

RISK SHARING AND INFORMATION IN VILLAGE ECONOMIES

ETHAN LIGON

ABSTRACT. Arrangements for achieving efficient risk-sharing vary depending on the information available to agents in the economy. The usual Euler equation restricts efficient allocations in an economy which obeys the permanent income hypothesis, while efficient allocations in an economy with private information and long-term contracts satisfy a symmetric restriction, but *not* the Euler equation. Full insurance arrangements are unique in that they satisfy both restrictions.

We look at an environment in which it seems likely that long-term contracts play a role in mitigating the effects of private information: three village economies in South India. The evidence that consumption allocations satisfy the private information restriction is quite strong for households in two of the three villages; the evidence for the third village suggests that while consumptions for some households satisfies the private information restrictions, other households' consumptions obey the permanent income hypothesis.

Department of Agricultural and Resource Economics
and Member, Giannini Foundation
University of California
Berkeley, CA 94720-3310

JEL Classification: D82,D91,O12

Keywords: Consumption smoothing, Private information, Village economies

Date: November 1997.

This paper is derived from my dissertation, written at the University of Chicago. Thanks are due to my committee, especially Robert Townsend, and to the many colleagues and seminar audiences who have given helpful comments. I am, of course, responsible for any remaining errors.

1. INTRODUCTION

This paper attempts to characterize and test the behavior of consumption in a class of private information models. Interest in this class of models is motivated by the failure of ‘permanent income’ models and complete markets models, two workhorses of economic theory, to adequately account for the role which individual income plays in determining consumption, particularly in the village economies of developing countries.

Although many other researchers have appealed to private information in order to explain the shortcomings of existing models, this research is among the first to actually test tightly specified versions of the complete markets and permanent income models against an explicit private information alternative.

Economists’ models can often be classified according to what markets are present in them. In this taxonomy, there is a locus of models near Friedman’s work on the permanent income hypothesis. These are models which have a full set of spot and credit markets, and which we will assign to the permanent income regime. This regime includes important contributions by Muth (1960), Bewley (1977) and Hall (1978). Consumption in these permanent income models depends on the nature of the *individual* income processes of agents.

A second cluster of models owes its position in the taxonomy to Arrow and Debreu. These models, characterized by complete contingent claims markets, generally suppose a larger set of markets than in the permanent income models. It is an implication of these models that individuals’ consumption depends only on the *aggregate* income process, not on individual income; the consumption insurance that results is reasonably called *full* insurance [cf. Diamond (1967), Wilson (1968)].

The third set of models belong to the private information regime. These are models which, in the presence of full information, *would* deliver the complete set of contingent claims markets promised by the Arrow-Debreu class of models. What distinguishes these models from the complete markets models, however, is that not all information is known to all agents, ruling out some contracting possibilities due to the problems of adverse selection and moral hazard. The optimal consumption allocations in such environments will fail to achieve Pareto optimality; the best attainable allocation will be second best, or constrained-efficient.

It is probably not the case that models belonging to a particular regime are right while the others are wrong. The question, rather, is this: if we take a particular actual economy, which regime provides the best description of that economy? Differences in technology and information structure alone are enough to suggest that different economies may be best described by different regimes.

This paper offers a consumption based means of determining which regime best describes a given economy. Subject to separability in preferences between actions and consumptions, we derive

a partial, but fairly general, characterization of constrained efficient consumption allocations in environments with private information due to Rogerson (1985). Furthermore, this characterization can be tested in the absence of further knowledge of both the information structure and production technology. Under a particular parameterization of preferences, the hypothesis that consumption allocations are indeed constrained-efficient yields a restriction which can be tested against similar restrictions that provide partial characterizations of the full-insurance and permanent income regimes.

Using panel data from three villages in rural South India, we test which regime best describes the data. The private information regime clearly outperforms the full insurance and permanent income alternatives in two of three villages; there is evidence that different households may belong to different regimes in the third. Such differences across households or villages are not at odds with the theory; differences in preferences or technology may mean that different households or villages are best modeled by different regimes.

Section 2 sets up some notation, and makes precise what we mean by a permanent income economy, a full information economy, and a private information economy. Along with this precision comes a set of mutually exclusive restrictions, which makes clear how one might distinguish full insurance and permanent income economies on the basis of observations of consumption. In Section 2.3 we derive a restriction which characterizes a fairly large class of constrained-efficient consumption allocations, and which, when contrasted with the restrictions derived in Sections 2.1 and 2.2 will allow us to judge which regime best fits the data. Section 3 discusses the data, methodology, estimation procedure, and results, in that order. Section 4 concludes.

2. THREE MODELS

In this section we set out to derive distinct sets of restrictions on intertemporal consumption allocations under uncertainty. We wish to characterize three different model regimes: the permanent income regime, the full insurance regime, and the private information regime. We obtain a general set of restrictions for each regime under the maintained hypothesis that agents' von Neumann-Morgenstern preferences are additively separable in consumption and action.¹

¹Generalizations to non-separable preferences are straightforward for the full information and permanent income regimes. For the private information regime, one must make some additional assumptions regarding the information structure of the economy. Roughly, in order for it to be convenient to generalize the private information restrictions, it must be the case that each agent's marginal utility of consumption be publicly observable. While we find ourselves in excellent company in assuming separable preferences, violations of this assumption may, of course, have dire consequences. While we will not describe these here, readers are referred to Gjesdal (1982) for an understanding of the role that nonseparabilities might play in models with private information, and Browning and Meghir (1981) for an empirical example in which nonseparabilities appear to be important.

The environment. The environment has a principal (sometimes labeled agent 0), and a finite set of agents indexed by $i = 1, 2, \dots, n$. Time is discrete and is indexed by $t = 0, 1, 2, \dots, T$. The random state in period t is denoted by ω_t and is drawn from some finite, discrete set of possible states Ω_t .

In each period t agent i takes some action $a_{it} \in A$ and consumes some quantity $c_{it} \in C$. Each of these affects i 's contemporaneous utility, and the actions taken by agents may affect the probability distribution of the state. We denote the history of actions for agent i by a_i^t , and denote the actions of all agents at time t by a_t .

The history of the economy is initially given by h_{-1} , and for $t = 0, 1, \dots, T$ is recursively defined by $h_t = (h_{t-1}, \omega_t, a_t)$. Let the set of possible histories from 0 to t be denoted by H_t , while the collection of possible histories of all lengths is H . The probability of a particular history h_t being realized is given by $\Pr(h_t)$.

Preferences. Principal and agents have von Neumann-Morgenstern preferences of the form

$$(1) \quad u_i = \sum_{t=0}^T \beta^t [U_i(c_{it}) - V_i(a_{it})]$$

where $U_i : C \rightarrow \mathbb{R}$ and $V_i : A \rightarrow \mathbb{R}$ are continuously differentiable, strictly increasing, and weakly concave. The discount factor β is assumed to be common to all agents (including the principal), and is taken to lie in the interval $(0, 1)$.

Technology. In each period t , each agent i produces some quantity x_{it} of the consumption good. Actions taken by the agents affect production because x_{it} depends on the realization of ω_t , the distribution of which depends on actions taken earlier in the period.

In addition, agents may have access to a linear technology mapping a single unit of the consumption good at $t-1$ into $1+r_t$ units of the consumption good at t . We might think of this technology as some sort of storage; however, because storage is usually positive, it's somewhat more natural to think of this technology as some sort of credit.

2.1. Permanent Income. Various formulations of the permanent income hypothesis have been tested in a wide variety of environments, both in the United States and in a number of developing countries [see Deaton (1992) for an excellent survey]. Alone of the model regimes described in this paper, the permanent income hypothesis is essentially derived from a partial equilibrium model; in order to express the hypothesis, we need only solve the problem of a single household which takes prices as given. In this tradition, the formulation of the permanent income hypothesis given here presumes universal access to some risk-free asset which bears a rate of return r_t .

Let the amount of the consumption good saved (or borrowed, if negative) by agent i at date t after history h_{t-1} be denoted by $b_{it}(h_{t-1})$. Take as given the the initial quantity of the asset held

by the agent, $b_{i0}(h_{-1})$. The agent solves

$$(2) \quad \max_{\{c_{it}(h_t), b_{i,t+1}(h_t), a_{it}(h_{t-1})\}_{h_t \in H}} \sum_{t=0}^T \beta^t \sum_{h_t \in H_t} \Pr(h_t) [U_i(c_{it}(h_t)) - V_i(a_{it}(h_{t-1}))]$$

such that the agent's resource constraint is respected,

$$(3) \quad c_{it}(h_t) = x_{it}(\omega_t) + (1 + r_t)b_{it}(h_{t-1}) - b_{i,t+1}(h_t)$$

for all $h_t \in H$, and subject to the restriction that the agent pays (and is repaid) all of his debts,

$$(4) \quad b_{i,T+1}(h_T) = 0 \quad \text{for all } h_T \in H_T.$$

The solution to this problem is partially characterized by the familiar Euler equation,

$$(5) \quad \mathbb{E} \left(\frac{U_i'(c_{i,t+1}(h_{t+1}))}{U_i'(c_{it}(h_t))} \middle| h_t \right) = \frac{1}{\beta(1 + r_t)},$$

where $\mathbb{E}(\cdot | h_t)$ denotes expectations conditioned on the history h_t . From (5), we can see that agents will share a common intertemporal marginal rate of substitution equal to the common marginal rate of transformation. Notice, however, that so long as there is cross-sectional variation in the realization of individual output x_{it} , there will be variation in the marginal rate of substitution across states of nature. This is just another way of saying that there are missing insurance markets in this formulation of a permanent income economy.

It is worth discussing some of the qualitative features of consumption implied by the Euler equation (5). In particular, suppose that the right hand side of (5) is equal to one (assume, that is, that the interest rate is constant and equal to the common rate of time preference). Then the marginal utility of consumption is a martingale process. We shall refer to this as the 'strict' version of the permanent income hypothesis. This is closely related to Hall's (1978) observation; if preferences are quadratic, then consumption itself is a martingale, an observation which seems at least roughly in accord with the available evidence (Deaton 1992).

The notion that the marginal utility of consumption follows a martingale is the implication of the permanent income hypothesis which has received the most attention in the empirical literature. However, if we focus solely on the time series behavior of consumption, we actually miss a much stronger implication of (5); namely, that agents' intertemporal marginal rates of substitution should be perfectly correlated across agents at a point in time. If all agents possess, say, identical iso-elastic utility functions, then (5) implies that expected consumption growth is the same for all agents at every period. The fact that it is only *expected* changes in agents' marginal utility of consumption that move in lockstep implies that there will be growing cross-sectional consumption inequality (Deaton and Paxson 1994; Lucas 1991; Phelan 1990).

2.2. Full Information. The task of finding the set of equilibrium allocations in our model with full information is equivalent to finding the set of Pareto optimal allocations. Suppose that the initial allocation is such that agents one through n have ex ante utility given by \underline{u}_i , $i = 1, 2, \dots, n$. Now we can just use the definition of a Pareto optimum to find Pareto optimal allocations by having the principal (agent 0) maximize her utility subject to providing the remaining n agents at least their reservation utilities $\{\underline{u}_i\}$, given some initial aggregate savings $\{b_{00}(h_{-1})\}$. Here we pose the problem as one in which the principal chooses a level of aggregate savings $b_{0,t+1}(h_t)$ and a set of contingent transfers, $\tau_{it}(h_t)$, to make to each of the other agents after every history. Formally, she also chooses the consumptions and actions of each of the agents; however, her choices for these quantities will agree precisely with what the agents themselves would have chosen, given transfers $\{\tau_{it}(h_t)\}$. We can write this problem as

$$(6) \quad \max_{\{\{\tau_{it}(h_t), c_{it}(h_t), a_{it}(h_{t-1})\}_{i=0}^n, b_{0,t+1}(h_t)\}_{h_t \in H}} \sum_{t=0}^T \beta^t \sum_{h_t \in H_t} \Pr(h_t) [U_0(c_{0t}(h_t)) - V_0(a_{0t}(h_{t-1}))]$$

subject to a set of resource constraints,

$$(7) \quad c_{0t}(h_t) = (1 + r_t)b_{0t}(h_{t-1}) + x_{0t}(\omega_t) - b_{0,t+1}(h_t) - \sum_{i=1}^n \tau_{it}(h_t)$$

for all $h_t \in H$, and

$$(8) \quad c_{it}(h_t) = x_{it}(\omega_t) + \tau_{it}(h_t)$$

for all agents $i = 1, 2, \dots, n$ and for all $h_t \in H$. There is also a requirement that all debt be repaid:

$$(9) \quad b_{0,T+1}(h_T) = 0 \quad \text{for all } h_T \in H_T,$$

and that reservation utilities are respected:

$$(10) \quad \sum_{t=0}^T \beta^t \sum_{h_t \in H_t} \Pr(h_t) [U_i(c_{it}(h_t)) - V_i(a_{it}(h_{t-1}))] \geq \underline{u}_i \quad i = 1, \dots, n.$$

The interest rate r_t that appears in the equations may be exogenously given, if the economy is an open one; or, if the economy is closed, may be endogenously determined. In the latter case, there would be an additional constraint due to the fact that outside resources cannot be brought into the economy. This would serve to close the model, and allow us to solve for r_t . However, we will not pursue this here.

Rather than providing a complete characterization of the Pareto optimal allocation, we'll concentrate on the implications of Pareto optimality for the intertemporal and inter-agent (or cross

sectional) allocation of consumption. Combining first-order conditions gives us

$$(11) \quad \frac{U_0'(c_{0,t+j}(h_{t+j}))}{U_0'(c_{0t}(h_t))} \frac{U_i'(c_{it}(h_t))}{U_i'(c_{i,t+j}(h_{t+j}))} = 1$$

for all $j = 0, 1, \dots$, which can be interpreted as a statement that marginal utilities of consumption remain fixed in proportion across agents. The fact that (11) involves no expectations operator is the hallmark of full insurance. There is no uncertainty in the ratio of future marginal rates of substitution because any fluctuations in income, whether individual or aggregate, are shared according to (11).

Finally, combining (11) with the first-order conditions with respect to $b_{0,t+1}(h_t)$ yields the Euler condition (5). Equation (5) holds, of course, for any model in which there is unrestricted access to credit markets, as in the previous section. Another way to see this is by noting that (5) would hold even if the $\{\tau_{it}(h_t)\}$ in the problem given above, rather than being a choice variable, were exogenously given. Note, however, that (11) would *not* generally hold in this latter case, which gives us a way to classify economies as belonging to either the permanent income or full insurance regimes.

2.3. Private Information. To solve for efficient allocations in the private information regime, we imagine that the principal maximizes own utility (6) subject to satisfying resource constraints (7) and (8), debt repayment (9), and subject to satisfying utility promises (10). This problem is identical to the full information problem to this point. The difference is that the actions $\{a_{it}(h_{t-1})\}$ chosen by the principal must also be incentive compatible; that is, the agent should find it in his best interest to follow the recommendations of the principal, given the compensation schedule $\{c_{it}(h_t)\}$ chosen by the principal:

$$(12) \quad \{a_{it}(h_{t-1})\} \in \operatorname{argmax}_{\{a_{it}(h_{t-1})\}} \sum_{t=0}^T \beta^t \sum_{h_t \in H_t} \Pr(h_t) [U_i(c_{it}(h_t)) - V_i(a_{it}(h_{t-1}))].$$

This restriction must hold for each agent. In our discussion of the problem of private information thus far we've emphasized the possibility that the agent takes some unobserved action which affects the distribution of output. Equation (12) certainly captures this possibility. However, it's important to recognize that (12) also encompasses the case in which the agent simply has better information regarding the distribution of output than does the principal. The action recommended by the principal to the agent in this case will be to make a truthful report; (12) is then a requirement that truth-telling must be in the agent's best interest.

Of course, the requirement of incentive compatibility greatly complicates the problem facing the principal, and generally leads to contracts which are arbitrarily complicated in the sense that they may depend on all the information in the environment (Holmström 1979). For the class of problems

treated here, however, there's a particular dimension along which these contracts are uncomplicated, because the agent is indifferent to changes in the contract along this dimension. In particular, imagine that the principal modifies the agent's objective function in (12) by taking away some utility y in one period, but then returning y/β utils (with certainty) in the following period. It's easy to see that this reallocation doesn't affect the incentive compatibility constraints (12) at all. The same cannot be said for the principal's problem. To see this, rather than assigning consumptions $c_{it}(h_t)$, which determine contemporaneous utility from consumption $U_i(c_{it}(h_t))$, let the principal solve the equivalent problem of assigning utility from consumption $z_{it}(h_t) = U_i(c_{it}(h_t))$. The key to rewriting the problem is to notice that the cost (in terms of the consumption good) of providing z utils to agent i is just the inverse utility function of agent i evaluated at z , or $U_i^{-1}(z)$. Thus the problem becomes

(13)

$$\max_{\{y_{i,t+1}(h_t), z_{it}(h_t), a_{it}(h_{t-1})\}_{i=0}^n, c_{0t}(h_t), b_{0,t+1}(h_t)\}_{h_t \in H}} \sum_{t=0}^T \beta^t \sum_{h_t \in H_t} \Pr(h_t) [U_0(c_{0t}(h_t)) - V_0(a_{0t}(h_{t-1}))]$$

subject to the resource constraints

$$(14) \quad c_{0t}(h_t) = (1 + r_t)b_{0t}(h_{t-1}) - b_{0,t+1}(h_t) + x_{0t}(\omega_t) + \sum_{i=1}^n [x_{it}(\omega_t) - U_i^{-1}(z_{it}(h_t) + y_{it}(h_{t-1})/\beta - y_{i,t+1}(h_t))]$$

for all $h_t \in H$ and subject also to the requirement that debt be repaid (9) and that reservation utilities are respected (10). Finally, of course, recommended actions must satisfy (12) in order to be incentive compatible. Initial assets $b_{00}(h_{-1})$ and initial utility reschedulings $\{y_{i0}(h_{-1})\}$ are taken as given.

For present purposes, we are concerned only with the principal's optimal intertemporal scheduling of utility assignments, reflected by her choice of $y_{i,t+1}(h_t)$. This scheduling is very simple, since the principal's choice of $y_{i,t+1}(h_t)$ doesn't affect incentive compatibility conditions for this problem. First-order conditions with respect to the $y_{i,t+1}(h_t)$ and $c_{0t}(h_t)$ yield

$$U_0'(c_{0t}(h_t))U_i^{-1'}(z_{it}(h_t)) = \mathbb{E} [U_0'(c_{0,t+1}(h_{t+1}))U_i^{-1'}(z_{i,t+1}(h_{t+1})) | h_t].$$

Applying the inverse function theorem, we have

$$(15) \quad \mathbb{E} \left[\frac{U_0'(c_{0,t+1}(h_{t+1}))}{U_0'(c_{0t}(h_t))} \frac{U_i'(c_{it}(h_t))}{U_i'(c_{i,t+1}(h_{t+1}))} \middle| h_t \right] = 1.$$

This is quite similar to the characterization of cross sectional marginal utilities of consumption (11) derived for the full information case, except that this version of that restriction holds only in

<i>Equation</i>	$E\left(\frac{U_i'(c_{i,t+1}(h_{t+1}))}{U_i'(c_{it}(h_t))} \mid h_t\right) = \frac{1}{\beta(1+r_t)}$	$\frac{U_0'(c_{0t}(h_t))}{U_i'(c_{it}(h_t))} = E\left(\frac{U_0'(c_{0,t+1}(h_{t+1}))}{U_i'(c_{i,t+1}(h_{t+1}))} \mid h_t\right)$
Full Information	×	×
Private Information	—	×
Permanent Income	×	—

TABLE I. Distinctions between models.

expectation. Note that this does *not* imply, as one might suppose, that ratios of marginal utilities are equal across agents in expectation; quite the opposite. Except for degenerate cases (of which full insurance is the leading example), intertemporal marginal rates of substitution will vary across agents. Note that this prediction is quite counter to what we see in either the full insurance or permanent income regimes (characterized by the Euler equation (5)), in both of which intertemporal marginal rates of substitution are constant across agents.

Table I summarizes our results— we’ve derived two conditions, unique combinations of which characterize each of three different regimes. The permanent income regime is characterized by (5) but *not* (15); the private information regime is characterized by (15) but *not* (5); and the full insurance regime is characterized by both (5) *and* (15).

To expand on the economic intuition behind the distinctions of Table I, consider first the permanent income regime, in which agents use noncontingent credit to smooth their consumption over time. Because agents all face the same prices, agents’ expected intertemporal marginal rates of substitution are all equal in cross-section. However, because consumption is not insured, these marginal rates of substitution will vary in cross-section *ex post*. The full insurance regime adds markets which insure away this uncertainty, so that intertemporal marginal rates of substitution are equal both *ex ante* and *ex post*. In the private information regime, the principal controls *ex post* variation in agents’ marginal rates of substitution, and uses this as a tool in order to help induce the agent to undertake possibly costly actions.

Finally, it’s worth noting that, because consumption in the regimes we’ve described is only partially characterized by various combinations of equations (5) and (15), there may be other models which would satisfy these restrictions aside from the ones we consider here. Further, these unknown alternative models may not have anything to do with information, and may feature allocations which are far from constrained-efficient.

3. INFERENCE AND ESTIMATION

3.1. Methods. The general idea is to test which of equations (5) or (15) provides a better fit to the observed (*ex post*) consumption data. If both seem to provide a reasonable fit, then we cannot reject the possibility that allocations are well characterized by a full insurance regime.

We test the simplest form of these equations, termed the ‘strict’ form above. Briefly, we can derive the strict form of each restriction by assuming that the principal is risk neutral. With this assumption, the strict form of the private information hypothesis is immediate; because $U_0'(c)$ is just some constant, equation (15) simplifies to become²

$$(16) \quad \mathbb{E}_t \left[\frac{U_i'(c_{it})}{U_i'(c_{i,t+1})} \right] = 1.$$

In the case of the permanent income economy, the strict form follows because the Euler equation (5) must be satisfied for every agent. If agent zero (or indeed, any agent) is risk neutral, then this agent’s marginal utility of consumption must be a constant, so that

$$\mathbb{E}_t \left[\frac{U_0'(c_{0,t+1})}{U_0'(c_{0t})} \right] = 1 = \frac{1}{\beta(1+r_t)}.$$

This equation is only satisfied when $r_t = 1/\beta - 1$, and so permits us to solve for the interest rate.³ Accordingly, the Euler equation for each agent becomes

$$(17) \quad \mathbb{E}_t \left[\frac{U_i'(c_{i,t+1})}{U_i'(c_{it})} \right] = 1.$$

Comparing these two results, the reader will note that the two differ only in that the operand of one is the reciprocal of the other; that is, substituting $\frac{U_i'(c_{it})}{U_i'(c_{i,t+1})}$ from (16) for its reciprocal in (17) yields the characterization of the private information regime. It is critical to note that despite their close relationship, Jensen’s inequality implies that (16) and (17) are very different restrictions, and imply very different behavior.

We will exploit this difference and this symmetry by selecting the popular constant relative risk aversion (CRRA) parameterization of utility from consumption $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$. Parameterized versions of equations (17) and (16) are, respectively,

$$(18) \quad \mathbb{E}_t \left[\left(\frac{c_{i,t+1}}{c_{it}} \right)^{-\gamma} - 1 \right] = 0$$

and

$$(19) \quad \mathbb{E}_t \left[\left(\frac{c_{i,t+1}}{c_{it}} \right)^{\gamma} - 1 \right] = 0$$

²Because we now turn our attention to consumption data given some fixed, realized history, we will henceforth suppress references to histories in our notation unless otherwise noted. Thus, $c_{it}(h_t)$ becomes c_{it} , and $\mathbb{E}(\cdot|h_t)$ becomes $\mathbb{E}_t(\cdot)$.

³Of course, if the agent’s consumption must lie in some bounded set C , then of course the Euler equation (5) will only hold when consumption lies on the interior of this set. This is a ‘knife’s-edge’ assumption in an economy in which all agents are risk-neutral, but is reasonably robust in a closed economy when at least some agents have concave preferences.

where, in each of these parameterizations, γ is interpretable as the coefficient of relative risk aversion; hence we would expect $\gamma > 0$, since agents will be risk averse if and only if this is so.

The symmetry of the restrictions (18) and (19) suggests an easy way to choose between the hypotheses they characterize. We simply nest both of (18) and (19), estimating b_0 ⁴ in

$$(20) \quad \mathbf{E}_t \left[\left(\frac{c_{i,t+1}}{c_{it}} \right)^{b_0} - 1 \right] = 0.$$

Thus $b_0 \in \{-\gamma, \gamma\}$. If our estimate of b_0 is positive, then we will prefer the private information hypothesis; conversely, if we infer that b_0 is negative, then we will prefer the permanent income hypothesis as an explanation of the determination of consumption.

We have said little in this section regarding the full insurance hypothesis. It was earlier demonstrated that *both* (18) and (19) will hold under full insurance; clearly this can be strictly true only if there is no uncertainty (and no ex post variation) in the ratio of agents' intertemporal marginal rate of substitution to the principal's intertemporal marginal rate of substitution. In this case, and in the absence of measurement error, agents' will be able to forecast their share of future consumption without error, and we will be unable to estimate γ at all, due to the resulting singularity of the covariance matrix of forecast errors. However, if we are unable to distinguish between forecast errors evaluated at γ and $-\gamma$, for some reasonable guess of γ , then we may be unable to reject the hypothesis that the economy closely approximates the full-insurance regime.

3.2. Estimation. We remarked above that we could use the conditional moment restriction (20) to estimate the parameter b_0 , and to use this estimate in order to select either permanent income or private information as the most suitable explanation for the consumption allocations observed in our sample.

Let

$$\xi_{it} = \frac{c_{i,t+1}}{c_{it}},$$

and let

$$h(\xi_{it}, b) = \xi_{it}^b - 1$$

denote our estimate of agent i 's time t forecast error. Our conditional moment restriction for household i then becomes

$$(21) \quad \mathbf{E}_t h(\xi_{it}, b_0) = 0.$$

⁴The b_0 we use here to denote an unknown parameter is unrelated to the b_{it} we use to denote savings; this is the only sentence in this paper which mentions both quantities. Context ought to make clear which one we have in mind at any particular point.

	Village			
	Aurepalle	Shirapur	Kanzara	All
Consumption (’75 Rs)	302.44 (159.78)	354.08 (173.35)	354.06 (165.62)	337.03 (167.94)
Income (’75 Rs)	713.43 (697.23)	560.33 (417.31)	820.76 (679.56)	704.61 (626.09)
Landholdings (Acres)	0.7362 (0.8102)	0.7707 (1.0167)	0.7113 (1.0321)	0.7379 (0.9598)
Household Size	5.8409 (2.5101)	6.3145 (2.8394)	6.6042 (3.5399)	6.2625 (3.0317)
Rain Skewness	0.6500 (0.5593)	0.7779 (0.6137)	0.5050 (0.4875)	0.6374 (0.5293)

TABLE II. Means and standard deviations of instruments

Now let ζ_{it} denote a q dimensional vector of variables known to agent i at time t . Rational expectations implies that any such vector ζ_{it} ought to be orthogonal to agent i 's time t forecast error $h(\xi_{it}, b_0)$, so that the unconditional expected value of the product of this vectors with the true forecast error ought to be equal to zero, or

$$(22) \quad \text{E}[h(\xi_{it}, b_0) \cdot \zeta_{it}] = 0.$$

The $t-1$ one-period forecast error is in the time t information set, and hence ought to be independent of time t forecast errors.⁵

The data we use comes from three villages in the semi-arid tropics of India. These three villages are some of the six villages initially selected by the International Crops Research Institute of the Semi-arid Tropics (ICRISAT) for inclusion in ICRISAT's Village Level Studies (VLS), described in careful detail by Walker and Ryan (1990). Although numerous surveys have been conducted as a part of the VLS, we will chiefly rely on a ten year panel consisting of (after deletion of some households for which data is incomplete) about 35 households in each of the three villages. We will use only a subset of these data; data collected on consumption expenditures is suspect in the first and later years.

The instruments we employ include per capita household income and landholdings (a useful proxy for physical wealth). Household size is another instrument, chosen as a test of our presumption that intra-household allocations are efficient. All of the instruments mentioned thus far may be

⁵While we have no assurance that forecast errors are independent across agents, consistent estimates of b_0 may be had under considerably weaker conditions than independence. See Pakes (1994) for a general discussion of this issue, and an earlier working paper (Ligon 1996) for some Monte Carlo evidence on the consistency of the estimator used here.

Instrument Vector	Village			
	Aurepalle	Shirapur	Kanzara	All
	b ($\sigma(b)$) J	b ($\sigma(b)$) J	b ($\sigma(b)$) J	b ($\sigma(b)$) J
—	-0.4896 (0.3721) 0	1.3450* (0.3753) 0	0.8015* (0.3349) 0	0.3016* (0.0773) 0.6354
Income	-0.3660 (0.3357) 0.0142	1.3153* (0.3942) 0.0044	0.7710* (0.3729) 0.0000	1.1056* (0.2135) 0.8318
Income, Land	-0.6396* (0.3106) 0.0707	1.2932* (0.4289) 0.1544	0.6018 (0.3792) 0.0603	1.5063* (0.2143) 0.5041
Family Size, Income, Land	-0.7820* (0.3389) 0.0309	1.1314* (0.4156) 0.1532	0.6728 (0.3616) 0.0457	1.2458* (0.1868) 1.0769
Family Size, Income, Land, Rain, Avg Consumption	-0.5683* (0.1148) 0.2579	1.5764* (0.2355) 1.9342	0.6801* (0.1193) 1.3801	0.8137* (0.0938) 2.2527

TABLE III. Estimates of b_0 . Starred estimates are significant at the 95% level. Standard errors are reported in parentheses. Statistics reported under the columns labelled J are scaled to have an asymptotic χ^2 distribution.

affected primarily by shocks peculiar to the household; we shall also consider variables reflecting only aggregate shocks. A set of dummy variables designed to capture any village-level fixed effects was employed to this end, as was a measure of the timing of rainfall, so critical to dry-land agriculture. This last measure, suggested by Sharma and Singh, actually is a measure of the skewness of the rainfall distribution. Finally, the village average consumption at $t + 1$ was also used as a variable reflecting aggregate shocks. Time t realizations of all other instruments (save for the village dummies) are assumed to be known to the agents only following time t .

Aside from the dummy variables and the $t + 1$ realization of village consumption, each of the instruments we employ is predetermined. In a time series setting this predetermination would suffice to deliver a set of valid instruments; however, predetermination is not sufficient when we rely on $N \rightarrow \infty$ (Hayashi 1992). Accordingly, we will need to test the assertion that the instruments in question are valid; Hansen's test of overidentifying restrictions provides an appropriate statistic.

Results are summarized in Table III. The first three columns of results are for individual villages; the fourth column pools all three villages. The instrument set used is listed on the left of the table; instruments are in per capita terms where appropriate. In addition to the variables listed, each instrument set includes a constant, when the population used is a single village; a set of village dummies is substituted for this constant when the three villages are pooled.

The point estimates reported in Table III vary considerably, ranging in absolute value from 0.3016 to 1.5764. Nonetheless, estimates for each particular village tend to be fairly consistent

across instrument sets, ranging (again in absolute value) from 0.3660 to 0.7820 for Aurepalle; from 1.1314 to 1.5764 in Shirapur; and from 0.6018 to 0.8015 in Kanzara. This consistency does not seem to extend to the pooled estimates, which vary considerably.

While the within-village consistency of the absolute value of our estimates is reassuring, the most important thing to notice about Table III is that the *signs* of each of the estimates are consistent within a village. All of the estimates in Aurepalle are negative, while all of the estimates in Shirapur and Kanzara are positive. This indicates that in Aurepalle, the permanent income model provides a better fit to the data than does the private information model, while in the other two villages, the private information model fits the data better than does the permanent income model.

Although the signs of our estimates are consistent within each village, not all of the estimates are significantly different from zero at a 95 per cent level of confidence. Two of the five different estimates in each of Aurepalle and Kanzara are not significantly different from zero by this standard. Accordingly, in these two villages, we may suspect that consumptions are not too far from the full insurance Pareto optimum. Unfortunately, the full information hypothesis is at something of a disadvantage in this setting. If we were to adopt either of the permanent income or private information hypotheses as our null, and full information as the alternative hypothesis, we could never reject the null hypothesis using the framework presented here, at least without independent information on the true value of the preference parameter γ . The reason for this is that if the full information hypothesis is correct, then consumptions will move together, and our test will be completely lacking in power. Put more sharply, if there is full insurance, then there is no information in the cross-sectional behavior of consumption changes.

Our estimates of b_0 are only meaningful so long as we haven't committed some serious specification error or used invalid instruments. Fortunately, the techniques we've employed here provide an omnibus specification test; the value of the minimized criterion function defined above is asymptotically distributed χ_{q-1}^2 , where q is the number of moment restrictions we test.⁶ The smallest critical value (for $q = 2$, at a 95 per cent level of confidence) is about 3.841. Using this specification test, none of the J statistics reported in Table III is anywhere near even the level at which we'd reject the restrictions we test.

The accompanying figure plots the criterion function for each of the three hypotheses described above.⁷ This allows us to emphasize the point that the private information hypothesis is much preferred to the permanent income hypothesis over the entire positive orthant. The best estimate

⁶For the private information and permanent income hypotheses, q is simply equal to the number of instruments we employ. For the full information hypothesis, q is equal to twice the number of instruments.

⁷The instrument set used for these estimates includes a set of village dummies, household size, per capita household income and the valuation of landholdings, the skewness of rainfall, all at time t , and average village consumption at time $t + 1$. All villages are pooled.

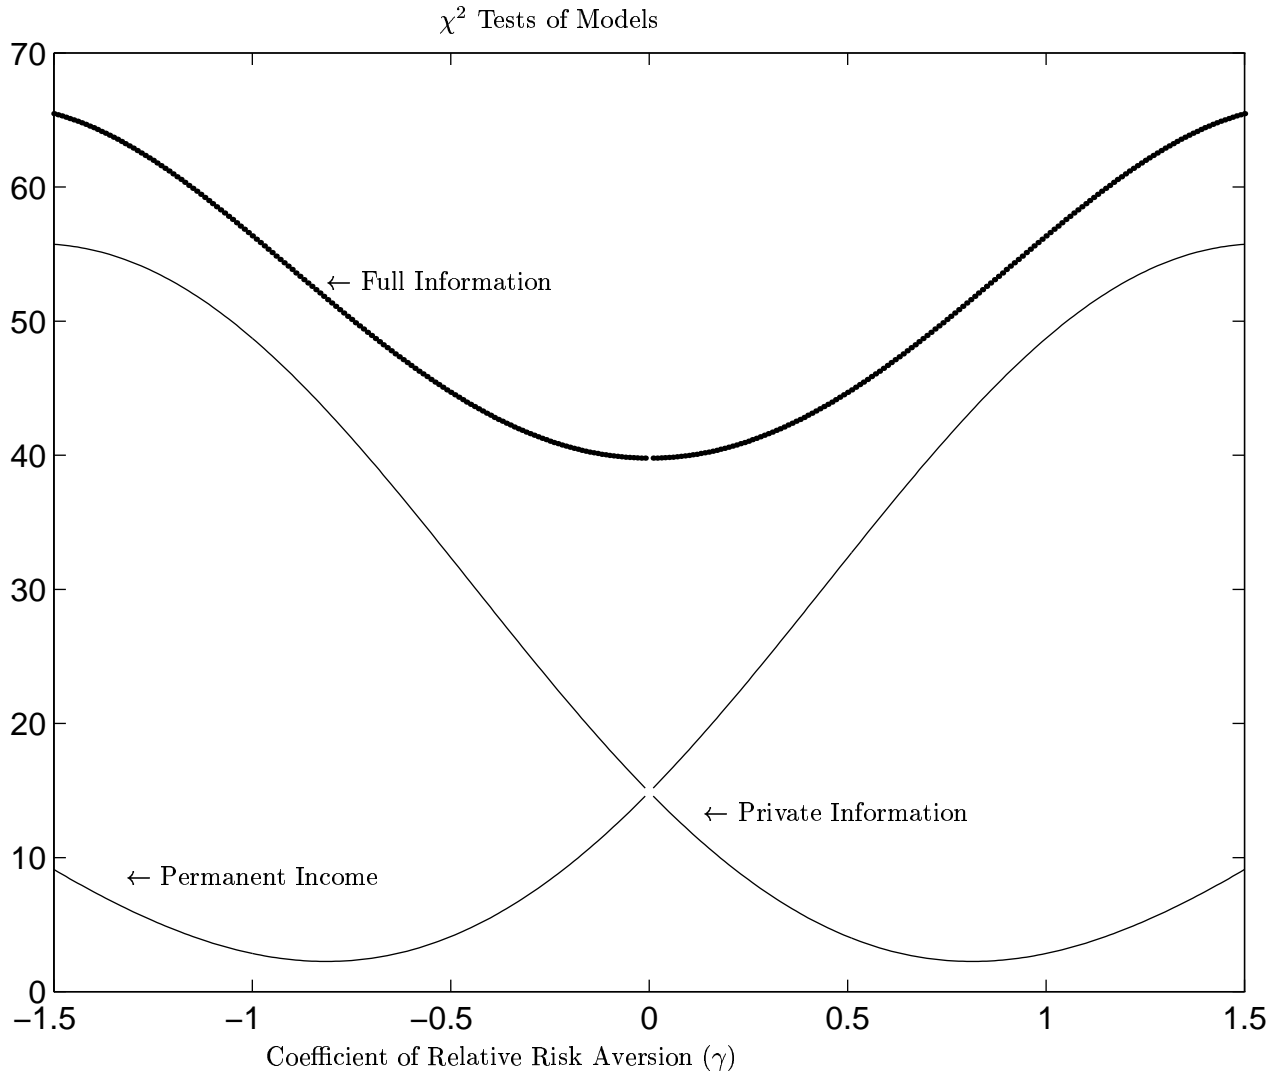


FIGURE 1. Comparison of permanent income, full information, and private information regimes using pooled data.

of γ under each regime occurs at the minimum of each curve; the curvature of each is related to the precision of the estimate, and the vertical distance between the curves for the permanent income hypothesis and the private information hypothesis provides another χ^2 statistic we can use to differentiate the two models given a particular value of γ .⁸ The permanent income regime is favored over the private information hypothesis only if agents are risk seeking.

Figure 1 makes it clear that our selection of a ‘best’ regime can only be made conditional on some prior regarding plausible values of risk aversion. After all, if agents are truly risk neutral, then they will have no need for the coinsurance implicit in our characterization of the private information

⁸We can’t quite use this vertical distance to test against the full information hypothesis, since the χ^2 statistic for this hypothesis is distributed with q more degrees of freedom than are the statistics for the other two hypotheses.

Village	γ	Permanent Income	Full Insurance	Private Information
Aurepalle	0.01	2.1541	5.8219	2.2897
	0.50	0.2862	7.3785	6.7598
	1.00	1.2837	11.6972	12.6704*
	1.50	4.1530	17.6440	18.4594*
Shirapur	0.01	19.5454*	5.4234	19.1801*
	0.50	27.2497*	7.4013	10.4732
	1.00	30.3724*	11.7898	4.2538
	1.50	28.0485*	14.8321	1.9699
Kanzara	0.01	4.9185	56.5458*	4.7103
	0.50	11.6241*	57.9435*	1.6009
	1.00	19.1927*	60.1808*	1.9795
	1.50	22.7672*	57.3151*	4.6137

TABLE IV. Comparison of χ^2 statistics for different models. Starred statistics are significant at 95% level.

regime. Similarly, we will never be able to use the test formulated here to decide between the joint hypothesis that agents have access to credit markets, but are risk seeking and the hypothesis that allocations are private information constrained efficient, and agents are risk averse.

However, if we are willing to take a stand on the magnitude of agents' relative risk aversion, we certainly can distinguish regimes. Table IV makes this clear, by reporting the χ^2 statistics associated with different regimes given various values of the relative risk aversion coefficient γ . Starred estimates imply that we can reject the associated regime with 95 per cent confidence. At $\gamma = 1$ for example (implying logarithmic utility), we can reject the null hypothesis that Aurepalle belongs to the private information regime in favor of some unspecified alternative, but we cannot reject the possibility that Aurepalle's economy is described by the full information or permanent income regimes.

Looking at the rejections in Table IV seems to turn up some very interesting patterns. Aurepalle is best characterized by the permanent income model, though we can't reject full information. We can reject the private information hypothesis in Aurepalle, at least for higher levels of risk aversion. This seems consistent with the institutional observation that there is a great deal of borrowing and lending within Aurepalle; if these credit transactions have few contingencies, then perhaps some version of the permanent income hypothesis ought to be expected to hold.

In sharp contrast, Shirapur is best characterized by the private information model. The permanent income hypothesis is very definitely rejected at every listed level of risk aversion. On the other hand, we cannot reject the middle ground of full information in Shirapur; in fact the J statistics (reported in Table III) for full information in Shirapur are strikingly similar to the J statistics for full information in Aurepalle.

The examples of Aurepalle and Shirapur might lead one to suppose that it's difficult to reject the full information model. Kanzara stands as a useful corrective to this supposition; full information is strongly rejected in Kanzara, at every level of risk aversion. The permanent income hypothesis is also firmly rejected, at all but the lowest levels of risk aversion. The private information hypothesis performs quite well, and cannot be rejected at any level of risk aversion.

This is probably not the ideal setting in which to try to pin down precise attitudes towards risk.⁹ However, it may be worth noting that the typical magnitude of relative risk aversion is toward the low end of estimates obtained by researchers using US data, at least for Aurepalle and Kanzara. It is interesting to speculate about the considerable differences in estimates of γ observed across villages. If we compare Shirapur and Kanzara, for example, we will see that the median estimate of γ in Shirapur, 1.3153, is nearly twice the median estimate in Kanzara, 0.6801. It is possible that this represents real, systematic differences in the preferences of households across the two villages. It is also possible, however, that the difference might be traceable to some observable difference between the two villages. A glance at Table II suggests that differences in expected income may be among the more salient differences between the two villages. Average household income per adult equivalent in Kanzara is 820 rupees, 46 per cent more than the average income of 560 rupees in Shirapur. If preferences do not really exhibit constant relative risk aversion, but rather slightly decreasing relative risk aversion, then our estimates of relative risk aversion should be lower for wealthier Kanzara.¹⁰

This supposition seems attractive on its face, but it neglects two other important facts. The first is that, while income in Kanzara is higher than in Shirapur, average consumption is nearly identical in the two villages. Since it is consumption, and not income, that is used in defining standard measures of risk aversion, this seems rather damning. The second problem is Aurepalle. Income (and consumption) are lowest in this village, yet the estimates of relative risk aversion are also the lowest. Differences in wealth combined with misspecification of preferences do not seem to account for the pattern of differences in relative risk aversion estimated across villages.

A more appealing explanation for the variation in estimated relative risk aversions across villages has to do with the possibility that some households within a particular village might be operating in a permanent income framework, while others might be operating in a private information

⁹There are two sorts of measurement error which are alleged to exist in the VLS consumption data. First, in the last years of the survey, certain categories of consumption appear to be underreported (Townsend 1994), and second, consumption from own production may be underreported (Gautam 1991). Evidence presented in Ligon (1996) suggests that estimates of b_0 are biased toward zero when either or both of these types of measurement error are present. Thus, while our inferences regarding regimes are if anything strengthened in the presence of measurement error, it may be that our estimates of γ shouldn't be taken too seriously.

¹⁰Using an altogether different test but the same dataset, Rosenzweig and Binswanger (1993) also find evidence that preferences in the VLS villages exhibit decreasing relative risk aversion. As in the present case, however, this hypothesis is not distinguishable from a hypothesis that households differ in their access to asset markets.

framework. There is certainly nothing in the theory to rule this out. This sort of heterogeneity certainly seems capable of invalidating the tests we conducted above. Furthermore, we have an example in which we seem to have made precisely this mistake. Since our results seem to indicate that Aurepalle is a permanent income village, while Shirapur and Kanzara are private information villages, pooling all three to arrive at a single estimate of b_0 is clearly a mistake. It seems possible that the somewhat confused looking pooled estimates in Table III are a direct consequence of this failure to respect the heterogeneity of the three villages.

If this explanation is correct, then we might suppose that Aurepalle, for example, while consisting mostly of households in the permanent income regime, has some proportion of households which actually belong to the private information regime. This heterogeneity would presumably bias our estimates of b_0 toward zero—we pursue this thesis further in Section 3.3.

3.3. Heterogeneous Regimes. Suppose that each household in our sample belongs to either the permanent income or private information regimes. We would like to find a method of estimation which would exploit equation (20) to not only estimate the preference parameter γ , but also to sort households into the two possible regimes. Combining the method of k -means (Pollard 1981; Pollard 1982b; Pollard 1982a) with the Generalized Method of Moments provides just such a method.

In order to implement the k -means part of the procedure, recall that if household i is a permanent income household, then

$$\mathbb{E}_t \xi_{i,t+1}^{-\gamma} - 1 = \mathbb{E}_t h(\xi_{it}, -\gamma) = 0,$$

while if the household is a private information household,

$$\mathbb{E}_t \xi_{i,t+1}^{\gamma} - 1 = \mathbb{E}_t h(\xi_{it}, \gamma) = 0.$$

We can use this difference to attempt to sort households into regimes. Define a set of indicator variables, one for each household i , by

$$(23) \quad \iota_i(\xi_{it}, b) = \begin{cases} 1 & \text{if } \|\sum_{t=1}^{T-1} h(\xi_{it}, b)\| < \|\sum_{t=1}^{T-1} h(\xi_{it}, -b)\| \\ 0 & \text{otherwise.} \end{cases}$$

Thus, so long as b is positive, we expect the indicator for household i to be one if household i is a private information household, and for the indicator to be zero otherwise.

The sample counterpart to our conditional moment conditions is

$$(24) \quad g(\xi, \zeta, b) = \frac{1}{N(T-1)} \sum_{i=1}^N \sum_{t=1}^{T-1} [\iota_i(\xi_{it}, b) h(\xi_{it}, b) + (1 - \iota_i(\xi_{it}, b)) h(\xi_{it}, -b)] \cdot \zeta_{it}.$$

We restrict the parameter space to be a compact subset of the positive real line, and use the generalized method of moments (Hansen 1982) to estimate b_0 (which may now be regarded as an estimate of the preference parameter γ , rather than as an estimator of either plus or minus γ), and set of indicators $\{\iota_i(\xi_{it}, b)\}$. Pollard (1981) establishes the consistency of the estimator of the preference parameter, while Pollard (1982a) establishes a central limit theorem for the estimator. Unfortunately, because each of these results relies on $N \rightarrow \infty$, it is not necessarily the case that the indicator variables produce the correct assignment of households to regimes, even asymptotically, though the assignments will be correct ‘on average,’ since $g(\xi, \zeta, \gamma)$ can be shown to converge to zero.

Application of this procedure gives the results in Table V for a variety of different instrument sets.¹¹ Estimated values of γ are fairly similar in Aurepalle and Shirapur for all but the last instrument set, ranging from 0.97 to 1.6. Estimates are somewhat lower in Kanzara, ranging from 0.69 to 1.0. Pooling data from all three villages yields estimates of γ ranging from 0.96 to 1.78. The magnitude of these estimated relative risk aversions is broadly consistent with experimental evidence collected by Binswanger (1981) in Aurepalle, and falls well within the range of estimated relative risk aversions from other microeconomic studies.

Several of the estimated values are not significantly different from zero; for these instrument sets we are unable to reject the hypothesis of full insurance. However, more complete instrument sets uniformly reject the full insurance hypothesis. The J statistics reported in Table V are asymptotically distributed χ^2 with degrees of freedom equal to the number of instruments minus one, and provide a test of our specification, which is not rejected for any instrument set.

This brings us to our main result, which is that most households in Shirapur and Kanzara appear to satisfy the private information restriction, while in Aurepalle a majority of households satisfy the permanent income restriction. In particular, in Aurepalle 12 of 33 households are assigned to the private information regime; in Shirapur all but one household is so assigned, while in Kanzara four households are assigned to the permanent income regime.¹² Thus, only in Aurepalle is there enough variation to make comparisons interesting. Perhaps surprisingly, regime assignment in Aurepalle does not seem to be associated with farm size except for the smallest landholders, all but one of whom belong to the permanent income regime. Nearly equal numbers of large farmers and landless households belong to the permanent income regime, and there are no large differences in the conditional means for any of the variables used in the estimation.

¹¹In addition to the variables mentioned in Table V, the instrument sets include a set of village dummy variables.

¹²For those interested in the particular identification of these households we list those households assigned to the permanent income regime. In Aurepalle, these are households 1, 5, 6, 9, 10, 30, 31, 32, 32, 33, 34, 35, 36, 37, 38, 41, 45, 46, 51, 56, 57, and 58; in Shirapur, household 50; in Kanzara, households 2, 35, 42, and 48.

Instrument Vector	Village			
	Aurepalle	Shirapur	Kanzara	All
	b ($\sigma(b)$) J	b ($\sigma(b)$) J	b ($\sigma(b)$) J	b ($\sigma(b)$) J
—	1.3644 (1.6905) 0	1.3862* (0.4306) 0	0.9487 (0.5849) 0	1.5206* (0.3282) 0.0128
Income	1.5644 (1.2412) 0.0409	1.4300* (0.4230) 0.0003	0.9956* (0.5144) 0.0165	1.7800* (0.3057) 0.0756
Income, Land	1.1644 (0.8172) 0.2718	1.4531* (0.4492) 0.1431	0.7612 (0.4409) 0.0689	1.7156* (0.3180) 0.3868
Family Size, Income, Land	1.0656* (0.4601) 0.0442	1.2250* (0.4278) 0.1840	0.6875 (0.4641) 0.0400	1.4469* (0.3031) 1.4179
Family Size, Income, Land, Rain, Avg Consumption	0.9675* (0.2080) 0.2074	1.5938* (0.2584) 1.7231	0.8237* (0.1507) 0.6804	0.9644* (0.1199) 0.9524

TABLE V. Estimates of risk aversion with heterogeneous regimes. Starred estimates are significant at the 95% level. Standard errors are reported in parentheses. Statistics reported under the columns labelled J are scaled to have an asymptotic χ^2 distribution.

4. CONCLUSION

We've developed a set of restrictions that allows one to distinguish among consumption allocations that obey the permanent income hypothesis, that exhibit full insurance, and that are constrained by private information. Tests of these restrictions in three villages in India indicate that of these three possible regimes, the private information restriction seems to provide the best explanation of consumption for nearly all households in two of these three villages. In the third village, more households have allocations broadly consistent with the permanent income restriction, but there is evidence that different households belong to different regimes in this village.

The inference that private information plays a key role in the villages studied here has some rather profound implications. The contractual "inter-linkages" which some authors (Bardhan 1980) have held to be a distinguishing characteristic of developing economies are a ubiquitous feature of these models, as is the implication that the economy ought to exhibit growing inequality over time. These models imply a very rich theory of distribution which begs for further investigation.

Several other recent empirical studies have used these same data to conduct various tests of risk sharing, and it seems appropriate to compare our findings with some of these earlier results. Townsend (1994) asks whether consumption allocations in these villages might not exhibit full insurance. While he finds a quite considerable amount of insurance, he rejects the hypothesis of full insurance; households' consumption responds to idiosyncratic income shocks (see also Ravallion

and Chaudhuri (1997)). Lim (1992) addresses the question of whether credit markets alone can explain the degree of smoothing observed in the VLS villages, and concludes that while the amount of smoothing is less than full, it is greater than could be delivered by credit markets alone. Private information constrained allocations certainly seem capable of delivering the observed amount of insurance.

A great deal of theoretical attention has been paid in recent years to the possibility that private information may play an important role in village settings. Despite the fact that much of this theory may have been inspired by stylized facts and institutions observed in village settings, whether or not this body of theory has any real application in development is, in our view, an empirical question. There are too many difficult issues for us to claim that this research has settled the question, but the evidence we present here should provide some encouragement to theorists. Future empirical work should take seriously the possibility that private information plays an important role in shaping the allocations and institutions of village economies.

REFERENCES

- Bardhan, P. K. (1980). Interlocking factor markets and agrarian development: A review of the issues. *Oxford Economic Papers* 32, 82–98.
- Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory* 16, 252–92.
- Binswanger, H. P. (1981). Attitudes toward risk: Theoretical implications of an experiment in rural India. *The Economic Journal* 91, 867–890.
- Browning, M. and C. Meghir (1981). The effects of male and female labor supply on commodity demands. *Econometrica* 59(4), 925–951.
- Deaton, A. (1992). *Understanding Consumption*. Oxford: Clarendon Press.
- Deaton, A. and C. Paxson (1994). Intertemporal choice and inequality. *Journal of Political Economy* 102(3), 437–467.
- Diamond, P. (1967). The role of stock markets in a general equilibrium model with technological uncertainty. *The American Economic Review* 57, 759–776.
- Gautam, M. (1991). Measurement errors and consumption analysis: Implications of using alternative estimates for ICRISAT VLS data. Manuscript.
- Gjesdal, F. (1982). Information and incentives: The agency information problem. *Review of Economic Studies* 49, 373–390.
- Hall, R. E. (1978). Stochastic implications of the life cycle–permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86, 971–987.
- Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica* 50, 1029–1054.
- Hayashi, F. (1992). Comment: On the estimation of panel-data models with serial correlation when instruments are not strictly exogenous. *Journal of Business and Economic Statistics* 10(1), 15–17.
- Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics* 10, 74–91.
- Ligon, E. (1996). Risk-sharing and information: Theory and measurement in village economies. <http://are.berkeley.edu/~ligon/Papers/reid.ps>.
- Lim, Y. (1992). *Disentangling Permanent Income from Risk Sharing: A General Equilibrium Perspective on Credit Markets in Rural South India*. Ph. D. thesis, University of Chicago.
- Lucas, Jr., R. E. (1991). On efficiency and distribution. Manuscript.

- Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association* 55.
- Pakes, A. (1994). Dynamic structural models: Problems and prospects. In C. A. Sims (Ed.), *Advances in Econometrics: Sixth World Congress*, Volume 1, Chapter 5. Cambridge: Cambridge University Press.
- Phelan, C. J. (1990). *Incentives, Insurance, and the Variability of Consumption and Leisure*. Ph. D. thesis, University of Chicago.
- Pollard, D. (1981). Strong consistency of k -means clustering. *Annals of Statistics* 9(1), 135–140.
- Pollard, D. (1982a). A central limit theorem for k -means clustering. *Annals of Probability* 10(4), 919–926.
- Pollard, D. (1982b). Quantization and the method of k -means. *IEEE Transactions on Information Theory* IT-28(2), 199–205.
- Ravallion, M. and S. Chaudhuri (1997). Risk and insurance in village India: Comment. *Econometrica* 65(1), 171–184.
- Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica* 53, 69–76.
- Rosenzweig, M. R. and H. P. Binswanger (1993). Wealth, weather risk and the composition and profitability of agricultural investments. *Economic Journal* 103, 56–78.
- Sharma, K. C. and R. P. Singh (undated). Rainfall and crop production in three villages of India's SAT. Undated manuscript.
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.
- Walker, T. S. and J. G. Ryan (1990). *Village and Household Economies in India's Semi-arid Tropics*. John Hopkins.
- Wilson, R. (1968). The theory of syndicates. *Econometrica* 36, 119–132.

AGRICULTURAL AND RESOURCE ECONOMICS UNIVERSITY OF CALIFORNIA, BERKELEY, CA 94720-3310