that make him exactly indifferent between staying in or leaving the network. By the envelope theorem, the Lagrange multiplier on the agent's promise-keeping constraint uniquely describes the efficient combination of $c_{it}, u_{i,t+1}$. Moreover, under limited commitment the agent's lagged inverse marginal utility fully captures the Lagrange multiplier on the promise-keeping constraint. So $\frac{1}{v'(c_{i,t-1})}$ captures all the information from time $t-1$ and earlier that is relevant in predicting time t consumption, c_{it}.

Since, as discussed below, the same sufficiency result is obtained under moral hazard (with the additional, stronger implication of an Inverse Euler equation), and, with isoelastic utility an indistinguishable result holds under the PIH[7], if we are unable to reject sufficiency of LIMU in a given setting, this does not tell us whether limited commitment, moral hazard, or borrowing-lending is a more plausible alternative. Before moving on to moral hazard, I discuss a stronger implication of limited commitment.

3.3.2 Amnesia

A stronger implication of limited commitment, which does not hold under moral hazard or PIH, is what Kocherlakota calls "amnesia:" when limited commitment binds for household i, consumption $y_r + \tau_{irt}$ and promised future utility $u_{ir,t+1}$ are pinned down by the requirement that the household be just indifferent between staying in and leaving the network, and that the utility value of current and future consumption be equated at the margin,

$$v(y_r + \tau_{irt}) + \beta u_{ir,t+1} = u^t_{aut}(y_r)$$

$$v'(y_r + \tau_{irt}) = -\left(\frac{\partial u_N(\mathbf{u}_{r,t+1})}{\partial u_{ir,t+1}}\right)^{-1}$$

independent of the time t promised value u_{it}. Thus the household's old promised value is forgotten when limited commitment binds. Kocherlakota suggested the following procedure to test for amnesia: find the network member(s) with the lowest growth in consumption between periods $t-1$ and t. Ignoring measurement error in consumption for now (see Section 5), define

$$B_t \equiv \min_{j=1,\ldots,N}\{v'(c_{jt-1})/v'(c_{jt})\}$$

Those for whom $v'(c_{jt-1})/v'(c_{jt}) > B_t$, by construction, had binding limited commitment constraints—otherwise their consumption would have been fully smoothed between periods $t-1$ and t. Those with $v'(c_{jt-1})/v'(c_{jt}) = B_t$ were not constrained, and therefore did

[7] If there was no interpersonal insurance, but households could borrow and save at a riskless rate, the Euler equation $\mathbb{E}_t u'(c_{i,t}) = \frac{1}{\beta r} u'(c_{i,t-1})$ implies that, conditional on past marginal utility (rather than inverse marginal utility), no information in the household's information set at time $t-1$ should predict time t consumption (Hall 1978). With isoelastic utility, in a log specification the sufficient statistic under limited commitment, $\frac{1}{u'(c_{i,t-1})}$, cannot be distinguished from the sufficient statistic under autarky, $u'(c_{i,t-1})$.

achieve efficient intertemporal consumption smoothing. Define the sets of constrained and unconstrained households

$$\begin{aligned} C_t &\equiv \{j : v'(c_{jt-1})/v'(c_{jt}) > B_t\} \\ U_t &\equiv \{j : v'(c_{jt-1})/v'(c_{jt}) = B_t\} \end{aligned}$$

Amnesia implies that, for any constrained household $j \in C_t$, $\frac{1}{v'(c_{j,t-1})}$ should not predict c_{jt}, given current income y_{jt}. That is, if we estimate the regression

$$\ln c_{it} = \alpha_1 \ln c_{it-1} + \alpha_2 \ln y_{it} + \delta_v + \varepsilon_{it} \tag{6}$$

for those $i \in C_t$, limited commitment implies, since the households are constrained, $\alpha_1 = 0$: the old promises are forgotten. This test is implemented, and the results discussed, in Section 4.

3.4 Moral hazard

Under a moral hazard model, the agent must be given incentives to do something—exert effort, invest, or report output—which cannot be directly observed or contracted on. The "hidden action" moral hazard model includes the cases where the thing to be done occurs before output is realized and affects the expected level of output, e.g. effort or investment. As discussed in the introduction, this model has also been widely used to explain imperfect insurance in developing and developed countries. The Inverse Euler equation implication of moral hazard-constrained insurance (Rogerson 1985) has been used to test the moral hazard model against the PIH. I show below that a weaker implication, sufficiency of LIMU, holds under moral hazard even if the distribution of output depends on actions taken in past periods as well as the current period. Fernandes and Phelan (2000) show that when the distribution of income depends on past as well as current effort, two additional "threat-keeping" constraints are also added to the planner's problem. These constraints enforce an upper bound on a household's expected utility from today on if the household disobeys today's effort recommendation, whether they obey or disobey today. Even with such technological linkages between periods, we get the following implication of the moral hazard model:

Proposition 3 *When insurance is constrained by moral hazard, conditional on the time t shadow price of resources η_t, LIMU $\left(\frac{1}{u'(c_{i,t-1})}\right)$ is a sufficient statistic for household i's time t inverse marginal utility.*

Proof. *In appendix.* ■

Together these imply that, conditional on η_t, time $t-1$ inverse marginal utility is a sufficient statistic for all $t-1$ information for forecasting time t consumption. We get this result because the moral hazard-constrained model shares with the limited commitment model the feature that income is observed. As a result, the planner directly controls consumption and marginal utility. Moreover, the temptation preventing full insurance (reneging, shirking) is evaluated at the same level(s) of consumption and marginal utility that the household actually realizes[8]. Therefore, expected marginal utility can be expressed as a function of the past only via past marginal utility. Thus sufficiency of lagged inverse marginal utility holds under moral hazard as under limited commitment.

3.4.1 Testable implication

The fact that, under either limited commitment or moral hazard, the village's shadow cost of resources and LIMU should together be a sufficient statistic for the past in forecasting consumption implies that the prediction errors

$$\hat{\varepsilon}_{it} \equiv \frac{1}{v'(c_{it})} - \mathbb{E}\left(c_{it} | \frac{1}{v'(c_{i,t-1})}, \eta_t\right) \tag{7}$$

should be uncorrelated with past income, a finding that contrasts with the prediction of the hidden income model discussed below. When utility is CRRA,

$$\ln\left(\frac{1}{v'(c_{i,t-1})}\right) = \alpha \ln c_{i,t-1}$$

[8] This result does not rely on the assumption that utility is separable in consumption and leisure, since the principal knows what level of effort the agent will choose to exert.

and η_t is a village-year effect, so expected consumption $\mathbb{E}\left(c_{it}|\frac{1}{v'(c_{i,t-1})},\eta_t\right)$ is just the predicted value of c_{it} from the regression

$$c_{ivt} = \alpha \ln c_{iv,t-1} + \delta_{vt} + \varepsilon_{ivt}$$

and the residual is

$$\hat{\varepsilon}_{it} \equiv c_{it} - \hat{\alpha} \ln c_{iv,t-1} - \hat{\delta}_{vt}. \tag{8}$$

Ligon (1998) and Attanasio and Pavoni (2009) test for asymmetric information regarding agents' choice of actions (moral hazard) using GMM approaches. The test proposed here has the advantage of accommodating nonparametric estimates of the utility function, rather than requiring the specification of a functional form. Of course, in the event that the restrictions imposed by GMM are correct, they may provide more powerful tests, but such assumptions are difficult to test and may result in incorrect conclusions. The results of the regression-based test are discussed in Section 5.

3.5 Hidden income

As well as issues of commitment and of *ex ante* information (moral hazard), *ex post* informational asymmetries may also restrict the type of (implicit or explicit) contracts that agents can enter into, and thereby restrict insurance. Namely, it may be that income is not observable by the community, and households must be given incentives to report it. It turns out that such *ex post* informational asymmetries cause the sufficiency result of limited commitment and moral hazard to break down.

Assume now that agents can commit to the insurance arrangement—there are no *ex post* participation constraints. Potentially $S(S-1)$ incentive-compatibility constraints are added to the planner's problem:

$$\begin{aligned} v(y_r + \tau_{irt}) + \beta u_{ir,t+1} &\geq v(y_r + \tau_{ir',t}) + \beta u_{ir',t+1} \\ r' &\in \mathcal{S} \backslash y_r \end{aligned}$$

These constraints require that a household realizing any of the S income levels must not gain by claiming any of the $S-1$ other possible levels. Thomas and Worrall (1990) show that only the $S-1$ local downward constraints, which require that an agent getting income y_r not prefer to claim the slightly lower income y_{r-1}, will be binding at the optimum.

The first-order conditions of the problem imply:

Proposition 4 *When agents can commit to the insurance agreement, and effort is contractible, but output is hidden, forecasts using only* $\frac{1}{v'(y_r+\tau_{r,t-1})}$ *and* η_t *will over-predict consumption for households with the lowest time* $t-1$ *income realizations, and the degree of overprediction will decline with the level of time* $t-1$ *income (controlling for an interaction between time* $t-1$ *income and the aggregate shock* η_t*).*

Proof. *In appendix.* ■

The intuition for this result, as noted above, is that, in the limited commitment and moral hazard cases, unlike the hidden income case, the temptation of someone with high output to claim a lower level of output is not a relevant constraint, and so lagged inverse marginal utility is a sufficient statistic for forecasting current consumption. When income is private information, the optimal schedule of transfers and promised utilities distorts the tradeoff between current consumption and future expected utility, with households announcing low incomes being penalized more in terms of future utility than current consumption. Aggregate risk may matter because if the network receives a positive income shock, there is a potentially countervailing effect: all agents consume more than would have been predicted using past marginal utility.

In the limited commitment and moral hazard cases, unlike the hidden income case, the temptation of someone with high output to claim a lower level of output is not a relevant constraint, and so lagged inverse marginal utility is a sufficient statistic for forecasting current consumption, including in the case of aggregate risk, because last period's promise is the only past information which determines how the aggregate shock is divided among the households.

Of course, utility promises u_{irt} are unobserved objects, but at the optimum, increasing τ_{irt} or $u_{ir,t+1}$ must be equally effective in satisfying promise-keeping constraints, so promises are monotonically increasing in past output (as is time t consumption $y_r + \tau_{rt}$, although transfers τ_{rt} are decreasing in output, as they must be to provide partial insurance to those with low incomes (Thomas and Worrall 1990)). So if the residuals defined in (7) are regressed on lagged income:

$$\hat{\varepsilon}_{it} = \alpha + \beta y_{i,t-1} + u_{it} \tag{9}$$

hidden income predicts $\alpha < 0, \beta > 0$.

Equation (9) suggests a way to distinguish the hidden income model from moral hazard/limited commitment: finding $\alpha < 0, \beta > 0$ (overprediction at the bottom) is evidence in favor of hidden income. If we are unable to reject $\alpha = \beta = 0$, this is evidence for either limited commitment, moral hazard or autarky, which can then be distinguished based on the amnesia test discussed above and the inverse Euler implication of moral hazard.

3.5.1 An additional implication of hidden income: insufficiency of LIMU is less when income is less variable

An additional prediction of the hidden income model is that a reduction in the variability of a household's income process will have the effect of making truth-telling constraints less binding, which in turn implies a reduced wedge between LIMU and expected promised utility.

Proposition 5 *A decrease in variability of the income process (in the sense of that the new distribution is second-order stochastically dominated by the old distribution, keeping the probability of each income realization the same) reduces the degree to which LIMU over-predicts current inverse marginal utility for low-lagged income households.*

Proof. *In appendix.* ■

More predictable income reduces the degree of overprediction at the bottom

3.6 Summary: Distinguishing barriers to insurance

The preceding discussion suggests four tests that, in combination, can distinguish among limited commitment, hidden income, and autarky:

1. Sufficiency of $\frac{1}{v'(c_{i,t-1})}$: under limited commitment, moral hazard, or borrowing-lending (PIH), $\mathbb{E}\left(\frac{1}{v'(c_{it})}|\frac{1}{v'(c_{i,t-1})}, \eta_t, x_{i,t-s}\right) = \mathbb{E}\left(\frac{1}{v'(c_{it})}|\frac{1}{v'(c_{i,t-1})}, \eta_t\right), \forall x_{i,t-s}, s > 0$

2. Amnesia: under limited commitment, $\mathbb{E}\left(\frac{1}{v'(c_{it})}|\frac{1}{v'(c_{i,t-1})},\eta_t,y_{it}\right)=\mathbb{E}\left(\frac{1}{v'(c_{it})}|\eta_t,y_{it}\right)$ if household i is constrained

3. Overprediction at the bottom: under hidden income, when $a_t \leq a_{t-1}$,

$$\mathbb{E}\left(\left[\frac{1}{v'(c_{it})}-\mathbb{E}\left(\frac{1}{v'(c_{it})}|\frac{1}{u'(c_{i,t-1})},\eta_t\right)\right]\Bigg|y_{i,t-1}=0\right) < 0$$

and

$$\frac{d}{dy_{i,t-1}}\left(\frac{1}{v'(c_{it})}-\mathbb{E}\left(\frac{1}{v'(c_{it})}|\frac{1}{u'(c_{i,t-1})},\eta_t\right)\right) > 0$$

4. Inverse Euler equation: under moral hazard

$$\frac{1}{v'(c_{i,t-1})}=\mathbb{E}_{t-1}\left(\frac{1}{v'(c_{it})}\right)$$

These tests are summarized in the following table:

	Autarky/PIH	Limited com.	Moral hazard	Hidden inc.
Sufficiency of $\ln c_{t-1}$	✓	✓	✓	
Amnesia		✓		
Overprediction at the bottom				✓
Inverse Euler			✓	

4 Data and Estimation

4.1 Data

Data are from the 1999-2005 waves of the Townsend Thai Monthly Survey, which covers 16 villages in central and northeastern Thailand, 4 each in four provinces, two in the central region near Bangkok and two in the northeast. In each village, 45 households were initially selected at random and reinterviewed each year. (See Townsend et al. (1997) for details.) Detailed data were collected on households' demographic composition and their income, including farms, businesses, and wage employment. Information was also collected on household expenditure, using detailed bi-weekly and monthly surveys. Thus expenditure is likely to be quite well-measured in this dataset, relative to datasets which measure expenditure over a longer recall period and/or which collect information on only a subset of expenditures.

A total of 531 households appear in all 84 months of the survey period used here, out of an original 670 who were interviewed in January 1999. I focus on the continuously-observed sample so that changes in a household's rank in the PCE distribution are not due to migration in and out of the survey. Smaller households and those whose head is engaged in rice farming or construction are most likely not to be continuously observed, while corn and livestock farmers are more likely to be continuously observed. This degree of missing data is a concern; however, imputing income and expenditure data for missing household-months based on village, year, occupation and baseline demographic variables, and running the analysis on this sample, yields results similar to the results for the continuously-observed sample.[9]

Summary statistics are reported in Table 3. Average reported monthly per capita expenditure was 5,213 2002 baht (approximately 124 2002 US dollars[10]; all following references to baht refer to 2002 baht). Average reported monthly income per capita is higher than expenditure, 8,981 baht, reflecting investment (and also perhaps savings).

Another strength of the Townsend Thai Monthly Survey data is that households are asked separately about gifts and transfers (both in money and in-kind) from organizations, from households in the village, and from households outside of the village. All of these types of transfers are prevalent: gifts given to other households in the same village equal 5.4% of average expenditure, while gifts from others in the same village equal 9% of average expenditure. Gifts/remittances given to those outside the household's village equal 17.5% of average expenditure, and gifts/remittances received from those outside the village equal 27.7% of average expenditure. Moreover, these numbers exclude transfers embodied in interest-free, low-cost and flexible loans, which are prevalent in these villages, as well as in other settings ((Platteau and Abraham 1987), (Udry 1994), (Fafchamps and Lund 2003)) The significant magnitude of intra-village transfers is direct evidence that within-village insurance is important, while transfers made with those outside the village may constitute

[9] Results available on request.
[10] The exchange rate in 2002 was approximately 42 baht=$1.

a source of unobserved income.

Finally, using data from rain gauges located in each village, yielding a measure of total rainfall in each village in each month between 1999 and 2003, quarterly rainfall variables (deviations from the provincial average in that quarter over the entire period) were constructed following Paxson (1992):

$$R_{qvt} - \bar{R}_{qp}, \left(R_{qvt} - \bar{R}_{qp}\right)^2,$$
$$q = 1, 2, 3, 4$$

These variables will be used for tests of the hidden income model reported below.

4.2 Testing sufficiency of lagged inverse marginal utility

Under limited commitment, moral hazard, or autarky, consumption and investment should only depend on the past through $\frac{1}{v'(c_{i,t-1})}$. If households' consumptions are described by efficient, limited-commitment-constrained insurance, we should find $\alpha \neq 0, \beta = 0$ in

$$\ln c_{it} = \alpha \ln c_{i,t-1} + \beta' X_{i,t-1} + \delta_v + \varepsilon_{it}$$

where $X_{i,t-s}$ is any information dated $t-1$ or before. Table 4 presents the results of this test. While lagged inverse marginal utility (i.e. lagged consumption) is significantly predictive of current per capita consumption (col 1), lagged log income is also a significant predictor of consumption ($p < .001$) in the full sample (col. 2). The result is unchanged when the top and bottom 5% of per capita expenditure (by year) are dropped (col 4). This suggests that neither limited commitment or moral hazard alone can explain the failure of full insurance in these villages.

4.3 Testing amnesia

Table 5 presents tests of the amnesia prediction of the limited commitment model. If there is measurement error in expenditure, following Kocherlakota's procedure exactly—classifying as constrained every household in a village who had consumption growth above the village minimum—would result in every household but one in each village appearing constrained.

In fact, many of these households will be unconstrained, and including them in the set of households for whom amnesia is predicted will introduce bias toward rejecting the predictions of limited commitment. To address this, in columns 1 through 4, interaction terms between $\ln \frac{1}{v'(c_{i,t-1})}$ and indicators for the quartile of the village distribution of consumption growth between $t-1$ and t into which the household fell; and similar interaction terms with $\ln(y_{i,t})$ are added to (6). That is, estimate

$$\ln c_{it} = \alpha + \beta_1 \ln c_{i,t-1} + \sum_{q=2}^{4} \beta_q \ln c_{i,t-1} \times \mathbf{1}_q + \gamma_1 \ln y_{i,t} + \sum_{q=2}^{4} \gamma_q \ln y_{i,t} \times \mathbf{1}_q + \delta_v + \varepsilon_{it}$$

If past promises are forgotten, conditional on current income, for those who had highest consumption growth because these households had binding participation constraints, the sum of the coefficients on the interaction terms $(\beta_1 + \beta_q + \gamma_1 + \gamma_q)$ should be low and insignificant for higher quartiles of consumption growth and, since the main effect of $\ln \frac{1}{v'(c_{i,t-1})}$ is positive and significant $\beta_1 + \beta_4 + \gamma_1 + \gamma_4$ should be negative. In fact, these predictions are rejected—LIMU is strongly predictive of current consumption, conditional on current income, for households with the highest consumption growth, and the hypothesis that $\beta_1 + \beta_4 + \gamma_1 + \gamma_4$ equals zero is overwhelmingly rejected—suggesting again that limited commitment is not the (entire) explanation for incomplete insurance in these villages.

As a second test, columns 5 and 6 estimate (6) for households with above-median consumption growth, separately for villages where the variability of rainfall from year to year is high and villages where rainfall variability is low, based on monthly rainfall data from 1951-1985. Villages with high rainfall variance also had higher average income variance in every year but 2004, when the opposite is true—see Figure 1. If measurement error in expenditure is independent of the variance in incomes, then when high consumption growth is observed in high-rainfall-variance (HRV) villages, it is more likely to be due to a high income realization resulting in a binding participation constraint. In low-rainfall-variance (LRV) villages, high consumption growth is more likely to be due to measurement error. This suggests that, if limited commitment is the true model, the amnesia prediction should do better in HRV villages, i.e. the coefficient on $\ln \frac{1}{v'(c_{i,t-1})}$ in column 6 should be less than in column 5. In fact, the point estimate for HRV villages is lower than for LRV villages, and the two estimates are not statistically different ($p = .66$). Therefore, both the sufficiency and amnesia predictions of the limited commitment model are strongly rejected.

4.4 Testing hidden income: insufficiency of LIMU and predicitve power of income

Table 6 presents the results of the tests that under hidden income LIMU will overpredict consumption for those households whose promises decreased, i.e. who had low income in the previous period, while under moral hazard or limited commitment, the prediction errors will be uncorrelated with last-period income because LIMU is a sufficient statistic for history, hence no additional lagged information will contain predictive power. Consistent with the hidden income prediction, when the prediction errors (8) are regressed on lagged income (and lagged income and lagged income squared interacted with the aggregate shock measure η_t) the slope is positive and significant while the intercept is significantly negative (column 1). Since the dependent variable is a regression residual, which has mean zero by construction, the slope and intercept are not independent tests. The joint hypothesis that $\alpha = 0, \beta = 0$ is rejected at the .0001 level. Column 2 repeats this test without the aggregate shock interaction terms, to show that the overprediction result in fact holds unconditionally; i.e., the potential countervailing effect of increased aggregate resources does not undo the overprediction result. Again, the joint hypothesis that $\alpha = 0, \beta = 0$ is rejected at the .0001 level.

Therefore, there is evidence that hidden income constraints cause those with low past income to receive less current consumption than predicted by LIMU. This suggests that insurance is constrained by the need to provide incentives to high-income households to truthfully reveal that income. Robustness checks, discussed in Section 5, suggest that this finding is not driven by measurement error or misspecification. However, before discussing these tests, I present two tests of the prediction that households with easier-to-predict income processes should display less departure from sufficiency of LIMU.

4.5 Testing hidden income: departure from sufficiency and predictive power of rainfall

If the primary barrier to insurance is the inability of the community to directly observe households' incomes, and this barrier is manifested through insufficiency of LIMU, households whose income processes are less uncertain, because they are predicted by observed factors, or simply less variable, should display less insufficiency of LIMU.

As a first test of this prediction, I regressed income on the rainfall variables $R_{qvt} - \bar{R}_{qp}$ and $\left(R_{qvt} - \bar{R}_{qp}\right)^2$ separately for households in each of 10 occupational categories. The R^2 from this regression was interacted with lagged income. (The R^2s appear in Table A1.) Table 7a shows the results of regressing the prediction errors (8) on lagged income and its interaction with the occupation-specific rainfall R_o^2:

$$\hat{\varepsilon}_{iot} = \alpha + \beta y_{io,t-1} + \gamma y_{io,t-1} \times R_o^2 + \delta R_o^2 + u_{it}$$

If insufficiency of LIMU is reduced when a household's income is easier to forecast, $\gamma < 0, \delta > 0$.In fact, this is the case: there is less insufficiency of LIMU (in the sense of less correlation of the residuals with lagged income), when rainfall R^2 is higher, and the parameter estimates imply that if $R^2 = 1$ (income were perfectly predicted by rainfall), insufficiency of LIMU would be completely eliminated.

As a second test, for each household, I calculate variance of income, after removing the component of income predicted by the rainfall variables and occupation-year dummies; i.e. that part which should be difficult to forecast. I split the sample according to whether this variance is above or below-median. The prediction here is that there should be less insufficiency of LIMU for the low-variance sample. Table 7, Panel B shows the results. Both in terms of the point estimates and the chi-squared test of joint significance, the high-variance sample displays greater insufficiency of LIMU: $\alpha^{high} < \alpha^{low}, \beta^{high} < \beta^{low}$, and $\chi^2_{high} > \chi^2_{low}$.

5 Robustness checks

In practice, two further issues have to be considered in distinguishing among different insurance regimes: agents' utility functions are not known, and consumption is measured with error. Both of these, if not accounted for, can result in biased inference about the nature of the barrier to full insurance.

5.1 Measurement Error in Expenditure

5.1.1 Classical measurement error

If expenditure is measured with classical error, its coefficients in (9) will be attenuated toward zero. This will result in biased predictions of consumption using LIMU. To see what form the bias will take, note that we want to estimate the part of consumption that is unexplained by LIMU or a village-year effect:

$$\varepsilon_{ivt} = \ln c_{ivt} - \delta_{vt} - \beta \ln c_{iv,t-1}$$

Assume an error-ridden measure of consumption is observed,

$$\tilde{c}_{iv,t-1} = c_{iv,t-1} \cdot \nu_{iv,t-1}$$

where the measurement error $\nu_{iv,t-1}$ is uncorrelated with true time $t-1$ consumption, $c_{iv,t-1}$, or true time t consumption, $c_{iv,t}$.

The estimated prediction error is constructed using observed lagged consumption $\tilde{c}_{iv,t-1}$, and the estimates of β and δ:

$$\hat{\varepsilon}_{ivt} = \ln c_{ivt} - \hat{\delta}_{vt} - \hat{\beta} \ln \tilde{c}_{iv,t-1}$$

Assume the true DGP is limited commitment or moral hazard, that is

$$corr(\ln c_{ivt} - \delta_{vt} - \beta \ln c_{iv,t-1}, y_{iv,t-1}) = 0$$

so that, since $\mathbb{E}\nu_{iv,t-1}y_{iv,t-1} = 0$,

$$corr(\ln c_{ivt} - \delta_{vt} - \beta \ln \tilde{c}_{iv,t-1}, y_{iv,t-1}) = 0.$$

Will we correctly not reject the null of no correlation? Potentially not, because $\hat{\beta}$ is biased downward:

$$\begin{aligned}
p\lim \hat{\beta} &= \beta \left(1 - \frac{\sigma_\nu^2}{\sigma_c^2 + \sigma_\nu^2} \right) \\
\hat{\varepsilon}_{ivt} &= \ln c_{ivt} - \hat{\delta}_{vt} - \hat{\beta} \ln \tilde{c}_{iv,t-1} \\
&= \underbrace{\ln c_{ivt} - \hat{\delta}_{vt} - \beta \ln \tilde{c}_{iv,t-1}}_{\text{uncorrelated w/ } y_{iv,t-1}} + \underbrace{\ln \tilde{c}_{iv,t-1} \frac{\sigma_\nu^2}{\sigma_c^2 + \sigma_\nu^2}}_{\text{+ correlated w/ } y_{iv,t-1}}
\end{aligned}$$

That is, we may conclude wrongly that $corr(\hat{\varepsilon}_{ivt}, y_{iv,t-1}) > 0$ when consumption is measured with classical error.

However, for classical error, there is a straightforward solution. If β is estimated using the second lag of consumption as an instrument for the first lag, we obtain a consistent estimate:

$$\begin{aligned}
p\lim \hat{\beta}_{IV} &= \frac{cov(\ln \tilde{c}_{iv,t-2}, \ln \tilde{c}_{ivt})}{cov(\ln \tilde{c}_{iv,t-2}, \ln \tilde{c}_{iv,t-1})} \\
&= \beta \left(1 - \frac{cov(\nu_{t-2}, \nu_{t-1})}{cov(\ln \tilde{c}_{iv,t-2}, \ln \tilde{c}_{iv,t-1})} \right)
\end{aligned}$$

If $cov(\nu_{t-2}, \nu_{t-1}) = 0$ (classical measurement error) and the true DGP is limited commitment or moral hazard, we will correctly conclude

$$corr(\hat{\varepsilon}_{ivt}^{IV}, y_{iv,t-1}) = 0.$$

Columns 3 and 4 of table 6 show that instrumenting $\ln c_{iv,t-1}$ with $\ln c_{iv,t-2}$ does not overturn the finding that the prediction residuals are negative at low levels of lagged income: the null that the slope and the intercept in (9) are both 0 is rejected at the 1% level.

5.2 Non-classical measurement error

Using longer lags of consumption as an instrument will not address measurement error which is correlated over time. Because total consumption is extrapolated from components of expenditure, correlated measurement error is not implausible: if in all periods good x is a larger or smaller fraction of household i's expenditure than for the average household in the Socio-Economic Survey, this may introduce measurement error which is correlated over time. A possible solution in this case is to move lagged consumption from the right- to the left-hand side of the equation of interest, and test overidentifying restrictions on the reduced form equations for $\ln c_{it}$ and $\ln c_{i,t-1}$. If lagged income affects current consumption only through lagged consumption, then all components of lagged income, or any other lagged information $x_{i,t-s}$ which predicts lagged income, should satisfy the restriction

$$\frac{d\ln c_{it}}{dx_{i,t-s}} / \frac{d\ln c_{i,t-1}}{dx_{i,t-s}} = K, \forall x_{i,t-s}$$

That is, a unit change in the instrument $x_{i,t-s}$ should have the same relative effect on current versus lagged consumption as a unit change in another instrument $x'_{i,t-s}$.

Under the null of limited commitment/moral hazard, consumption depends on a household's initial Pareto weight and its subsequent income realizations. (While under limited commitment or moral hazard, lagged income does not belong in the structural equation for consumption, it appears in the reduced form because y_{is} depends on c_{is}.) Three lags of income are significant predictors of c_{it}, so write

$$\ln c_{it} = \sum_{s=1}^{3} \alpha_s y_{i,t-s} + \hat{\lambda}_0 + \varepsilon_{it}$$

where $\hat{\lambda}_0$ is a measure of the household's Pareto weight as of 1999: the household's rank in the 1999 per-capita consumption distribution for the village.

Since lagged income appears in the reduced form for consumption, lags of total income cannot be used to generate overidentifying restrictions. Instead, I test whether the *composition* of lagged income matters for predicting current consumption, beyond its effect on lagged consumption. In particular, while different types of income (earned income versus remittances) may convey different information about effort, or different information about the household's prospects in autarky, under limited commitment or moral hazard that information will be completely encoded in lagged consumption. Under hidden income, in contrast, the components of income will also matter through the direct effect of lagged income on current consumption. So in the reduced-form regressions

$$\begin{aligned}
\ln c_{it} &= \sum_{s=1}^{3}[\pi_{1Es} y_{earned,i,t-s} + \pi_{1Rs} y_{remit,i,t-s}] + \hat{\lambda}_{i0} + \varepsilon_{it} \\
\ln c_{i,t-1} &= \sum_{s=1}^{3}[\pi_{2Es} y_{earned,i,t-s} + \pi_{2Rs} y_{remit,i,t-s}] + \hat{\lambda}_{i0} + \varepsilon_{i,t-1}
\end{aligned}$$

if lagged income does not directly affect current consumption, we should find $\frac{\pi_{1Es}}{\pi_{2Es}} = \frac{\pi s}{\pi_{2Rs}}, s = 1, 2, 3$. These 3 overidentifying restrictions can be used to test whether the rejection of limited commitment is only due to measurement error. *[Results to come.]*

5.3 Specification of u()

The test of hidden income proposed above is to test whether $\varepsilon_t \perp y_{t-1}$ in

$$\ln\left(\frac{1}{v'(c_t)}\right) = \delta_t + \ln\left(\frac{1}{v'(c_t)}\right) + \varepsilon_t \tag{10}$$

However, since the form of $v()$ is unknown, the approach above was to approximate it with the isolelastic function

$$v(c_{it}) = \frac{c^{1-\rho}}{1-\rho}$$

and test $\hat{\varepsilon}_t \perp y_{t-1}$ in

$$\ln(c_{it}) = \delta_{vt} + \ln(c_{t-1}) + \hat{\varepsilon}_t \tag{11}$$

This raises the question, if the true error ε_t satisfies $\varepsilon_t \perp y_{t-1}$ in (10), will testing $\hat{\varepsilon}_t \perp y_{t-1}$ in (11) yield the correct conclusion? Nonparametrically estimating $\frac{1}{v'(c)}$ avoids the need to make an assumption about the form of the utility function. In order to correct for measurement error as well, a nonparametric IV approach seems most appropriate.

One possible approach is the nonparametric 2SLS approach of Newey and Powell (2003) to estimate

$$f\left(c_{it}\right)=\delta_{vt}+f\left(\tilde{c}_{t-1}\right)+\hat{\varepsilon}_t$$

where $\tilde{c}_{t-1}$ is estimated using a nonparametric first stage with c_{t-2} as an instrument. However, consistency of this estimator requires that $f()$ and its derivatives are bounded in the tails, if $\tilde{c}_{-1}$ is not bounded. Since in this context $f()$ is a marginal utility function which may tend to infinity as consumption tends to zero, this is an unappealing assumption. ,Newey and Powell's approach also requires the conditional mean zero assumption:

$$\mathbb{E}\left(\varepsilon_t|\tilde{c}_{-2}\right)=0$$

which is stronger than the assumption needed for linear IV:

$$corr\left(\varepsilon_t,\tilde{c}_{-2}\right)=0$$

Fortunately, inspection of the nonparametric first stage between $\ln(\tilde{c}_{-1})$ and $\ln(\tilde{c}_{-2})$ shows it to be nearly linear, suggesting that a linear first stage may be a suitable approach. That is, the strategy is to nonparametrically estimate $f()$, using a 5-knot spline[11], in

$$\ln\left(\tilde{c}_t\right)=\eta_{vt}+f\left(\hat{c}_{t-1}\right)+\tilde{\varepsilon}_t$$

where $f\left(\hat{c}_{t-1}\right)$ is instrumented with $f\left(\hat{c}_{t-2}\right)$.

Table 9 shows that the results when $\frac{1}{v'(c)}$ is estimated nonparametrically, sufficiency of LIMU is once again rejected, at the 1% level in the OLS regression and at the 5% level in the IV regression.

6 Conclusion

This paper suggests a method that can be used to test whether either of two workhorse models of endogenously incomplete insurance, limited commitment or moral hazard, is consistent with the relationship between consumption and other information in long panel

[11] Results are not sensitive to the number of knots used. (Results available on request.)

data. If information from "the past" helps to forecast current consumption, conditional on one lag of inverse marginal utility, neither limited commitment or moral hazard can fully explain incomplete insurance. However, if a household's past income helps to forecast current consumption, in the sense that prediction errors ignoring past income are positive when past income was low, this is consistent with a model in which household's cannot directly observe one another's income.

Measurement error in right-hand side variables, which is commonly seen as a threat to power (causing underrejection of the null), is a particular concern with tests of this type, because mismeasurement of the proposed sufficient statistic (here, lagged inverse marginal utility) can distort the size of the test, causing *over*rejection of the null, if those variables which are excluded under the null hypothesis are correlated with the true value of the proposed sufficient statistic. This concern is addressed here with instrumental variables and by testing overidentifying restrictions on the reduced forms for the left- and right-hand-side variables.

Results from a 7-year monthly panel of households in rural Thailand are inconsistent with pure moral hazard or limited commitment, and suggest that hidden income plays a role in constraining households from achieving full risk sharing. This suggests that policies which make it easier (harder) for villagers to infer one another's incomes may improve (worsen) risk sharing. Changes that improve observability of income could include dissemination of crop or other price information; changes that worsen observability could include access to larger, anonymous markets; diversification of occupations within a village; electronic payments of remittances or for business transactions; seasonal migration; and private rather than group banking. Since policies that have the potential to worsen observability of income may also raise the average *level* of income, this is not to suggest that such policies be avoided. However, when possible they should be designed with consideration of the consequences for informal insurance in settings where formal insurance is limited or absent.

References

Ai, H. and F. Yang (2007, August). Private information, limited commitment, and risk sharing.

Albarran, P. and O. P. Attanasio (2003). Limited commitment and crowding out of private transfers: evidence from a randomised experiment. *The Economic Journal 113*, C77–C85.

Attanasio, O. and N. Pavoni (2007, March). Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption. NBER Working Paper 12994.

Attanasio, O. and J.-V. Rios-Rull (2000). Consumption smoothing in island economies: Can public insurance reduce welfare? *European Economic Review 44*, 1225–1258.

Blundell, R., L. Pistaferri, and I. Preston (2008, February). Consumption inequality and partial insurance.

Cashdan, E. A. (1985). Coping with risk: Reciprocity among the Basarwa of Northern Botswana. *Man 20*, 454–474.

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

Cox, D., Z. Eser, and E. Jimenez (1998). Motives for private transfers over the life cycle: An analytical framework and evidence for Peru. *Journal of Development Economics 55*, 57–80.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica 59*, 1221–1248.

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

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

Fernandes, A. and C. Phelan (2000). A recursive formulation for repeated agency with history dependence. *Journal of Economic Theory 91*, 223–247.

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

Grossman, S. J. and O. D. Hart (1983). An analysis of the principal-agent problem. *Econometrica 51*, 7–45.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy 86*, 971–987.

Hayashi, F., J. Altonji, and L. Kotlikoff (1996). Risk-sharing between and within families. *Econometrica 64*, 261–294.

Kaplan, G. (2006). The cross-sectional implications of incomplete markets: Private information or limited enforcement? NYU Working Paper.

Karaivanov, A. and R. M. Townsend (2008, September). Firm dynamics and finance: Distinguishing information regimes. Manuscript.

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

Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *Review of Economic Studies 63*, 595–609.

Koeppl, T. V. . (2006). Differentiability of the efficient frontier when commitment to risk sharing is limited. *Topics in Macroeconomics 6*, 1–8.

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

Ligon, E., J. P. Thomas, and T. Worrall (2000). Mutual insurance, individual savings, and limited commitment. *Review of Economic Dynamics 3*, 216–246.

Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies 69*, 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*, 59–118.

Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics 12*, 380–391.

Morduch, J. (1995). Income smoothing and consumption smoothing. *Journal of Economic Perspectives 9*, 103–114.

Olken, B. A. (2005). Revealed community equivalence scales. *Journal of Public Economics 89*, 545–566.

Paulson, A. L., R. M. Townsend, and A. Karaivanov (2006). Distinguishing limited liability from moral hazard in a model of entrepreneurship. *Journal of Political Economy 114*, 100–144.

Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in thailand. *American Economic Review 82*, 15–33.

Platteau, J.-P. (1991). *Social security in developing countries*, Chapter Traditional Systems of Social Security and Hunger Insurance: past achievements and modern challenges, pp. 112–163. Oxford: Clarendon Press.

Platteau, J.-P. and A. Abraham (1987). An inquiry into quasi-credit contracts: The role of reciprocal credit and interlinked deals in small-scale fishing communities. *Journal of Development Studies 23*, 461–490.

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

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

Rosenzweig, M. R. and K. I. Wolpin (1993). Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in India. *Journal of Political Economy 101*, 223–244.

Scott, J. C. (1976). *The Moral Economy of the Peasant*. Oxford: Basil Blackwell.

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

Suri, T. (2005, September). Spillovers in village consumption: Testing the extent of partial insurance.

Thomas, J. and T. Worrall (1990). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory 51*, 367–390.

Townsend, R. M. (1982). Optimal multiperiod contracts and the gain from enduring relationships under private information. *Journal of Political Economy 90*, 1166–1186.

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

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

Townsend, R. M. (2007a). Townsend Thai project household annual resurvey, 1998-2006. [Computer file]. 2nd Release. Chicago, IL: Social Sciences Computing Services, the University of Chicago [Producer & distributor].

Townsend, R. M. (2007b). Townsend Thai project initial household survey, 1997. [Computer file]. 2nd Release. Chicago, IL: Social Sciences Computing Services, the University of Chicago [Producer & distributor].

Townsend, R. M. and R. A. E. Mueller (1998). Mechanism design and village economies: From credit to tenancy to cropping groups. *Review of Economic Dynamics 1*, 119–172.

Townsend, R. M., A. L. Paulson, S. Sakuntasathien, T. J. Lee, and M. B. 1997. (1997). *Questionnaire Design and Data Collection for NICHD Grant 'Risk, Insurance and the Family' and NSF Grants.* Univ. Chicago.

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

7 Appendix 1: Proofs

Define the N-dimensional vector of household incomes at $t, h_t = \{y_{it}\}_{i=1}^N$, and the history $(h_1, ..., h_t) \equiv h^t$.

7.1 Proof of Proposition 1: Full insurance rules out rank-reversals and dependence of consumption on income

Let λ_{it} be the multiplier on household i's promise-keeping constraint, and η_t be the multiplier on the village's time t budget constraint. Solving (OBJ) subject to the promise-keeping constraints (λ_{it}) and the village's budget constraint (η_t) yields the following first-order conditions for transfers, promised utility, and assets:

Proof. The FOCs are ■

$\tau_{it}(h^t)$:

$$\eta_t(h^t) = \lambda_{it} \Pr(h^t) v'(y_{it} + \tau_{it}(h^t)) \tag{12}$$

$u_{i,t+1}(h_t)$:

$$\Pr(h^t) \frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} = -\Pr(h^t)\lambda_{it}, \forall h^t, i < N \tag{13}$$

a_{t+1} :

$$\Pr(h^t) \frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial a_{t+1}(h^t)} = \eta_t(h^t) \tag{14}$$

and the envelope conditions:

$$\frac{\partial u_N(\mathbf{u}_t(h^{t-1}), a_t(h^{t-1}), \mathbf{e}')}{\partial u_{i,t}(h^{t-1})} = -\lambda_{it}, \forall i < N \tag{15}$$

$$\frac{\partial u_N(\mathbf{u}_t(h^{t-1}), a_t(h^{t-1}), \mathbf{e}')}{\partial a_t(h^{t-1})} = \eta_{t-1}(h^{t-1}) \tag{16}$$

The FOCs for transfers for households i and N imply

$$\frac{\lambda_{Nt}}{\lambda_{it}} = \frac{v'(y_{it} + \tau_{it}(h^t))}{v'(y_{Nt} + \tau_{Nt}(h^t))}$$

So that

$$c_{it} \equiv y_{it} + \tau_{it}(h^t) = v'^{-1}\left(\frac{\lambda_{Nt}}{\lambda_{it}} v'(y_{Nt} + \tau_{Nt}(h^t))\right) \tag{17}$$

Substituting into the law of motion for assets,

$$R^{-1}a_{t+1} = a_t + \sum_{i=1}^{N} y_{it} - \sum_{i=1}^{N} v'^{-1}\left(\frac{\lambda_{Nt}}{\lambda_{it}} v'(y_{Nt} + \tau_{Nt}(h^t))\right) \tag{18}$$

which is a single equation in c_{Nt}, i.e. c_{Nt} depends only on the aggregate endowment, and not on h^t or $\{y_{it}\}$. Then (17) implies that for all households, c_{it} depends only on the aggregate endowment.

Using (13) and (15), $\lambda_{it} = \lambda_{it+1} = \lambda_i, \forall i, t$.

Further, for all i, j in the network:

$$\begin{aligned} \frac{v'(y_r + \tau_{Nrt})}{v'(y_r + \tau_{irt})} &= \lambda_i, \forall r, t, i < N \\ \frac{v'(y_r + \tau_{jrt})}{v'(y_r + \tau_{irt})} &= \frac{\lambda_j}{\lambda_i} \end{aligned}$$

So if in the first period, household i consumes more than household j, this will be the case in all subsequent periods, and vice versa. Therefore under full insurance the ordering of initial multipliers λ_{i0} or equivalently initial promises u_{i0} will determine the ordering of household i in the consumption distribution in all periods.■

7.2 Proof of Proposition 2: Lagged inverse marginal utility is a sufficient statistic under limited commitment

Let $\frac{\lambda_{it}}{\Pr(h^t)}$ be the multiplier on household i's promise-keeping constraint, and $\frac{\eta_t(h^t)}{\Pr(h^t)}$ be the multiplier on the village's time t budget constraint after history h^t. Using the stationarity of the problem, $\Pr(h^t|\mathbf{u}(h^{t-1}), a(h^{t-1}), \mathbf{e}) = \Pr(h^t|h^{t-1}) = \Pr(h_t)$, so probabilities are written conditional only on the time t realization h_t. Let $\phi_{it}(h^t)$ be the multiplier on household i's participation constraint.

Assume that there is at least one realization h_t such that no household's participation constraint is binding: this guarantees differentiability of the planner's value function (Koeppl 2006). Solving (OBJ) subject to the promise-keeping constraints (λ_{it}), the participation constraints (5) and the village's budget constraint (η_t) yields the following first-order conditions for transfers, promised utility, and assets:

$\tau_{it}(h^t)$:

$$\eta_t(h^t) = (\lambda_{it} + \phi_{it}(h^t))v'(y_{it} + \tau_{it}(h^t)) \tag{19}$$

$u_{i,t+1}(h^t)$:

$$\Pr(h^t)\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} = -\Pr(h^t)\lambda_{it} - \phi_{it}(h^t), \forall h^t, i < N \tag{20}$$

$a_{t+1}(h^t)$:

$$\Pr(h^t)\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial a_{t+1}(h^t)} = \eta_t(h^t) \tag{21}$$

and the envelope conditions for current promises (15) and assets (16):

$$\frac{\partial u_N(\mathbf{u}_t(h^{t-1}), a_t(h^{t-1}), \mathbf{e}')}{\partial u_{i,t}(h^{t-1})} = -\lambda_{it}, \forall i < N$$

$$\frac{\partial u_N(\mathbf{u}_t(h^{t-1}), a_t(h^{t-1}), \mathbf{e}')}{\partial a_t(h^{t-1})} = \eta_{t-1}(h^{t-1})$$

It will be helpful to use the following result:

Lemma 6 *The double (y_{it}, η_t) is a sufficient statistic for the N-vector of income realizations h^t in determining household i's transfer: $\tau_{it}(h^t) = \tau_{it}(y_{it}, \eta_t)$*

Proof. Note that, when $\phi_{it}(h^t) > 0$, i.e. household i's participation constraint is binding, (19) and (15) imply that the household's transfer and future promise are set to make the household exactly indifferent between staying in the network or defaulting, and to equate the cost of providing the current transfer τ and future promise u, irrespective of the income realizations of other households in the network:

$$v(y_r + \tau_{it}(h^t)) + \beta u_{i,t+1}(h^t) = u^t_{aut}(y_r)$$

$$v'(y_r + \tau_{it}(h^t)) = -\left(\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)}\right)^{-1}$$

so $\tau_{it}(h^t|\phi_{it}(h^t) > 0) = \tau_{it}(y_{it}, \eta_t)$. And, when $\phi_{it}(h^t) = 0$, i.e. household i's participation constraint is not binding, (19) and (15) imply that $\frac{1}{v'(y_{it}+\tau_{it}(h^t))} = \frac{\eta_t(h^t)}{\lambda_{it}}$, so, again, $\tau_{it}(h^t|\phi_{it}(h^t) = 0) = \tau_{it}(y_{it}, \eta_t)$. ■

This allows us to write $\tau_{it}(y_{it}, \eta_t)$ for $\tau_{it}(h^t)$. Using the FOCs for $\tau_{it}(y_{it}, \eta_t)$ and $u_{i,t+1}(h^t)$:

$$\begin{aligned}
\eta_t(h^t) &= \Pr(h_t)\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} v'(y_{it} + \tau_{it}(y_{it}, \eta_t)) \\
&= \Pr(y_{it}, \eta_t) v'(y_{it} + \tau_{it}(y_{it}, \eta_t)) \Pr(h_t|y_{it}, \eta_t) \frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)}
\end{aligned}$$

since $\Pr(y_{it}, \eta_t) \Pr(h_t|y_{it}, \eta_t) = \Pr(h_t \cap (y_{it}, \eta_t)) = \Pr(h_t \cap (\eta_t(h^t)))$. This says that inverse marginal utility, weighted by the shadow price of resources scaled by the probability of (y_{it}, η_t), is equal to the gradient of the planner's value function with respect to household i's time $t+1$ promised utility weighted by the probability of the N-vector of income realizations h_t, given (y_{it}, η_t):

$$\frac{\eta_t(h^t)}{\Pr(y_{it}, \eta_t) v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} = \Pr(h_t | y_{it}, \eta_t) \frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} \tag{22}$$

Note that

$$\sum_{h_t | \Pr(h_t|y_{it},\eta_t)>0} \left(\frac{\eta_t(h^t)}{\Pr(y_{it}, \eta_t) v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} \right) = \frac{\Pr(y_{it}, \eta_t)^{-1}}{v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} \sum_{h_t | \Pr(h_t|y_{it},\eta_t)>0} \eta_t(h^t)$$

since the term $\frac{\Pr(y_{it},\eta_t)^{-1}}{v'(y_{it}+\tau_{it}(y_{it},\eta_t))}$ does not depend on h^t: $\Pr(y_{it}, \eta_t)$ is the unconditional probability that (y_{it}, η_t) occurs.

Summing (22) over all time t realizations h_t such that $\Pr(h_t|y_{it}, \eta_t) > 0$ gives

$$\begin{aligned} & \frac{\Pr(y_{it}, \eta_t)^{-1}}{v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} \sum_{h_t | \Pr(h_t|y_{it},\eta_t)>0} \eta_t(h^t) \\ = & \sum_{h_t | \Pr(h_t|y_{it},\eta_t)>0} \Pr(h_t|y_{it}, \eta_t) \frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} \\ = & \mathbb{E} \left(\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} \middle| y_{it}, \eta_t \right) \end{aligned}$$

So that

$$\frac{1}{v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} \sum_{h_t | \Pr(h_t|y_{it},\eta_t)>0} \eta_t(h^t) = \Pr(y_{it}, \eta_t) \mathbb{E} \left(\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} \middle| y_{it}, \eta_t \right)$$

Summing over all realizations of (y_{it}, η_t) gives

$$\sum_{y_{it},\eta_t} \frac{1}{v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} \sum_{h_t | \Pr(h_t|y_{it},\eta_t)>0} \eta_t(h^t) = \mathbb{E} \left(\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} \right)$$

or

$$\sum_{h_t} \left(\frac{\eta_t(h^t)}{v'(y_{it} + \tau_{it}(y_{it}, \eta_t))} \right) = \mathbb{E} \left(\frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} \right)$$

So, using the time $t+1$ envelope condition for $u_{i,t+1}(h^t)$, (15):

$$\frac{\eta_t(h^t)}{v'(y_{it} + \tau_{it}(h^t))} = \Pr(h_t) \frac{\partial u_N(\mathbf{u}_{t+1}(h^t), a_{t+1}(h^t), \mathbf{e})}{\partial u_{i,t+1}(h^t)} = \mathbb{E}_{h^{t+1}} \left(\lambda_{i,t+1} | h^t \right)$$

lagging by one period and using the FOC for $\tau_{it}(h^t) = \tau_{it}(y_{it}, \eta_t)$,

$$\mathbb{E}_{h^t}\left(\frac{\lambda_{it}}{\eta_t(h^t)}|h^{t-1}, \eta_t(h^t)\right) = \frac{1}{v'(c_{i,t-1}(h^{t-1}))} = \frac{1}{\eta_t(h^t)}\left(\lambda_{i,t-1} + \frac{\phi_{ir,t-1}(y_{i,t-1})}{\Pr(h^t|h^{t-1})}\right).$$

Starting from the multiplier on the initial promise-keeping constraint, λ_{i0},

$$\begin{aligned}\mathbb{E}_{h^t}\left(\frac{1}{v'(c_{it}(h^t))}|\eta_t(h^t)\right) &= \\ \frac{1}{v'(c_{i,t-1}(h^{t-1}))} &= \frac{\lambda_{i,t-1}(h^{t-1})}{\eta_t(h^t)} = \lambda_{i0} + \sum_{\tau=1}^{t-1}\frac{\phi_{ir,t-\tau}(y_{i,t-\tau})}{p(y_\tau)\eta_\tau}\end{aligned}$$

Lagged inverse marginal utility, conditional on the current shadow price of resources $\eta_t(h^t)$, captures all past information relevant to forecasting current marginal utility of consumption.■

The need to condition on $\eta_t(h^t)$ is due to the fact that both $\frac{1}{v'(c_{it})}$ and η_t are martingales ($\mathbb{E}\frac{1}{v'(c_{it})} = \frac{1}{v'(c_{i,t-1})}; \mathbb{E}\eta_t = \eta_{t-1}$). It is not generally the case that the unconditional expectation $\mathbb{E}\frac{\eta_t}{v'(c_{it})} = \frac{\eta_{t-1}}{v'(c_{i,t-1})}$, but, conditioning on $\eta_t, \mathbb{E}\left(\frac{\eta_t}{v'(c_{it})}|\eta_t\right) = \frac{\eta_t}{v'(c_{i,t-1})}$.

7.3 Proof of Proposition 3: Under moral hazard, lagged inverse marginal utility is a sufficient statistic for current consumption

The proof proceeds in two steps. First, to show that the difference between the multipliers on the household's time t promise- and threat-keeping constraints equals expected time t inverse marginal utility. Second, that the expected difference between the multipliers on the household's time t promise- and threat-keeping constraints equals time t inverse marginal utility; the difference is a random walk (conditional on the time t budget multiplier, η_t).

Again let λ_{it} be the multiplier on household i's promise-keeping constraint, and η_t be the multiplier on the village's time t budget constraint. Let ζ_{it} be the multiplier on household i's incentive-compatibility constraint. (Since there are only two possible effort levels and utility is separable in consumption and effort, the result of Grossman and Hart (1983) implies that the incentive-compatibilty constraint will be binding at the optimum.)

The planner's problem is now

$$\begin{aligned}&u_N(\mathbf{u}_t, \hat{\mathbf{u}}_t, a_t|\mathbf{e}') \equiv \max_{\{\boldsymbol{\tau}_{rt}\},\{\mathbf{u}_{r,t+1}\},\{\hat{\mathbf{u}}_{r,t+1}\}} \\ &\sum_{r=1}^{S} p_{r11}v(y_r + \tau_{Nrt}) - c(1) + \beta\mathbb{E}_{\{y\}}u_N(\mathbf{u}_{t+1}, \hat{\mathbf{u}}_{t+1}, a_{t+1}|\mathbf{e})\end{aligned}$$

subject to the promise-keeping constraints:

$$\sum_{r=1}^{S} p_{r11}[v(y_r + \tau_{irt}) - c(1) + \beta u_{ir,t+1}] \geq u_{it}, i < N \qquad (\lambda_{it})$$

the law of motion for assets:

$$R^{-1}a_{t+1} = a_t - \sum_{i=1}^{N} \tau_{irt} \qquad (\eta_t)$$

the incentive-compatibility constraints:

$$\begin{aligned} & \sum_{r=1}^{S} p_{r11}[v\left(y_r + \tau_{irt}\right) + \beta u_{ir,t+1}] - c(1) \qquad (\zeta_{it}) \\ = & \sum_{r=1}^{S} p_{r10}[v\left(y_r + \tau_{irt}\right) + \beta \hat{u}_{ir,t+1}] - c(0) \end{aligned}$$

threat-keeping 1: if the household disobeyed yesterday but obeys today, they don't get more than $\hat{u}_{it}$:

$$\sum_{r=1}^{S} p_{r10}[v(y_r + \tau_{irt}) - c(1) + \beta u_{ir,t+1}] \leq \hat{u}_{it}, i < N \qquad (\psi_{1it})$$

threat-keeping 2: if the household disobeyed yesterday and disobeys today, they don't get more than $\hat{u}_{it}$:

$$\sum_{r=1}^{S} p_{r00}[v(y_r + \tau_{irt}) - c(0) + \beta \hat{u}_{ir,t+1}] \leq \hat{u}_{it}, i < N \qquad (\psi_{2it})$$

The FOCs are:

τ_{irt} :

$$\frac{\eta_t p_{r11}^{-1}}{v'(y_r + \tau_{irt})} = \lambda_{it} + \frac{p_{r11} - p_{r01}}{p_{r11}}\zeta_{it} - \frac{p_{r10}}{p_{r11}}\psi_{1it} - \frac{p_{r00}}{p_{r11}}\psi_{2it}$$

$u_{ir,t+1}$:

$$-\mathbb{E}_{\{y_{-i}|y_i\}}\frac{\partial u_N(\cdot,\cdot,\cdot|\mathbf{e})}{\partial u_{ir,t+1}} = \lambda_{it} + \zeta_{it} - \frac{p_{r10}}{p_{r11}}\psi_{1it}$$

$\hat{u}_{ir,t+1}$:

$$-\mathbb{E}_{\{y_{-i}|y_i\}}\frac{\partial u_N(\cdot,\cdot,\cdot|\mathbf{e})}{\partial \hat{u}_{ir,t+1}} = -\frac{p_{r01}}{p_{r11}}\zeta_{it} - \frac{p_{r00}}{p_{r11}}\psi_{2it}$$

a_{t+1} :

$$\mathbb{E}_{\{y\}}\frac{\partial u_N(\cdot,\cdot,\cdot|\mathbf{e})}{\partial a_{t+1}} = \eta_t$$

and the envelope conditions:

$$
\begin{aligned}
-\frac{\partial u_{Nt}(\mathbf{u}_t, \hat{\mathbf{u}}_t, a_t|\mathbf{e}')}{\partial u_{irt}} &= \lambda_{it} \\
-\frac{\partial u_{Nt}(\mathbf{u}_t, \hat{\mathbf{u}}_t, a_t|\mathbf{e}')}{\partial \hat{u}_{irt}} &= \psi_{1it} + \psi_{2it} \\
\frac{\partial u_{Nt}(\mathbf{u}_t, \hat{\mathbf{u}}_t, a_t|\mathbf{e}')}{\partial a_t} &= \eta_t
\end{aligned}
$$

Multiplying the FOC for each τ_{irt} by p_{r11} and summing gives

$$
\eta_t \mathbb{E}\left(\frac{1}{v'(y_r + \tau_{irt})}|\eta_t\right) = \lambda_{it} - (\psi_{1it} + \psi_{2it})
$$

Expected inverse marginal utility at t equals the difference $\lambda_{it} - (\psi_{1it} + \psi_{2it})$ (Step 1) Adding the FOCs for $u_{ir,t+1}$ and $\hat{u}_{ir,t+1}$ gives:

$$
\mathbb{E}_{\{y_{-i}|y_i\}}\left(\frac{-\partial u_N(\mathbf{u}_{t+1}, \hat{\mathbf{u}}_{t+1}, a_{t+1}|\mathbf{e})}{\partial u_{ir,t+1}} - \frac{\partial u_N(\mathbf{u}_{t+1}, \hat{\mathbf{u}}_{t+1}, a_{t+1}|\mathbf{e})}{\partial \hat{u}_{ir,t+1}}\right)
$$

$$
= \underbrace{\lambda_{it} + \zeta_{it} - \frac{p_{r10}}{p_{r11}}\psi_{1it}}_{u_{ir,t+1}} + \underbrace{\left(-\frac{p_{r01}}{p_{r11}}\zeta_{it} - \frac{p_{r00}}{p_{r11}}\psi_{2it}\right)}_{\hat{u}_{ir,t+1}}
$$

$$
\begin{aligned}
&= \lambda_{it} + \frac{p_{r11} - p_{r01}}{p_{r11}}\zeta_{it} - \frac{p_{r10}}{p_{r11}}\psi_{1it} - \frac{p_{r00}}{p_{r11}}\psi_{2it} \\
&= \frac{\eta_t}{v'(y_r + \tau_{irt})}
\end{aligned}
$$

Lagging this by one period,

$$
\frac{\eta_{t-1}}{v'(y_{i,t-1} + \tau_{i,t-1})} = \mathbb{E}_{\{y\}}\frac{-\partial u_N(\mathbf{u}_t, \hat{\mathbf{u}}_t, a_t|\mathbf{e}')}{\partial u_{it}} - \frac{\partial u_N(\mathbf{u}_t, \hat{\mathbf{u}}_t, a_t|\mathbf{e}')}{\partial \hat{u}_{it}}
$$

So that, using the time t envelope conditions for u_{it} and $\hat{u}_{it}$:

$$
\frac{\eta_{t-1}}{v'(y_{i,t-1} + \tau_{i,t-1})} = \lambda_{it} - (\psi_{1it} + \psi_{2it})
$$

Using Step 1, this implies

$$
\frac{1}{v'(y_{i,t-1} + \tau_{i,t-1})} = \frac{\eta_t}{\eta_{t-1}}\mathbb{E}\left(\frac{1}{v'(y_{it} + \tau_{it})}|\eta_t\right)
$$

Inverse marginal utility times the budget multiplier is a random walk (given the time t budget multiplier).

LIMU is a sufficient statistic for past information in forecasting consumption.■

7.4 Proof of proposition 4: With hidden income, lagged inverse marginal utility over-predicts consumption for low-lagged income households

Let λ_{it} be the multiplier on household i's promise-keeping constraint, η_t the multiplier on the budget constraint, and ζ_{irt} the multiplier on the truth-telling constraint when $y_t = y_r$. The FOCs are:

τ_{irt} :

$$\eta_t = (p_{ree'}\lambda_t + \xi_{irt})\, v'(y_r + \tau_{rt}) - \xi_{i,r+1,t} v'(y_{r+1} + \tau_{irt})$$

$u_{ir,t+1}$:

$$p_{ree'}\mathbb{E}_{\{y_{-i}|y_i\}} \frac{-\partial u_N(\mathbf{u}_{t+1}, a_{t+1}, \mathbf{e})}{\partial u_{ir,t+1}} = p_{ree'}\lambda_t + \xi_{irt} - \xi_{i,r+1,t}$$

a_{t+1} :

$$-\mathbb{E}_{\{y\}} \frac{\partial u_N(\mathbf{u}_{t+1}, a_{t+1}, \mathbf{e})}{\partial a_{t+1}} = -\eta_t$$

envelope conditions:

$$\begin{aligned} \frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e}')}{\partial u_{it}} &= -\lambda_{it} \\ \frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e}')}{\partial a_t} &= \eta_t \end{aligned}$$

The lagged promise-keeping multiplier, λ_{t-1}, is a sufficient statistic for history, since the FOC for $u_{ir,t+1}$ and the envelope condition for u_{it} imply

$$\mathbb{E}\left(\lambda_{i,t+1}|\eta_{t+1}\right) = \lambda_{it} + \frac{\xi_{irt} - \xi_{i,r+1,t}}{p(y_t)}$$

lagging one period,

$$\mathbb{E}\left(\lambda_{it}|\eta_t\right) = \lambda_{i,t-1} + \frac{\xi_{ir,t-1} - \xi_{i,r+1,t-1}}{p(y_{t-1})}.$$

The FOC for transfers at $t-1$ implies that

$$\lambda_{i,t-1} = \frac{1}{v'(y_r + \tau_{ir,t-1})} \times \left(1 - \frac{\xi_{ir,t-1} v'(y_r + \tau_{irt}) - \xi_{i,r+1,t-1} v'(y_{r+1} + \tau_{irt})}{\eta_{t-1} p(y_{t-1})}\right) \tag{23}$$

Since $\mathbb{E}(\lambda_{it}|\eta_t) = \lambda_{i,t-1}$,

$$\mathbb{E}(\lambda_{it}|\eta_t) = \frac{1}{v'(y_r + \tau_{r,t-1})} \times \left(1 - \frac{\xi_{ir,t-1} v'(y_r + \tau_{rt}) - \xi_{i,r+1,t-1} v'(y_{r+1} + \tau_{rt})}{\eta_{t-1} p(y_{t-1})}\right)$$

Using the envelope condition for u_{it}, the time $t-1$ FOC for u_{it} can be written

$$\frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it}} - \frac{\partial u_N(\mathbf{u}_{t-1}, a_{t-1}, \mathbf{e})}{\partial u_{i,t-1}} = \frac{\xi_{i,r,t-1} - \xi_{i,r+1,t-1}}{p_{ree'}}$$

First, assume no aggregate uncertainty: $a_t = a_{t-1}$

Since $u_N(\mathbf{u}_t, a_t, \mathbf{e})$ is concave in each u_{it}, when a household's promise decreases ($u_{it} < u_{i,t-1}$),

then

$$\frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it}} > \frac{\partial u_N(\mathbf{u}_{t-1}, a_t, \mathbf{e})}{\partial u_{i,t-1}},$$

so $\xi_{ir,t-1} > \xi_{i,r+1,t-1}$: truth-telling constraints bind more at lower than higher output levels.

Then, since $v'(y_r + \tau_{rt}) > v'(y_{r+1} + \tau_{rt})$,

$$\xi_{ir,t-1} v'(y_{ir} + \tau_{ir,t-1}) > \xi_{i,r+1,t-1} v'(y_{ir+1} + \tau_{ir,t-1})$$

so

$$\mathbb{E}(\lambda_{it}|h^{t-1}) < \frac{1}{v'(y_{ir} + \tau_{ir,t-1})}$$

LIMU over-predicts λ_{it} when the household's promise decreased between $t-1$ and t. Promises are unobserved, but truth-telling implies that promises are an increasing function of income, so low-y_{t-1} households will get less consumption at t than predicted using lagged inverse marginal utility.

However, if $a_t > a_{t-1}$, there is an offsetting effect:

$$\begin{aligned}\frac{\partial^2 u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it} \partial a_t} &\neq 0 \Rightarrow \\ \frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it}} &\neq \frac{\partial u_N(\mathbf{u}_t, a_{t-1}, \mathbf{e})}{\partial u_{i,t-1}}\end{aligned}$$

However, we can sign this effect: by the envelope condition for u_{it}:

$$\frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it}} = -\lambda_{it}$$

So

$$\begin{aligned}\frac{\partial^2 u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it} \partial a_t} &= -\frac{\partial \lambda_{it}}{\partial a_t} \\ sgn\left(-\frac{\partial \lambda_{it}}{\partial a_t}\right) &= sgn\left(\frac{\partial \lambda_{it}}{\partial \eta_t}\right)\end{aligned}$$

Using the formula for λ_{it}:

$$\frac{\partial \lambda_{it}}{\partial \eta_t} = \frac{1}{v'(y_r + \tau_{irt})} \times$$
$$\frac{\partial}{\partial \eta_t}\left(1 - \frac{\xi_{irt} v'(y_r + \tau_{irt}) - \xi_{i,r+1,t} v'(y_{r+1} + \tau_{irt})}{\eta_t p(y_r)}\right)$$

$$sgn\left(\frac{\partial \lambda_{it}}{\partial \eta_t}\right) = sgn\left(\xi_{irt} v'(y_r + \tau_{irt}) - \xi_{i,r+1,t} v'(y_{r+1} + \tau_{irt})\right)$$

That is, when $u_{it} < u_{i,t-1}$,

$$\frac{\partial^2 u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it} \partial a_t} > 0$$

so the extent of "overprediction at the bottom" is reduced the greater is $\Delta a_t \equiv a_t - a_{t-1}$.■

7.5 Proof of proposition 5: Less variable income processes display a reduced wedge between LIMU and current inverse marginal utility:

Using (23):

$$\mathbb{E}(\lambda_{it}|\eta_t) = \frac{1}{v'(y_q + \tau_{iq,t-1})} \times$$
$$\left(1 - \frac{\xi_{iqt-1} v'(y_q + \tau_{iqt-1}) - \xi_{iq+1t-1} v'(y_{q+1} + \tau_{iqt+1})}{\eta_{t-1} p_{qee'}}\right)$$

Define

$$\theta(y_q) \equiv 1 - \frac{\xi_{iqt-1} v'(y_q + \tau_{iqt-1}) - \xi_{iq+1t-1} v'(y_{q+1} + \tau_{iqt+1})}{\eta_{t-1} p_{qee'}}$$

the term that measures the "wedge" between λ_{it} and $\frac{1}{v'(y_q + \tau_{iq,t-1})}$.

Take the expectation of $\theta(y_q)$, given that y_q was below the average level of income $\bar{y}$:

$$\mathbb{E}\left[\theta(y_q)|y_q < \bar{y}\right] =$$
$$\sum_{q: y_q < \bar{y}} \left[1 - \frac{\xi_{iqt-1} v'(y_q + \tau_{iqt-1}) - \xi_{iq+1t-1} v'(y_{q+1} + \tau_{iqt+1})}{\eta_{t-1} p_{qee'}}\right]$$

Fixing the probability of each income realization, $p_{qee'}$, a SOSD reduction in variability will reduce

$$\mathbb{E}\left[v'(y_q + \tau_{iqt-1}) - v'(y_{q+1} + \tau_{iqt+1})\right]$$

since income levels are closer together, and will reduce

$$\mathbb{E}\left[\xi_{ir,t-1} - \xi_{i,r+1,t-1}\right]$$

since

$$\frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it}} - \frac{\partial u_N(\mathbf{u}_{t-1}, a_{t-1}, \mathbf{e})}{\partial u_{i,t-1}} = \frac{\xi_{i,r,t-1} - \xi_{i,r+1,t-1}}{p_{ree'}}$$

a reduction in the amount of uncertainty about the household's income moves u_{it} and $u_{i,t-1}$closer together, on average (insurance improves), and, by the concavity of the planner's value function, this in turn reduces the gap $\frac{\partial u_N(\mathbf{u}_t, a_t, \mathbf{e})}{\partial u_{it}} - \frac{\partial u_N(\mathbf{u}_{t-1}, a_{t-1}, \mathbf{e})}{\partial u_{i,t-1}}$ (which remains negative since the household's promise is falling).

Therefore, $\mathbb{E}\left[\theta(y_q)|y_q < \bar{y}\right] \to 1$ as the variability of y decreases,so that the amount of additional information contained in y_{t-1} falls.■

Table 1: Consumption smoothing at the individual and village level

	Δlog(household PCE)	Δlog(household PCE)	Δ(avg log household PCE)
	(1)	(2)	(3)
Δlog(household income)	.0441***	.0423***	
	[.0083]	[.0085]	
Δ(avg log household income)			0.0348
			[.0364]
Village-year fixed effect?	No	Yes	
F(95,2577) statistic	-	2.125	-
P value	-	0.0000	-
Observations	2674	2674	96
R-squared	0.0105	0.0823	0.0096

Note: Robust standard errors (corrected for household-level clustering) in brackets. All variables are in 2002 Thai baht.
* p<.1, ** p<.05, *** p<.01

Table 2

A: Correlations in per capita expenditure over time, level and rank

Rank in village PCE distribution	2005	2004	2003	2002	2001	2000	1999
2005	1.000						
2004	0.643	1.000					
2003	0.645	0.658	1.000				
2002	0.565	0.681	0.680	1.000			
2001	0.453	0.549	0.591	0.589	1.000		
2000	0.354	0.409	0.436	0.437	0.539	1.000	
1999	0.375	0.442	0.466	0.459	0.525	0.824	1.000

B: Changes in PCE rank vs. changes in income

	OLS (LHS var: change in PCE rank)	Ordered probit (LHS var: direction of change)
Change in ln(income)	.527***	.1478***
	[.1414]	[.0225]
	9.757	*8.465*
R-Squared	0.0052	
N	2674	2674

Notes: In panel B, standard errors in brackets.

Table 3: Summary statistics

	531-HH panel mean	Non-continuously observed HHs difference	N
Monthly income	8981.224	-2624.627	670
Monthly expenditure	5213.472	-1108.721***	670
Monthly income, resids	32.443	-163.756	670
Monthly expenditure, resids	67.416	-570.84	670
Gifts given			
Gifts to orgs in village	33.714	-9.813	670
Gifts to orgs not in village	53.749	-29.063**	670
Gifts given for events in village	103.219	-35.550***	670
Gifts given for events not in village	220.117	-140.576***	670
Other gifts to HHs in village	147.317	-29.854	670
Other gifts/remittances to HHs not in village	637.198	-96.868	670
Gifts received			
Gifts from orgs in village	36.105	-20.002**	670
Gifts from orgs not in village	38.963	10.82	670
Gifts rec'd for events in village	316.862	-213.653***	670
Gifts rec'd for events not in village	80.068	9.976	670
Other gifts from HHs in village	118.129	-20.575	670
Other gifts/remittances from HHs not in village	1327.131	-253.376	670
Household size	4.525	-0.663***	669
Adult equivalents	3.786	-0.638***	669
Number adult men	1.382	-0.324***	669
Number adult women	1.552	-0.247***	669
Occupation of household head (baseline)			
Rice farmer	0.355	0.116*	667
Corn farmer	0.098	-0.062*	667
Orchard farmer	0.017	0.005	667
Other crop farmer	0.043	-0.036*	667
Shrimp/fish farmer	0.036	-0.021	667
Livestock farmer	0.089	-0.082***	667
Construction	0.015	0.036*	667
Non-ag wage labor	0.119	0.033	667
Ag wage labor	0.051	0.007	667
Other	0.074	0.013	667

Notes: All baht-denominated variables were converted to 2002 baht using the Thai Ministry of Trade's Rural

Table 4: Testing sufficiency of lagged inverse marginal utility

	Full sample		Drop top & bottom 10% of PCE	
	(1)	(2)	(3)	(4)
ln(LIMU)	.7386***	.7126***	.6215***	.5952***
	[.0208]	[.023]	[.0212]	[.0233]
Lagged income		.0424***		.0378***
		[.007]		[.0068]
Village Fes	Yes	Yes	Yes	Yes
R-squared	0.6645	0.6687	0.62	0.6299
Observations	3186	2845	2874	2573

Notes: Robust standard errors in brackets. Ln(LIMU) is proportional to ln(c_t-1).

Table 5: Testing Amnesia

	Full Sample		Drop top & bottom 5% of PCE		Low rainfall variance	High rainfall variance
	(1)	(2)	(3)	(4)	(5)	(6)
ln(LIMU)	0.727***	0.669***	0.584***	0.554***	0.949***	0.933***
	[0.011]	[0.012]	[0.012]	[0.012]	[0.027]	[0.024]
ln(LIMU)X25						
ln(LIMU)X50						
ln(LIMU)X75						
ln(current income)	0.052***	0.050***	0.038***	0.037***	0.030*	-0.004
	[0.001]	[0.001]	[0.001]	[0.001]	[0.012]	[0.012]
ln(income)X25	0.084***	0.080***	0.064***	0.063***		
	[0.002]	[0.001]	[0.001]	[0.001]		
ln(income)X50	0.138***	0.132***	0.103***	0.100***		
	[0.002]	[0.002]	[0.002]	[0.002]		
ln(income)X75		0.107***		0.067***		
		[0.007]		[0.005]		
Chi-squared (High=Low)						0.20
p						(0.658)
Fixed effects	Village	Village	Village	Village	Village	Village
Sample	Full	Full	Middle 90% by PCE	Middle 90% by PCE	HHs with above-median growth in PCE, low variance villages	HHs with above-median growth in PCE, high variance villages
R-squared	0.85	0.86	0.82	0.83	0.70	0.74
N	3186	2860	2874	2589	665	811

Table 6: Testing the hidden income model

LHS=Prediction residuals from a regression of ln(c_t) on ln(c_t-1) and a village-year effect.

	OLS		IV	
	(1)	(2)	(3)	(4)
Constant (a)	-.5406***	-.4839***	-.2301***	-.2123***
	[.0691]	[.0694]	[.0668]	[.0576]
Lagged log income (b)	.0509***	.0453***	.0224***	.0205***
	[.0061]	[.0063]	[.0059]	[.0052]
Control for aggregate shock interactions?	Yes	No	Yes	No
Chi-square stat (a<0, b>0)	81.47	54.84	19.11	19.40
p value	(0.000)	(0.000)	(0.000)	(0.000)
Observations	2781	2781	2322	2322

Notes: Bootstrapped standard errors in brackets. All regressions include a village-year fixed effect. Chi-square stat is the statistic for the test that the slope>0, intercept<0. p-value in parentheses.

Table 7: Testing the hidden income model

Panel A: Rainfall R-squared

LHS=Prediction residuals from a regression of ln(c_t) on ln(c_t-

Constant (a1)	-0.70
	[0.132]
Rainfall R-squared (a2)	0.681
	[0.509]
Lagged log income (b1)	0.061
	[0.012]
Lagged log incomeXrainfall R-squared (b2)	-0.046
	[0.047]
Aggregate shock interactions?	Yes
a1+a2	-0.015
Chi-square stat (a1+a2=0)	0.00
p value	(0.970)
b1+b2	0.015
Chi-square stat (b1+b2=0)	0.159
p value	(0.690)
Observations	2499

Notes: Bootstrapped standard errors in brackets. Rainfall R-

Table 7: Testing the hidden income model

Panel B: Split by variance of income

LHS=Prediction residuals from a regression of ln(c_t) on ln(c_t-1) and a village-year effect.

	High variance	Low variance
	(1)	(2)
Constant (a)	-0.49	-0.406
	[0.087]	[0.089]
Lagged log income (b)	0.047	0.037
	[0.008]	[0.008]
Control for aggregate shock interactions?	Yes	Yes
Chi-square stat (a<0, b>0)	56.96	22.03
p value	(0.000)	(0.000)
Observations	1387	1394

Notes: Bootstrapped standard errors in brackets. Chi-square stat is the statistic for the test that the slope>0, intercept<0. p-value in parentheses.

Table 9: Testing the hidden income model, nonparametric u()

LHS=Prediction residuals from a regression of ln(c_t) on f(c_t-1) and a village-year effect.

	OLS	IV
	(1)	(2)
Constant (a)	-0.37	-0.1411
	[0.0643]	[0.0668]
Lagged log income (b)	0.034	0.014
	[0.0059]	[0.0060]
Control for aggregate shock interactions?	Yes	Yes
Chi-square stat (a<0, b>0)	33.86	7.30
p value	(0.000)	(0.026)
Observations	2781	2322

Notes: Std errors bootstrapped (50 replications) to account for the generated regressor in brackets. LHS variable is prediction residuals from OLS or IV regression of ln(c_t) on f(c_t-1) and a village-year effect. Column (1) uses the nonparametric spline estimate of f(c_t-1) as an explanatory variable to form the predicted value of ln(c_t); column (2) instruments this nonparametric estimate with its lag, f(c_t-2). Chi-square stat is the statistic for the test that the slope>0, intercept<0. p-value in parentheses.

Table A1: Predicting income with rainfall

Occupation	R2	N
Rice farmer	0.386	752
Construction	0.292	32
Orchard farmer	0.222	36
Shrimp/fish farmer	0.195	76
Ag wage labor	0.143	108
Livestock	0.142	188
Other crop farmer	0.120	92
Non-ag wage labor	0.116	252
Other	0.100	156
Corn farmer	0.088	208

Notes: R2 is the R-squared of annual income on quarterly income deviations and squared deviations, plus province-fixed effects. N is the number of household-year observations.