WORKING PAPER NO. 824

RISK-SHARING AND INFORMATION:
THEORY AND MEASUREMENT IN VILLAGE ECONOMIES

by

Ethan Ligon

RISK-SHARING AND INFORMATION: THEORY AND MEASUREMENT 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 nearly 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 in two of the three villages; the evidence for the third village is mixed.

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 may be the first to actually test tightly specified versions of the complete markets and permanent income models against an explicit private information alternative.

1.1. **Three economic regimes.** 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; 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). For the purposes of this research, this last contribution is of particular interest because of Hall's seminal role in proposing to test the implications (derived from the Euler equations) of the permanent income model. 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. In such models, consumption allocations may depend on all public information in a manner which is highly sensitive to the precise production technology and information structure of the economy (Holmström 1979).

1.2. **Consumption based tests of regimes.** It is probably not the case that models belonging to 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.

Hall's initial contribution to the permanent income literature spawned a voluminous empirical literature devoted to testing implications of the permanent income model (variously defined), including Hansen and Singleton (1982), Hayashi (1987), Zeldes (1989), and Runkle (1991).

The literature testing the full-insurance implications of complete markets models is comparatively recent; although Leme (1984) provides some early, informal evidence that nations do not fully insure each other, Mace (1991) and Cochrane (1991) were the first to explicitly seek a test of full insurance using data on individuals in the United States. Work by Altug and Miller (1990) and Hayashi, Altonji, and Kotlikoff (1996) also deserves special mention. Tests of full insurance in the villages of various developing countries have been undertaken by, most notably, Rosenzweig (1988), Townsend (1992), Morduch (1991), and Deaton (1990). Tests of full insurance within extended families in the United States have been conducted by Altonji, Hayashi, and Kotlikoff (1989). Lim (1992) deserves special mention for attempting to discriminate between the insurance implications of a particular full-insurance, Arrow-Debreu model and a permanent income model.

There have been few attempts to test the implications that models of the private information regime hold for consumption. This is at least in part due to the difficulty, alluded to above, of specifying the manner in which consumption ought to depend on other observables. Because

allocations in general may depend on all public information, there may be a serious specification problem for an econometrician who observes only a subset of the public information set, and who does not know what remaining information is public and what private. In the absence of detailed knowledge of the primitives of the economy under consideration, the hypothesis that private information plays a role in determining consumption may often play the role of an ill-specified alternative to some well-specified null hypothesis.

Among the tests of private information models that *have* been conducted, Phelan (1990) shows that the increase in consumption inequality in the United States is consistent with that predicted by a well-specified private information model, but is not able to distinguish very well between private information and permanent income regimes. Green and Oh (1991) describe some observable implications of private information models which might serve to distinguish them from permanent income models, but they do not examine any actual economies. The implications that they derive are both different from and weaker than those derived here.

1.3. **Organization.** 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. 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 can test which regime best describes the data. The private information regime clearly outperforms the alternatives in two of three villages; there is evidence that different households may belong to different regimes in the third. Such differences across households are 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 3 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 3.4 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 3.2 and 3.3 will allow us to judge which regime best fits the data. Section 4 discusses the data, methodology, estimation procedure, and results, in that order. Section 6 extends the methods of Section 4 to accomodate heterogeneous regimes, different interpretations of the role of possible intermediaries

(or principals) in the village, as well as addressing the role of aggregate risk and measurement error in our estimation procedure. Section 7 concludes.

2. Three Villages

It would be nice to know if the rather stylized models of this paper bear any resemblance to an actual village economy. The actual economies we take as examples are 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). 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 of the survey.

These villages are very poor by western standards. The median annual per capita income, averaged over the three villages and expressed in 1975 dollars, is only about $50; the share of this income spent on food is in excess of 60 per cent. This income includes income from all sources save transfers, including payments made in kind. Reflecting the importance of foodstuffs, agricultural land occupies a dominant position in the portfolio of assets held by villagers; the share of land in total assets ranges from 63 to 80 per cent (Walker and Ryan 1990). Agricultural produce accounts for over 70 per cent of all income in each of the villages; income from labor accounts for the bulk of the remainder.

Although all three villages rely on agriculture for the bulk of their income, the agriculture of the three villages is varied. In the first village, Aurepalle, the chief crops are sorghum, millet, castor, and paddy. In the second village, Shirapur, sorghum dominates production. Kanzara, the third village, uses a variety of intercropping techniques to grow cotton, pigeon-pea, and mung beans, and grows a hybrid sorghum as a sole crop. To some extent, these different cropping patterns may reflect differences in sophistication and access to outside markets; witness the intercropping of Kanzara, or the cash crops of castor and cotton in Aurepalle and Kanzara respectively. However, to a greater extent crop choice reflects exogenously given aspects of the environment each village finds itself in. Despite the fact that all three villages are located in central India, they differ considerably by soil type, rainfall, and of course geographic location. This last point is salient because different villages lie in different states of India, and thus are subject to differing government policies.

Differing physical environments and differing choices of cropping patterns mean that different tasks must be undertaken by farmers and farm workers. Some crops, particularly paddy, are

particularly labor-intensive; some tasks, such as weeding, may be difficult to monitor. In addition to strictly agricultural tasks, villagers practice a variety of trades and produce a variety of handicrafts.

The three villages studied here seem an excellent choice for testing models of insurance, because households in these villages face a great deal of idiosyncratic income risk. Of 33 households in Aurepalle continuously sampled over eight years, there were 17 instances in which a households per adult-equivalent income fell below half of the household's median income. There were also 17 such shortfalls in Shirapur, but for only 31 households; there were only eight such shortfalls for 36 households in rainfall assured Kanzara. One might suppose that these shortfalls tended to happen in response to some aggregate shock (which would limit the scope of mutual insurance), but they are fairly spread out across years. In Aurepalle, for example, although there were six shortfalls in 1976, there were three in each of 1977 and 1982, two in 1980, 1981, and 1983, and one in 1979. In Shirapur, there were four shortfalls in 1980, three in each of 1976, 1979, and 1983, two in 1982, and one in each of 1977 and 1981. Kanzara's eight shortfalls are spread out almost uniformly over 1978–81.

More formal tests of income covariance also indicate that there is considerably idiosyncratic risk, although aggregate sources of risk are clearly non-negligible. The average of Spearman's rank correlation coefficients for household income is 0.1481 in Aurepalle, 0.0494 in Shirapur, and 0.2893 in Kanzara. Calculation of these statistics permits us to construct a nonparametric test of the hypothesis that there is no cross-sectional correlation in realized incomes (Frees 1995); though this hypothesis is easily rejected at conventional levels of significance for each of the three villages, it seems that income correlations are rather small in magnitude, save perhaps for Kanzara.

If income risk creates scope for insurance mechanisms of some sort, is there evidence that such mechanisms are employed? In a first stab at answering this question, we compute the same rank correlation statistic we used a moment ago for income, but apply it instead to consumption. The statistics for consumption are surprisingly uniform across villages; 0.4604 for Aurepalle, 0.4905 for Shirapur, and 0.5059 for Kanzara. These are, again, significantly greater than zero, and also significantly greater than the corresponding statistics for income; there is definitely consumption smoothing going on.

There are two central questions we would like to answer about this consumption smoothing. Since the amount of smoothing we see is less than full insurance, we would like to know what model of consumption allocations correctly predicts the amount of smoothing observed in these villages. In this paper we pose this as a contest between models with full insurance, the permanent income hypothesis, and a class of models with private information.

The second question we would like to answer is how villagers accomplish the amount of consumption smoothing they do. There are a wide variety of possible institutional arrangements

which might support permanent income sorts of outcomes, and an at least equally large number which might support either full insurance or private information outcomes. In this paper, we will have considerably less success in addressing this second question than we will have in answering the first. Nonetheless, a description of some of institutional facets of consumption smoothing may be useful. A more detailed analysis of actual financial transactions is attempted in Section 6.1.

Several earlier studies have examined financial transactions using these same data. Bhende (1986) and Binswanger et al. (1985) use annual data on indebtedness to draw a picture of the supply and demand for credit in these villages; Binswanger and Rosenzweig (1986) use data on transactions to describe the sources and uses of credit. One of the conclusions one can draw from these papers is that there is that the actual financial arrangements employed in these villages are often rather complicated, and depend on a variety of contingencies. An obvious contingency which features in many descriptive accounts of credit involves the possibility of default. Default is common on loans from formal sources, but there is considerable evidence that both parties to such transactions actually expect default; though nominally loans, these transactions are actually more like transfers or subsidies from the government than they are like credit. Default is quite uncommon on financial transactions involving informal sources (which account for some 90 per cent of all loans received); however, loans may often be forgiven in whole or in part, or payments rescheduled. The addition of such contingencies as these makes these transactions look much more like bundles of contingent claims than the simple debt contracts usually called for under the permanent income hypothesis.

More evidence that actual contracts are considerably more complicated than simple debt is found in the fact that there are a variety of local names to describe different sorts of financial transactions, all of which tend to be compounded into the single category of "credit" in the data collected and described by economists. One particular sort of loan commonly seen in Aurepalle bears the local name *nagu*. A *nagu* loan is issued by a moneylender at the beginning of the cropping season, either for consumption purposes or to finance the purchase of crop inputs. The nominal interest rates on such loans are quite high; around 12 per cent per month. Loans (and rates of interest) are denominated in units of grain. The length of the crop season depends on the crop, and ranges from about three to about six months. The moneylender apparently does not attempt to closely monitor cropping operations during the growing season. At the end of the season, the crop is harvested, and must be threshed. At this time, the moneylender goes to the threshing floor, often in person, to claim a portion of the crop. If the crop is good, he will collect the original principal and interest; however, if the crop is poor, the moneylender may choose to either forgive the loan, or to reschedule payment. Binswanger et al. (1985) reports that most farmers are involved in a long term relation with a moneylender; at the time of their survey, the average length of relationship with the same moneylender was nearly a decade.

Though necessarily incomplete, this descriptive account of some of the financial arrangements made by villagers seems to be consistent with any of the three regimes we wish to test. There are commonly used contracts which exhibit features of both credit and insurance. There is evidently a considerable amount of consumption smoothing. To decide which regime best describes these village economies, we must answer harder questions than "is there credit," or "is there insurance," or even "how much insurance is there?"

3. 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]

In Section 3.1 we establish some notation and general structure common to all three models. Section 3.2 derives a partial characterization of consumption under the permanent income hypothesis, and makes some general remarks about the behavior of consumption in a permanent income world. Section 3.3 derives and discusses restrictions on consumption implied by full insurance, while Section 3.4 does the same for a broad class of models with private information.

3.1. Preliminaries.

The environment. The environment has a principal (sometimes labeled agent 0), and a finite set of agents indexed by $i = 1, 2, ..., n$. Time is discrete and is indexed by $t = 0, 1, 2, \ldots, T$. 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 of next period's 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 $t = 0, 1, \ldots, T$ 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

---

[1]Generalizations 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 with a taste for the macabre 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.

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

$$(1) \qquad u_i = \sum_{t=0}^{T} \beta^t \Big[U^i(c_{it}) - V^i(a_{it}) \Big]$$

where $U^i : C \to \mathbb{R}$ and $V^i : A \to \mathbb{R}$ are continuously differentiable, strictly increasing, and weakly concave. Note that we can think of the principal here as just another agent; in particular, the principal may take some action each period, just as the agents do. 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. 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 addition, agents 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 more somewhat more natural to think of this technology as some sort of credit.

3.2. Permanent Income.

The permanent income hypothesis enjoys particular stature among the hypotheses we consider in this paper. 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]. Another feature distinguishes the permanent income hypothesis from the other hypotheses I'll describe. 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's resources at each date-state are given by the sum of his savings, $(1 + r_t)b_{it}(h_{t-1})$, his current production of the consumption good, $x_{it}(\omega_t)$, and possibly some transfers from other households, $\tau_{it}(h_t)$. Given these resources, the agent must decide how much to save, consuming the rest, and must also decide what action to take.

Thus, the agent solves

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

such that the agent's resource constraint is respected,

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

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

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

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

$$(5) \qquad \mathrm{E}\left(\left. \frac{U^{i\prime}(c_{i,t+1}(h_{t+1}))}{U^{i\prime}(c_{it}(h_t))} \right| h_t \right) = \frac{1}{\beta(1 + r_t)}$$

where $\mathrm{E}(\cdot | h_t)$ 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).

3.3. **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 $\underline{u}_i$, $i = 1, 2, \ldots, n$. 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 savings $\{b_{i0}(h_{-1})\}_{i=0}^n$. Here we pose the problem as one in which the principal chooses 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, actions, and savings 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

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

subject to a set of resource constraints,

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

for all $h_t \in H$, and

$$(8) \qquad c_{it}(h_t) = (1 + r_t)b_{it}(h_{t-1}) + x_{it}(\omega_t) - b_{i,t+1}(h_t) + \tau_{it}(h_t)$$

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

$$(9) \qquad b_{i,T+1}(h_T) = 0 \qquad \text{for all } i = 0, 1, 2, \ldots, n \text{ and for all } h_T \in H_T.$$

In addition, reservation utilities must be respected:

$$(10) \qquad \sum_{t=0}^{T} \beta^t \sum_{h_t \in H_t} \Pr(h_t)[U^i(c_{it}(h_t)) - V^i(a_{it}(h_t))] \geq \underline{u}_i \qquad i = 1, \ldots, n.$$

The objective function is simply agent 0's discounted, expected utility, which she maximizes subject to providing utility of at least $\underline{u}_i$ to agent i, and subject to satisfying an intertemporal resource constraint given by (7), (8), and (9). The agents' 'flow' resource constraints (8) constrain the agent to consume no more than the sum of output $x_{it}(\omega_t)$, transfers $\tau_{it}(h_t)$ from the principal and savings $(1 + r_t)b_{it}(h_{t-1})$ from the previous period, minus savings this period. 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 with respect to $\{\tau_{it}(h_t)\}$ and $\{c_{it}(h_t)\}$ gives us

$$U^{0\prime}(c_{0t}(h_t)) = \lambda^i U^{i\prime}(c_{it}(h_t)) \qquad i = 1, \ldots, n; \, h_t \in H,$$

where the λ^i are the n Lagrangian multipliers associated with (10). Since this relationship holds for all histories, we can combine first order conditions for agent i after history h_t and after h_{t+j} to get

(11)
$$\frac{U^{0\prime}(c_{0,t+j}(h_{t+j}))}{U^{0\prime}(c_{0t}(h_t))} \frac{U^{i\prime}(c_{it}(h_t))}{U^{i\prime}(c_{i,t+j}(h_{t+j}))} = 1,$$

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).

In addition to (11), we have the first order conditions with respect to $b_{it}(h_{t-1})$, which yield the familiar Euler condition

(12)
$$U^{i\prime}(c_{it}(h_t)) = \beta(1 + r_t)\mathrm{E}[U^{i\prime}(c_{i,t+1}(h_{t+1})) \,|h_t].$$

Equation (12) 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 (12) 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.

3.4. **Private Information.** In this section we derive a characterization of any constrained efficient consumption allocation when the agent has private information regarding production. However, before proceeding to the derivation, some motivation is in order.

Consider an economy with moral hazard in which the principal and a single agent share information regarding output x_{it}, but only the agent knows his level of effort a_{it}. Suppose for further simplicity that there are only two periods: an initial period with a known, fixed endowment in which the agent takes some action which affects the distribution of output in the subsequent period. Let the state in the second period be given by an output action pair (x, a).

If there were no moral hazard problem, then the optimal allocation would be such that the agent's marginal rate of substitution would be equal to the principal's in every date and state. But in the second period the principal observes only the output x, and must infer the agent's action.

Both agent and principal are better off with more insurance rather than less, but any insurance arrangements the two make must be contingent only on public information in order to respect incentives.

Since the principal's information set is a subset of the agent's, the solution to the optimal provision of incentive compatible insurance is to insure so that marginal rates of substitution are equal *in expectation*, where expectations are conditioned on the principal's information set. Put another way, the principal, not knowing the action taken by the agent, cannot know precisely what the state is. The optimal allocation makes efficient use of the principal's coarser information set, and equalizes ratios of marginal utilities up to her imperfect knowledge of the true state.

In order to move beyond this simple example, we will think about the problem of finding constrained-efficient allocations as a simple extension of the problem of finding the fully efficient allocations of the full insurance regime. In particular, for any private information problem, we can use revelation principle arguments to justify simply adding some incentive compatibility constraints to the program given above (cf. Townsend (1987)).

The characterization promised above is a result due to Rogerson (1985). The idea behind the derivation is that the principal, under whatever contract, will neither make the agent worse off nor affect the agent's incentives if she takes y utils of consumption away today, so long as she gives back y/β utils of consumption with certainty tomorrow.

Proceeding more formally, 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 in the present instance is that the actions $\{a_{it}(h_t)\}$ 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:

$$(13) \qquad \{a_{it}(h_t)\} \in \operatorname*{argmax}_{\{a_{it}(h_t)\}} \sum_{t=0}^{T} \beta^t \sum_{h_t \in H_t} \Pr(h_t) \left[U^i(c_{it}(h_t)) + V^i(a_{it}(h_t)) \right].$$

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 at t which effects the distribution of output at $t+1$. Equation (13) certainly captures this possibility. However, it's important to recognize that (13) 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; (13) 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 will 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 (13) by taking away some utility y in one period, but then returning y/β utils (with certainty) in the following period. The modified incentive compatibility condition then becomes

$$(14) \quad \{a_{it}(h_t)\} \in \operatorname*{argmax}_{\{a_{it}(h_t)\}} \sum_{t=0}^{T} \beta^t \sum_{h_t \in H_t} \Pr(h_t)[y_{it}(h_{t-1})/\beta - y_{it}(h_{t-1})]$$
$$+ \sum_{t=0}^{T} \beta^t \sum_{h_t \in H_t} \Pr(h_t) \left[U^i(c_{it}(h_t)) + V^i(a_{it}(h_t)) \right].$$

Since this modification simply adds a constant to the objective function (the agent, recall, takes this intertemporal redistribution as given), the incentive compatibility constraints are utterly unaffected. 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[2]

$$(15) \quad \max_{\{\{y_{i,t+1}(h_t),z_{it}(h_t),a_{it}(h_t)\}_{i=0}^{n},c_{0t}(h_t),b_{0,t+1}(h_t)\}_{h_t \in H}} \sum_{t=0}^{T} \beta^t \sum_{h_t \in H_t} \Pr(h_t)[U^0(c_{0t}(h_t)) - V^0(a_{0t}(h_t))]$$

subject to the resource constraints

$$(16) \quad c_{0t}(h_t) = (1 + r_t)b_{0t}(h_{t-1}) - b_{0,t+1}(h_t) + x_{0t}(\omega_t)$$
$$+ \sum_{i=1}^{n} \left[x_{it}(\omega_t) - U^{i-1}(z_{it}(h_t) + y_{it}(h_{t-1})/\beta - y_{i,t+1}(h_t)) \right]$$

for all $h_t \in H$ and subject also to the requirement that debt be repaid:

$$(17) \qquad\qquad b_{0,T+1}(h_T) = 0 \qquad \text{for all } h_T \in H_T.$$

In addition, reservation utilities must be respected:

$$(18) \qquad \sum_{t=0}^{T} \beta^t \sum_{h_t \in H_t} \Pr(h_t)[z_{it}(h_t) - V^i(a_{it}(h_t))] \ge \underline{u}_i \qquad i = 1, \ldots, n;$$

---

[2]Note that we've also modified the earlier problem by permitting only the principal to use credit markets. There is no loss of generality, so long as the principal can either observe or infer $c_{it}(h_t)$.

and of course, recommended actions must satisfy (14) (equivalently, (13)) 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 order conditions for this problem. First order conditions with respect to the $y_{i,t+1}(h_t)$ and $c_{0t}(h_t)$ yield

$$U^{0\prime}(c_{0\tau}(h_t))U^{i^{-1}\prime}(z_{it}(h_t)) = \mathrm{E}\left[U^{0\prime}(c_{0,t+1}(h_{t+1}))U^{i^{-1}\prime}(z_{i,t+1}(h_{t+1})) \middle| h_t \right].$$

The inverse function theorem states that, for any differentiable, invertible function $f : \mathbb{R} \to \mathbb{R}$, such that $y = f(x)$:

$$f^{-1\prime}(y) = \frac{1}{f'(x)}$$

where $f'(x)$ exists. Applying this fact to the first-order conditions above, we have:

$$\frac{U^{0\prime}(c_{0t}(h_t))}{U^{i\prime}(c_{it}(h_t))} = \mathrm{E}\left[\frac{U^{0\prime}(c_{0,t+1}(h_{t+1}))}{U^{i\prime}(c_{i,t+1}(h_{t+1}))} \middle| h_t \right]$$

or

(19)
$$\mathrm{E}\left[\frac{U^{0\prime}(c_{0,t+1}(h_{t+1}))}{U^{0\prime}(c_{0t}(h_t))} \frac{U^{i\prime}(c_{it}(h_t))}{U^{i\prime}(c_{i,t+1}(h_{t+1}))} \middle| h_t \right] = 1.$$

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 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 (12)), in both of which intertemporal marginal rates of substitution are constant across agents.

A number of other interesting facts about private information allocations can be learned from equation (19). For example, suppose that the consumption awarded to agent i today is a function of the state. We can see from (19) that higher compensation today implies higher expected consumption relative to the principal in *all* future periods (such allocations are said to have 'memory' or to exhibit 'history dependence'). If future expected consumption changes each period, then so does the distribution of future utility; from this, a non-trivial theory of the distribution of wealth emerges (e.g., Atkeson and Lucas (1992), Thomas and Worrall (1990), Phelan (1995)).

Equation	$E\left(\frac{U^{i\prime}(c_{i,t+1}(h_{t+1}))}{U^{i\prime}(c_{it}(h_t))} \middle| h_t\right) = \frac{1}{\beta(1+r_t)}$	$\frac{U^{0\prime}(c_{0t}(h_t))}{U^{i\prime}(c_{it}(h_t))} = E\left[\frac{U^{0\prime}(c_{0,t+1}(h_{t+1})}{U^{i\prime}(c_{i,t+1}(h_{t+1})} \middle| h_t\right]$
Full Information	$\times$	$\times$
Private Information		$\times$
Permanent Income	$\times$	

TABLE I. Distinctions between models.

We have written rather blithely as though equation (19) characterized all multi-period economic models which incorporate private information. This is not true, of course; in particular, (19) characterizes only those allocations in which the principal controls the agents' consumption. Loosely, for (19) to apply it is sufficient for the principal to be able to specify a value for agents' marginal utility of consumption along every possible history. So, for example, the consumptions predicted by the model of Atkeson (1991) are not characterized by (19), because the agents in that model have private information regarding what of their income they consume and what part they invest. Equation (19) *does* characterize Atkeson and Lucas (1992), though the authors question the ability of real-world 'principals' to verify agents' consumption. Similarly, consumption in the models of Spear and Srivastava (1987), Thomas and Worrall (1990), and Phelan and Townsend (1991) all satisfy (19). Models which (perhaps implicitly) put limits on the space of controls available to the principal do not in general satisfy (19): the model of Stiglitz and Weiss (1981) and the example given by Townsend (1982) are instances of such models. Models which are not conveniently formulated as principal-agent models may not satisfy (19).

Summarizing, then, w've derived two conditions, unique combinations of which characterize each of three different regimes: the permanent income regime is characterized by (12) but *not* (19), the private information regime is characterized by (19) but *not* (12), and the full insurance regime is characterized by both (12) *and* (19). See Table I.

One potential objection should be dealt with before we proceed. Tests of the permanent income regime, several of which are cited above, have often tested versions of (12) against an alternative in which agents face some sort of borrowing constraint. When these constraints are binding, the implication that agents' intertemporal rates of marginal substitution are equal is lost. Can we distinguish private information allocations from some borrowing constrained version of a permanent income economy? The answer, on one level, is easy: borrowing constrained versions of permanent income models do not yield restrictions that look anything like (19). But more intuitively, the private information allocations we examine will typically feature agents who are not borrowing constrained at all, but rather are *savings* constrained!

To understand these savings constraints, consider the special case in which the principal is risk-neutral, and that the interest rate $r_t = 1/\beta - 1$.[3] We shall call this the 'strict' version of the private information model. This risk neutrality combined with equation (19) implies that

$$\mathrm{E}\left[\left.\frac{U^{i\prime}(c_{it}(h_t))}{U^{i\prime}(c_{i,t+1}(h_{t+1}))}\right| h_t\right] = 1.$$

Now since $1/x$ is a convex function, we can simply apply Jensen's inequality to this expression to bound the agents' intertemporal marginal rate of substitution from below:

$$\mathrm{E}\left[\left.\frac{U^{i\prime}(c_{i,t+1}(h_{t+1}))}{U^{i\prime}(c_{it}(h_t))}\right| h_t\right] \geq 1 = \frac{1}{\beta(1+r_t)}.$$

Since, for agents $i = 1, 2, \ldots, n$, the intertemporal marginal rate of substitution is greater than the intertemporal marginal rate of transformation, the agents will want to save at the prevailing interest rate; the principal's control over their consumption prevents them from doing so. The result that agents will be savings constrained extends to the case of a risk-averse principal and a rate of return equal to the principal's marginal rate of substitution, using a similar argument, but subject to a condition that consumption of principal and agent not be too highly (positively) correlated. We might expect high correlation if the principal faces a great deal of consumption uncertainty and if there are only a small number of agents; in this instance, an individual agent may provide the principal with significant consumption insurance.

The result that agents are savings constrained may appear to fly in the face of much of the conventional wisdom regarding credit in developing countries. However, in another guise, the result may actually seem rather familiar. Braverman and Stiglitz (1982) derive a similar result in a much less general model, but describe the result as one in which the principal subsidizes loans to the agent, causing the agent to borrow more than he otherwise would, and also causing to work harder than he would otherwise, simply to service his debt. Certainly the development literature is rife with anecdotes of peasants kept in perpetual debt by their patrons. This could be one way of satisfying equation (19).

4. INFERENCE AND ESTIMATION

4.1. **Methods.** The general idea is to test which of equations (5) or (19) 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.

---

[3]This pair of assumptions has been fairly standard in the literature which deals with repeated moral hazard. See for example Green (1987), Spear and Srivastava (1987), or Phelan and Townsend (1991). Indeed, so long as the principal is risk neutral, as she is in the vast bulk of the literature, and the distribution of output is stationary, then the assumption that the interest rate is equal to the principal's rate of time preference is crucial if the solution is to be bounded.

We begin by testing 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 $U^{0\prime}(c)$ is just some constant, equation (19) simplifies to become[4]

$$(20) \qquad \mathrm{E}_t \left[\frac{U^{i\prime}(c_{it})}{U^{i\prime}(c_{i,t+1})} \right] = 1.$$

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

$$\mathrm{E}_t \left[\frac{U^{0\prime}(c_{0,t+1})}{U^{0\prime}(c_{0t})} \right] = 1 = \frac{1}{\beta(1 + r_t)}.$$

This equation is only satisfied when $r_t = 1/\beta - 1$, and so permits us to solve for the interest rate.[5] Accordingly, the Euler equation for each agent becomes

$$(21) \qquad \mathrm{E}_t \left[\frac{U^{i\prime}(c_{i,t+1})}{U^{i\prime}(c_{it})} \right] = 1.$$

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

$$\frac{U^{i\prime}(c_{it})}{U^{i\prime}(c_{i,t+1})}$$

from (20) for its reciprocal in (21) yields the characterization of the private information regime. It is *critical* to note that, despite their close relationship, Jensen's inequality implies that (20) and (21) 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 (21) and (20) are, respectively,

$$(22) \qquad \mathrm{E}_t \left[\left(\frac{c_{i,t+1}}{c_{it}} \right)^{-\gamma} - 1 \right] = 0$$

---

[4]Because 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 $\mathrm{E}(\cdot|h_t)$ becomes $\mathrm{E}_t(\cdot)$.
[5]Of 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.

and

$$(23) \qquad \mathrm{E}_t\left[\left(\frac{c_{i,t+1}}{c_{it}}\right)^{\gamma} - 1\right] = 0$$

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

The symmetry of the restrictions (22) and (23) suggests an easy way to choose between the hypotheses they characterize. We simply nest both of (22) and (23), estimating b_0[6] in

$$(24) \qquad \mathrm{E}_t\left[\left(\frac{c_{i,t+1}}{c_{it}}\right)^{b_0} - 1\right] = 0.$$

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.

Green and Oh (1991), among others, make the point that any model which carries implications regarding the time series properties of consumption had better predict that consumption is close to a martingale in order to be consistent with the evidence. Since Hall (1978), it has been well known that the permanent income hypothesis is consistent with martingale consumption. However, it is not obvious from equations (22) and (23) that consumption under the permanent income and private information regimes possess similar time-series properties. Hall made the point that if preferences are quadratic and interest rates are constant, then we might expect consumption to follow a random walk under the permanent income regime; perhaps it suffices here to point out that if preferences are logarithmic ($\gamma = 1$) and the principal's consumption is constant, then (23) implies that consumption follows a martingale under the private information regime.

We have said little in this section regarding the full insurance hypothesis. It was earlier demonstrated that *both* (22) and (23) 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.

---

[6]The 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, and context ought to make clear which one we have in mind at any particular point.

4.2. **Estimation.** We remarked above that we could use the conditional moment restriction (24) 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. This section delves into the details of how this estimation is to be performed.

Let

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

and let

$$h(\xi_{i,t+1}, b) = \xi_{i,t+1}^b - 1$$

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

$$(25) \qquad\qquad\qquad E_t h(\xi_{i,t+1}, b) = 0.$$

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 $h(\xi_{i,t+1}, \gamma)$, so that the unconditional expected value of the product of this vectors with the true forecast error ought to be equal to zero, or

$$(26) \qquad\qquad\qquad E\left[h(\xi_{i,t+1}, b_0) \cdot \zeta_{it}\right] = 0.$$

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. We have no assurance, however, that forecast errors are independent across agents. To the contrary; we can easily imagine that aggregate unexpected events (e.g. drought, changes in government policy) might induce correlation across agents within the village at a point in time. That this possible conditional dependence across agents does not necessarily compromise the consistency of our estimator has been made clear by Pakes (1991); nonetheless, we will proceed cautiously, and will give an explicit treatment of the case in which errors are dependent in Section 6.3, where it is taken up at length.

There are three obvious ways to construct a sample counterpart to (26). Let our panel dataset consist of T observations on N households. We could construct N sets of q sample moments by averaging over time; we could construct $T-1$ sets of sample moments by averaging over households; or we could construct a single set of sample moments by averaging over both time and households. Constructing a single set of moment conditions is the most conservative approach when we are unsure of the degree of cross-sectional dependence in our sample. Because we ultimately care most about the sign of our estimate of b_0, and care little for the precise magnitude of the estimate, we

adopt this conservative route. See Pakes (1991) for an excellent discussion of the issues involved. This doubly-averaged sample counterpart to (26) is given by

$$(27) \qquad g(\xi, \zeta, b) = \frac{1}{N(T-1)} \sum_{i=1}^{N} \sum_{t=1}^{T-1} h(\xi_{i,t+1}, b) \cdot \zeta_{it}.$$

For the moment, we optimistically assume that forecast errors are independent in cross-section, and estimate the covariance matrix of $g(\xi, \zeta, b)$ as

$$(28) \qquad W(\xi, \zeta, b) = \frac{1}{N(T-1)} \sum_{i=1}^{N} \sum_{t=1}^{T-1} [h(\xi_{i,t+1}, b) \cdot \zeta_{it}][h(\xi_{i,t+1}, b) \cdot \zeta_{it}]'.$$

We assume that $W(\xi, \zeta, b_0)$ converges to some fixed matrix Ω.

Now consider the function

$$(29) \qquad J(\xi, \zeta, b) = N(T-1)g(\xi, \zeta, b)'W(\xi, \zeta, b)^{-1}g(\xi, \zeta, b)$$

and choose b to solve

$$(30) \qquad \min_{b \in \Gamma} J(\xi, \zeta, b),$$

where Γ is some compact subset of $\mathbb{R}$, and is assumed to contain b_0.

This estimator is closely related to the Generalized Method of Moments estimator of Hansen (1982), but differs in two important regards. The first difference has been discussed above; in thinking about the asymptotic behavior of the estimator defined here, we want to let both T and N approach infinity. The second difference has considerable practical import; in performing the minimization in (30), we are *simultaneously* computing both the sample moment conditions $g(\xi, \zeta, b)$ and the weighting matrix $W(\xi, \zeta, b)^{-1}$. Although different from, say, the two step estimator proposed by Hansen and Singleton (1982), this is a reasonably behaved estimator (Hansen, Heaton, and Yaron 1996); emerging Monte Carlo evidence suggests that its small sample properties are superior to those of the usual two step or iterated GMM estimators.[7]

---

[7]An interesting fact is that the two step estimator won't work as desired for the moment conditions proposed here. Consider the case in which the parameter space we're interested in is some compact real interval which includes the origin. By *examination* of (24), it's obvious that the GMM estimator of b_0 is zero. Since this clearly isn't the solution we're interested in, we exclude a small neighborhood of zero from the parameter space. If we were to employ the two-step method of Hansen and Singleton (1982), we would have to reparameterize equation (24) (by, say, dividing through by b) as in, e.g., Eichenbaum and Hansen (1990). The two step estimator is not invariant to such reparameterizations; the simultaneous estimator is. Also, it's easier—and asymptotically equivalent—to use the simultaneous criterion function given above; here, as b approaches zero, the determinant of the weighting matrix approaches infinity. In practice these two precisely balance each other out, and produce a perfectly well behaved criterion function over a suitably defined compact interval; see the example given in Figure 1.

Because $g(\xi, \zeta, b)$ is continuously twice differentiable in b, we obtain asymptotic normality and an expression for the variance of b:

$$\text{var}(b) = \left[g_b(\xi, \zeta, b)' \Omega^{-1} g_b(\xi, \zeta, b) \right]^{-1},$$

where $g_b(\xi, \zeta, b)$ denotes the partial derivative of $g(\xi, \zeta, b)$ with respect to b, and is given by

$$g_b(\xi, \zeta, b) = \frac{1}{N(T-1)} \sum_{i=1}^{N} \sum_{t=1}^{T-1} \frac{\partial h(\xi_{i,t+1}, b)}{\partial b} \cdot \zeta_{it}.$$

In addition to obtaining estimates of b_0 and the distribution of the estimator, the minimized value of the criterion function is asymptotically distributed χ^2_{q-1}. For $q > 1$ we can interpret the value of this statistic as a test of the $q - 1$ over-identifying restrictions, and think of this as a test of the underlying model.

5. RESULTS

The main result of the paper is that the private information regime provides a uniformly better fit to the data than does the permanent income regime in two of three villages (Shirapur and Kanzara). This finding seems remarkably robust, and is true for a variety of instrument sets. Although clearly consumption is best described by the private information regime, consumption is relatively well insured in Kanzara.

Some very interesting puzzles remain. The three villages in the sample differ with regard to how well the model fits the data. It may be that there is heterogeneity within villages, with some agents' consumption best characterized by the private information regime, and others best characterized by the permanent income regime. There is certainly nothing in the theory which prevents this, but it would certainly tend to bias the results of our estimation.

A word regarding the selection of instruments seems in order before proceeding. In our estimation framework, one's choice of an instrument set bears not only on estimation, but also upon the likelihood of rejecting the model. The best choice of a set of instruments for efficient estimation may well differ from that used to test the model; in the former case, one would want to choose variables used by the agents themselves in making their forecasts; in the latter case, one might well prefer to choose an instrument set whose elements are correlated with an agent's forecast errors only under some alternative hypothesis.

Some of the instruments employed have been discussed above; these 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 affected primarily by shocks peculiar to the household; we shall also consider variables reflecting only aggregate shocks. A set of dummy variables designed

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

TABLE II. Means and standard deviations of instruments

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

TABLE III. Estimates of risk aversion under the private information regime. 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.

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.
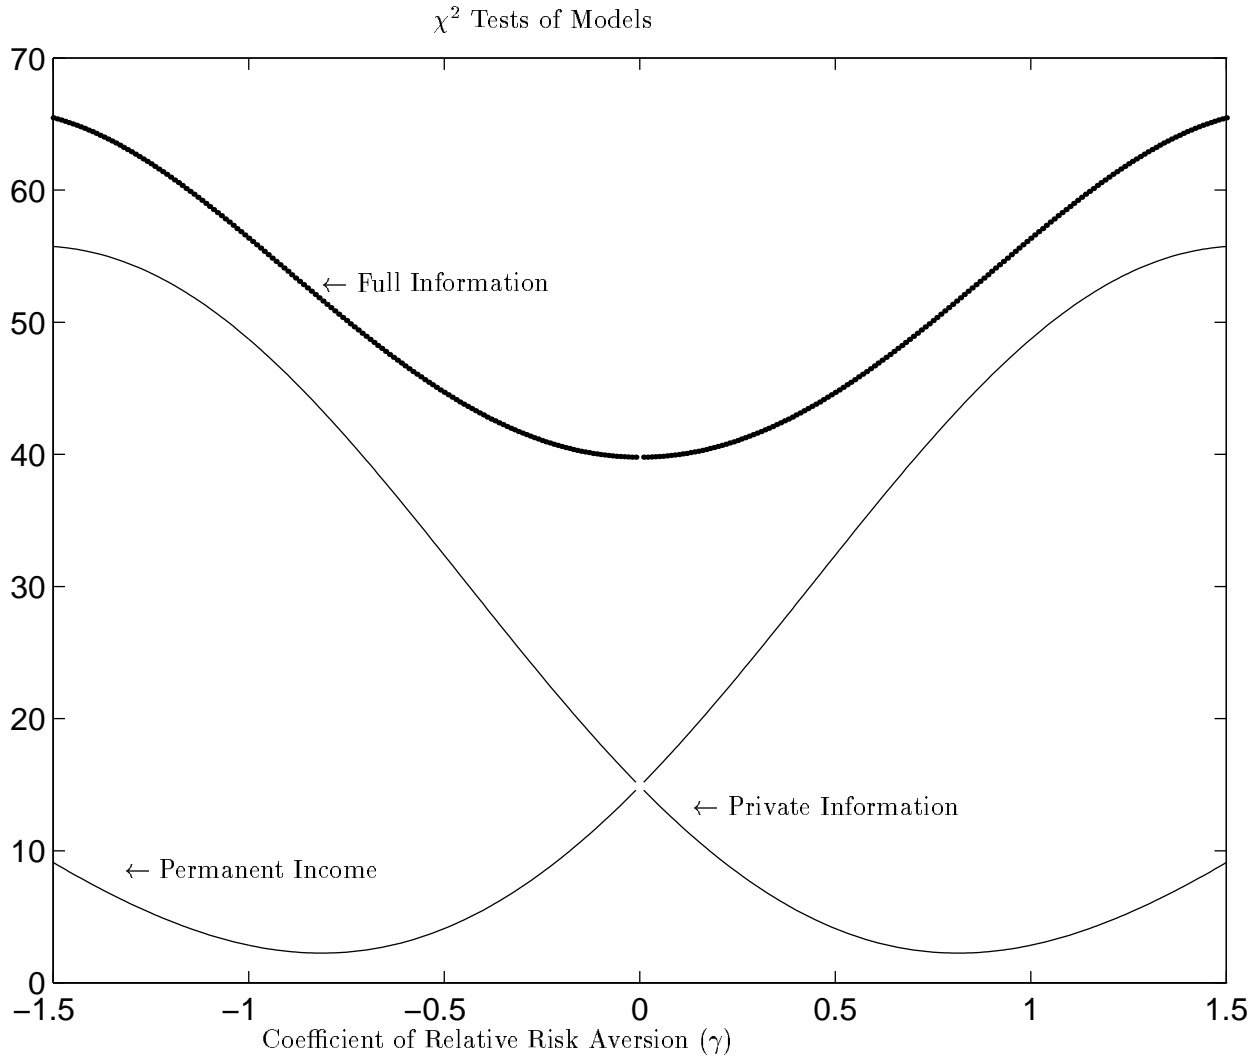
χ^2 Tests of Models

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

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 1.1313 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. 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^2_{q-1}, where q is the number of moment restrictions we test.[8] 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.[9] 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 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

---

[8]For 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.

[9]The 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.

to differentiate the two models given a particular value of γ.[10] 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 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 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 seems to be be 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, though again we are unable to reject the full information model. The permanent income hypothesis is very definitely rejected at every listed level of risk aversion. On the other hand, the we cannot reject the middle ground of full information in Shirapur; in fact the J statistics 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 dramatically 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.

---

[10] We 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

This is probably not the ideal setting in which to try to pin down precise attitudes towards risk. 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.[11]

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

---

[11] Using 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.

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 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 6.1.

6. Caveats and Extensions

6.1. 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 (24) 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

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

while if the household is a private information household,

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

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

$$(31) \qquad \iota_i(\xi_{i,t+1}, b) = \begin{cases} 1 & \text{if } ||\frac{1}{T}\sum_{t=1}^{T-1} h(\xi_{i,t+1}, b)|| < ||\sum_{t=1}^{T-1} h(\xi_{i,t+1}, -b)|| \\ 0 & \text{otherwise.} \end{cases}$$

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 was earlier given by equation (27). Let us redefine the sample moment in question, using instead

$$(32) \qquad g(\xi, \zeta, b) = \frac{1}{N(T-1)} \sum_{i=1}^{N} \sum_{t=1}^{T-1} [\iota_i(\xi_{i,t+1}, b) h(\xi_{i,t+1}, b) + (1 - \iota_i(\xi_{i,t+1}, b)) h(\xi_{i,t+1}, -b)] \cdot \zeta_{it}.$$

We also restrict the parameter space to be a compact subset of the positive real line. With these two modifications (and the obvious modification to our procedure for estimating the weighting matrix in equation (29)), estimation proceeds just as it did in Section 4.2.

This procedure leaves us with a point estimate b (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 $\{\iota_i(\xi_{i,t+1}, b)\}$. 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 $N \to \infty$, 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.[12]

Application of this procedure gives the results in Table V. After accounting for possible heterogeneity in regimes, estimates of risk aversion show considerably less variation across villages than previously. In particular, estimated values in Aurepalle and Shirapur now appear to be roughly similar, though estimates of γ in Kanzara remain relatively low. Pooled estimates seem to exhibit greater consistency across instrument sets, as well, providing some evidence that heterogeneity of regimes may be important.

How much heterogeneity is revealed by this method of estimation? In Aurepalle, twelve 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.[13] 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 perhaps 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,

---

[12] A detailed proof of these assertions is available from the author on request.

[13] For those interested in the particular identification of these households we list those households assigned to the permanent income regime. In Aurepalle, 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 b $(\sigma(b))$ J	Shirapur b $(\sigma(b))$ J	Kanzara b $(\sigma(b))$ J	All 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.

and there are no overwhelming differences in the conditional means for any of the variables used in the estimation.

A chief virtue of the methods we've used to test what regime different households belong to is that we haven't had to take a stand on precisely what set of contracts or other arrangements households are involved in. We simply look at outcomes. Nonetheless, any consumption allocation must in fact be effected by some set of contractual arrangements, and we might wonder if there appear to be systematic differences between the arrangements employed by permanent income households versus the arrangements of private information households.

We have data on realized transfers of goods and money for each household in our sample, though we know little about the households' partners for these transactions, and nothing at all about how any given transaction might be related to others (we cannot, for example, identify an outflow of cash as repayment for an earlier loan). On a very narrow reading of the models, we might expect that permanent income households would use only spot markets and markets for non-contingent credit. Transactions in our data are not, however, classified so that we can examine this hypothesis. The data do, however, identify particular payments and receipts as "credit" transactions. Table VI helps to classify these transactions for Aurepalle. Each transaction is classified according to whether it is a receipt or an outlay, what 'type'[14] the other party is, whether or not the transaction

---

[14] The table gives four possible categories for type. The data is only slightly richer. What we call "banks" is intended to include any formal credit transaction, and includes credit transactions with the government (including "fair price

is in cash or in kind, and finally whether the household is classified as belonging to the permanent income or private information regime. The numbers in the table report, first, the mean number of transactions of each sort per household (over the entire period of the survey); and second the mean value, in current rupees, of transactions of each sort (reported in parentheses). Thus, for example, the most common sort of credit transaction aside from 'other' was in kind receipts from moneylenders to permanent income households; these households averaged 19.30 such transactions over the ten year period 1975–84.

Probably the most striking feature of Table VI is that the overwhelming majority of transactions in Aurepalle are informal, with partners in the "moneylender" category. Because of the small number of households and the wide variety of sizes of transactions, it is somewhat difficult to know when differences across regimes are important. A casual examination of the data in Table VI suggests that private information households are slightly more likely to receive credit in cash than are permanent income households. An interesting difference is that total outlays by private information households are less than total receipts, while total outlays for permanent income households exceed receipts. This pattern is consistent with the idea that private information household's credit transactions may be interlinked with, say, labor transactions, while permanent income households are paying interest on loans received. The difference is really too slight to make too much of, however. Overall, there appear to be no overwhelming differences in either the number or magnitude of credit transactions across regimes.

Some interesting differences *do* emerge when we look at all financial transactions (Table VII). These are all recorded transactions that don't occur in spot markets, and chiefly comprise credit and transfers.

Private information households appear to take in considerably more receipts, both in number and magnitude. On the other side of the ledger, the private information households have a larger number of outlays than the permanent income households, but of smaller magnitude. Total outlays per household average 26,320 Rs. for the private information households, and slightly more for the permanent income households; 29,253 Rs. Transactions in Table VII are broken down by the type of good transferred. From this breakdown, we can see that the main source of the difference in total transactions across regimes comes from the receipt of food and clothing. Private information households receive 28 per cent more transactions of food and clothing than do private information households, each with an average value 19 per cent greater than the average food and clothing

shops"), commercial banks, and cooperatives. What we call "moneylenders" captures informal credit transactions in which interest is (perhaps only nominally) charged. The party may be any other household, or the household's landlord (if it is a tenant), tenant (if it is a landlord), or employer or employee. The category "Friends and relatives" is used for nominally interest free transfers from friends or relatives, but in which the principal is supposed to be repaid. The category "other" includes credit transactions with private shops, itinerant merchants, millers, and miscellaneous.

		Receipts		Outlays	
		Private Information	Permanent Income	Private Information	Permanent Income
Cash	Banks	0.8333 (6200.6000)	1.0500 (3265.2381)	1.1667 (2168.5714)	2.2500 (743.1111)
	Moneylenders	11.6667 (486.3214)	11.8000 (780.2076)	8.7500 (585.6190)	10.3500 (787.9662)
	Friends and Relatives	0.1667 (550.0000)	0.2000 (100.0000)	0.4167 (373.0000)	0.2500 (1604.0000)
	Other	8.0000 (430.6615)	5.8500 (364.2410)	9.4167 (350.5487)	7.9500 (332.6164)
	Total	20.6667 (695.7036)	18.9000 (782.3153)	19.7500 (562.5612)	20.8000 (618.8822)
In Kind	Banks	0.0833 (93.0000)	0.3500 (534.6071)	0.0833 (354.4000)	0.4000 (102.1875)
	Moneylenders	14.4167 (76.9038)	19.3000 (52.3469)	9.6667 (441.2441)	12.7000 (303.8982)
	Friends and Relatives	0.3333 (210.0000)	0.2500 (87.0100)	0.1667 (1795.0000)	0.1000 (135.6750)
	Other	24.5833 (92.3241)	17.9000 (94.5103)	5.7500 (351.6536)	3.8500 (80.5474)
	Total	39.4167 (87.6807)	37.8000 (77.0078)	15.6667 (422.3022)	17.0500 (247.7453)
Total		60.0833 (296.8203)	56.7000 (312.1103)	41.0833 (477.3096)	41.4500 (448.3826)

TABLE VI. Credit Transactions by Regime, Partner, and Type. The mean size of each transaction is reported in parentheses; the average number of transactions per household is not parenthetical.

	Receipts		Outlays	
	Private Information	Permanent Income	Private Information	Permanent Income
Cash	28.3333	25.4500	49.6667	44.5500
	(686.7426)	(619.4444)	(293.6785)	(450.7122)
Food and	39.4167	30.8000	12.8333	11.1500
Clothing	(93.3895)	(78.5478)	(193.4708)	(186.4544)
Other	10.5833	8.2500	1.6667	1.7000
Consumption	(333.2295)	(233.0436)	(436.3200)	(1407.7353)
Production	16.9167	18.6000	4.1667	3.8000
Inputs	(292.5325)	(93.2977)	(305.8200)	(310.7434)
Crop Output	0.0000	0.0000	16.0833	13.3500
	(NaN)	(NaN)	(450.8868)	(264.2719)
	95.2500	83.1000	84.4167	74.5500
Total	(331.9073)	(262.8409)	(311.8119)	(392.4909)

TABLE VII. Financial Transactions by Regime and Type. The mean size of each transaction is reported in parentheses; the average number of transactions per household is not parenthetical.

	Receipts		Outlays	
	Private Information	Permanent Income	Private Information	Permanent Income
Cash	21.0000	19.0500	20.0833	20.8000
	(750.5337)	(784.4231)	(575.9627)	(618.8822)
Food and	22.0833	20.7500	1.7500	3.8500
Clothing	(112.9974)	(82.3630)	(274.6857)	(149.4552)
Other	6.5833	4.1000	0.0000	0.0000
Consumption	(17.5038)	(57.8061)	(NaN)	(NaN)
Production	10.9167	13.3000	2.5833	2.5500
Inputs	(78.8385)	(73.4367)	(53.1613)	(69.2255)
Crop Output	0.0000	0.0000	11.7500	10.6500
	(NaN)	(NaN)	(607.2440)	(326.0216)
	60.5833	57.2000	36.1667	37.8500
Total	(317.4545)	(312.3428)	(534.2046)	(451.6990)

TABLE VIII. Credit Transactions by Regime and Type. The mean size of each transaction is reported in parentheses; the average number of transactions per household is not parenthetical.

receipt of permanent income households; on balance then, private information households receive 33 per cent more food and clothing. In contrast, outlays of food and clothing are very similar across regimes. Outlays of food and clothing are much fewer in number, but more than twice as large as receipts.

	Receipts		Outlays	
	Private Information	Permanent Income	Private Information	Permanent Income
Cash	6.0000	6.2500	29.0000	23.3000
	(586.5500)	(130.7392)	(100.7770)	(274.2791)
Food and	17.3333	10.0500	11.0833	7.3000
Clothing	(68.4084)	(70.6707)	(180.6474)	(205.9677)
Other	4.0000	4.1500	1.6667	1.7000
Consumption	(852.8615)	(406.1698)	(436.3200)	(1407.7353)
Production	6.0000	5.4500	1.5833	1.2500
Inputs	(681.3368)	(140.4367)	(718.0526)	(803.4400)
Crop Output	0.0000	0.0000	4.3333	2.7000
	(NaN)	(NaN)	(26.9183)	(20.7037)
	33.3333	25.9000	47.6667	36.2500
Total	(366.1354)	(153.6041)	(144.8700)	(313.0377)

TABLE IX. Transfer Transactions by Regime and Type. The mean size of each transaction is reported in parentheses; the average number of transactions per household is not parenthetical.

The fact that private information households have more and larger receipts of food and clothing is tantalizing, because it is clear that implementing the private information regime requires considerable control over consumption. Because food and clothing constitute the bulk of consumption, transfers of these goods seems to be the most direct instrument of control.

Because our earlier table on credit failed to turn up any large differences across regimes, it would be surprising if a different method of classifying credit transactions did turn up differences. Table VI classifies credit transactions according to the type of good given or received, and is unsurprising.

Between the two of them, credit transactions and transfer transactions account for nearly all financial transactions. Because there are no blatant differences in credit transactions across regimes, the differences in total financial transactions observed in Table VII must be accounted for principally by differences in transfers. Table IX displays these differences.

6.2. **Who is the Principal?** In Section 4, we assumed away the principal in our principal-agent model by supposing that the principal's marginal utility of consumption was constant (this produced the 'strict' restrictions of our permanent income and private information regimes). While we would not want to defend this assumption on the grounds of its realism, two other defenses are possible: either violations of the assumption are unimportant, or the principal in our models is reasonably regarded simply as a technical device to make the computation of optima straightforward.

In order to test the importance of omitting the principal, consider adding measures of the principal's marginal utility growth to the set of instruments employed in Table III. If we have misspecified the model by omitting the principal, then adding this instrument should result in significant increases in the J statistics presented in that table. The change in the J statistic can be interpreted

Instrument Vector	Village			
	Aurepalle b $(\sigma(b))$ J	Shirapur b $(\sigma(b))$ J	Kanzara b $(\sigma(b))$ J	All b $(\sigma(b))$ J
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, Avg. Consumption at $t+1$	-0.8844* (0.2775) *0.1040*	1.2969* (0.3944) *0.0309*	0.5781* (0.2185) *0.1019*	0.9750* (0.1696) *0.4926*
Family Size, Income, Land, Avg. Consumption growth	-0.7500* (0.1910) *0.2891*	1.3750* (0.3638) *0.0703*	0.6094* (0.2006) *0.1358*	0.9906* (0.1764) *0.3744*
Family Size, Income, Land, Residual Growth	-0.6969* (0.1506) *0.5704*	1.6687* (0.2723) *1.0705*	0.6063* (0.1654) *0.1836*	0.9656* (0.1646) *0.5446*
Family Size, Income, Land, Reciprocal of Residual Growth	-0.8969* (0.2230) *0.1594*	1.2969* (0.3036) *0.1921*	0.6531* (0.1420) *0.3008*	1.0344* (0.1942) *0.2699*

TABLE X. Estimates of b_0 using various proxies for the principal's consumption as instruments. Numbers in parentheses are estimated standard errors; statistics in the column labelled J are scaled so as to be asymptotically distributed χ^2.

as a likelihood ratio test of the restriction that the additional instrument is independent of agent's forecast errors (Newey and West 1987).

The VLS data does not, of course, identify any member of the sample as the principal whose consumption plays a role in the restriction imposed by the private information regime, and unfortunately, examination of the transactions of agents in our sample strongly suggest that the principal is not in our sample,[15] so that $\{c_{0t}\}$ is unobserved. However, it seems clear that the average consumption of the sample ought to be correlated with the unobserved consumption of the principal, if she exists, so we can employ this as an instrument instead of the principal's marginal utility growth. This seems a particularly appropriate under the permanent income regime, where changes in mean consumption would be related to changes in the interest rate. An alternative instrument for the principal's consumption is the difference between output and consumption in our sample. This measure seems particularly appropriate if we take the devices of the literature on private information seriously, as it puts the principal in the role of the residual claimant.

In Table X, we show results from estimating b_0 using the same restriction, equation (24), as we used earlier in Section 4, but adding to our basic set of idiosyncratic instruments a variety of proxies for the principal's marginal utility growth. These include simply the level of average household consumption at $t+1$ (which might be viewed as a direct test of the importance of aggregate

---

[15] The ICRISAT researchers, after drawing their stratified random sample of households, were reportedly pleased that no moneylenders had been selected.

uncertainty in determining consumption, since this amounts to a test of whether next period's aggregate consumption is in this period's information set), the growth of average consumption, growth in the aforementioned residual between output and consumption, and the reciprocal of this variable. Each of these is added separately to a set of instruments which also includes a village dummy, household size, household income per adult equivalent, and household operated land holdings per adult equivalent.

One thing to note about Table X is the fact that estimated parameters within each village display reasonable stability across instrument sets; there is much more variation across villages than within villages. Second, though all but one of the estimates is significant at the 95 per cent confidence level, none of the estimates which employ proxies for principal's consumption is significantly different from the parameter estimates employing only the base instrument set. The final observation we'd like to make has to do with the likelihood ratio tests of our specification described above. This test is constructed by taking the difference between the J statistics from instrument sets two through five minus the corresponding statistic from the base instrument set. The resulting statistic is asymptotically distributed χ_1^2. Thus, for example, the largest such statistic is found in Shirapur, for the instrument set which includes residual growth, and is equal to 1.0705-0.1532, or 0.9173. The critical value of this statistic for a confidence level of 95 per cent is roughly 3.841, so we are quite some ways from rejecting the null hypothesis for any of the proxies for principal's consumption.

Our use of the 'strict' versions of the private information and permanent income regimes is certainly unrealistic, since we can be reasonably confident that in fact interest rates change and aggregate consumption is somewhat unpredictable. Our results indicate that these complications are unimportant, in the sense that incorporating these factors wouldn't yield dramatic changes in our results for these particular data.

An interesting question of interpretation remains. Methods of decentralizing either the permanent income regime or the full insurance regime are well known, and interpretation of this decentralization is straightforward; the reader will readily agree that the principal in these cases is probably best thought of as nothing more than a useful fiction. This point is much less clear in the private information regime, because our understanding of how to decentralize such economies is still in its infancy. Still, something is known. Prescott and Townsend (1984) demonstrate how to decentralize private information economies. Although their methods apply most naturally to static economies, it seems clear from their work that private economies can in fact be decentralized quite generally. More recently, and specifically working with dynamic problems, Atkeson and Lucas (1992) are able to give an equilibrium interpretation to the optima of models similar to the present one. They are pessimistic about the usefulness of such a decentralization in the context of the US economy, because the commitment required to support such equilibria seems to them to be lacking,

but it seems clear that commitment to long-term contracts in the villages we study is not only possible, but common. Finally, Phelan (1995) provides an explicit decentralization which would apply to the present model. Because his work is also intended to be applied to the US economy, how best to interpret the equilibria of his model in the context of the village economies we study awaits further research.

When we think of previous attempts to model village institutions, economists have often chosen to use something like a principal-agent model in order to do so. Whether the model is designed to capture existing tenurial arrangements (Cheung 1969; Stiglitz 1974; Allen 1985), interlinkages of contracts (Bardhan 1980; Braverman and Stiglitz 1982; Bell 1988), or to analyze technology adoption (Bhaduri 1973; Basu 1989), principal-agent models may be said to have been invented with village economies in mind, and have proven to be immensely productive tools. Thus it is with some reluctance that we would consign the principal of our private information model to the dustbin of useful fictions. How might we salvage the principal? Even though we seem not to need her for the present empirical project, she might well play a useful role in a different setting. Accordingly, we will do our best to rehabilitate the principal using the data we've got, if only because to do so seems an interesting extension of the methods we've developed to this point.

One observes at least three sorts of relationships in these villages that seem to map nicely into the principal-agent framework described above. One of these is the relationship between landlord and tenant. Such sharecropping arrangements are, in fact, perhaps the prototypical principal-agent relationship, and a long and distinguished literature has attempted to explain the prevalence of sharecropping arrangements in developing countries.[16] Over half of the land holders in the three villages have been involved in some such arrangement at some point in time. Interestingly, at least in Aurepalle, very long term sharecropping arrangements used to be common; since then tenancy legislation has made this an outmoded contractual form. Sharecropping contracts now seldom last more than a single year, at least nominally; the details of the contract vary substantially (Jodha 1981).

A second sort of relationship is that between moneylender and borrower. Such contracts tend to involve numerous contingencies, are exclusive (borrowers cannot have outstanding loans from more than one moneylender), and are typically of very long duration; in Aurepalle, the mean contract length was nine and one half years (Binswanger et al. 1985). There are around eight moneylenders in Aurepalle, to whom about 60 per cent of the sample households are indebted. Most loans are extended and repaid in kind. No collateral is required, and the timing of repayment is variable.

---

[16] Stiglitz (1974) noted that sharecropping contracts help solve both incentive problems and risk-sharing; Allen (1985) shows that fixed arrangements can solve an enforcement problem as well.

Each moneylender lends only to a relatively small group of borrowers, and will act as their patron in any legal proceedings within the village.

The third type of relationship is between a farmer and what are termed regular farm servants (RFS's). An RFS is hired for a fixed period of time (one year in Aurepalle), under a verbal contract amazing for its specificity. Like tenancy, this is a contractual arrangement which seems to be of diminishing importance; an increasing proportion of labor is hired through daily markets in all three villages (Binswanger et al. 1984). Still, in the VLS sample, 11 households in Aurepalle and two households in each of Shirapur and Kanzara had at least one member who worked as an RFS or "attached servant" during the survey period; six households in Aurepalle and two households in Shirapur had at least one member who worked in this capacity over all ten years of the survey.

With this description, we have no shortage of prototypes for our putative principal. Indeed, we could do with a bit more of a shortage, for the roles we now assign our principal seem too many for a single actor to competently handle. There are several moneylenders in Aurepalle; there are many landlords in all three villages, and there are multiple households which employ regular farm servants in Aurepalle and Shirapur. Yet our model has allowed for only a single principal.

We will not dwell for long on the resolution of this many-principal difficulty; a lack of data makes speculation dangerous. We will note, however, that if the several principals are free to contract however they please among themselves, then their marginal utilities of consumption will grow at precisely the same rate; that is, the principals will be able to insure away all principal-specific risk. Let c_{0t} denote the total consumption of all principals at t. If principals' preferences are of the CES variety, then the consumption of *each* principal k at t will be a fixed proportion of this pooled consumption, say $\theta_k c_{0t}$. Now recall the general private information restriction, given by equation (19), and repeated here for convenience, but substituting in the sharing rule for principal k:

$$\mathrm{E}_t\left[\frac{U^{0\prime}(\theta_k c_{0,t+1})}{U^{0\prime}(\theta_k c_{0t})}\frac{U^{i\prime}(c_{it})}{U^{i\prime}(c_{i,t+1})}\right] = 1.$$

Assuming that principals and agents have identical CES preferences, this becomes

$$\mathrm{E}_t\left[\left(\frac{c_{i,t+1}/c_{it}}{c_{0t+1}/c_{0t}}\right)^\gamma\right] = 1.$$

Note that the share, θ_k, cancels out of this expression. The conclusion is that if the village is divided into a set of principals who have no private information, and a set of agents who have, then we can construct a single 'representative principal.'

Although this resolves the problem of multiple principals, it does nothing to overcome the problem that we don't observe the consumption of the principal. In other applications, this may not be a problem; here we will take seriously our interpretation of the principal as the residual claimant,

and approximate consumption by

$$c_{0t} = \sum_{i=1}^{n} (x_{it} - c_{it})$$

where x_{it} is agent i's output at t, and c_{it} is his consumption. Note that this amounts to assuming that the principal has no means of smoothing her consumption beyond scheduling rewards for the agents. If no smoothing by the principal represents one polar extreme, then our earlier 'strict' assumption of no variation in the principal's marginal utility represents the other. While the truth doubtless lies in some more moderate clime, it seems reasonable to suppose that tests at alternative extreme poles might give some indication of the sensitivity of our tests to the specification of the principal's consumption.

Since we observe only a sample of the village, it is clear that our estimate of the principal's consumption must suffer from measurement error of two sorts. First is just the sampling error we'd expect when estimating any variable which depends upon the entire population distributions of consumption and income. Though a potentially important problem, there is little we can do about it given the rather small size of the sample. It is to be hoped that this problem is of second order importance. The second source of measurement error also has to do with the fact that we only observe a portion of the population. If the *average* residual per agent is positive, then our estimate of c_{0t} will be increasing in the size of the sample. Thus, still assuming a positive average residual, we would expect to underestimate the principal's consumption. Fortunately under the parameterization we've adopted, we're ultimately concerned not with levels of consumption, but with rates of growth, which should be relatively immune to this problem of scale.

In addition to creating a testable restriction for the private information regime, we need to create a similar restriction for the permanent income regime. This will involve relaxing our 'strict' assumption that interest rates are constant and equal to the rate of time preference. Recall from above that the permanent income allocation is characterized by

$$(33) \qquad \mathrm{E}_t \left\{ \frac{U'(c_{i,t+1})}{U'(c_{it})} \beta(1 + r_t) \right\} = 1.$$

Equation (33) is identical to the moment restriction tested by Hansen and Singleton (1982). However, in our application we lack reliable data on the rate of return. This is due in part to simple deficiencies of the data, and due in part to the complicated nature of the contracts we observe in the villages under consideration. As suggested above, it is not at all clear that anything very comparable to simple loans exists in these environments, an observation that one might argue is due to the possibility that allocations are indeed information constrained. Since we cannot observe the price of a risk free one period bond, perhaps even in principle, we must find some other way to

pin down the intertemporal rate of marginal transformation. We needn't look far; our model tells us that the principal described above will set her intertemporal marginal rate of substitution equal to this quantity, even if allocations within the village are otherwise characterized by the private information regime.

Let

$$s_{0t} = \sum_{i=1}^{N} (x_{it} - c_{it})$$

denote the net aggregate savings (or borrowings, if negative) of the agents in our sample. We imagine the principal now as an intermediary on the other side of these transactions. Let b_{0t} denote any net savings by the principal outside the village, so that

$$c_{0t} = s_{0t} - b_{0,t+1} + (1 + r_t)b_{0t}$$

denotes the consumption of the principal, who stands on the other side of the transactions engaged in by agents $1 \ldots N$. This is one way of imposing market clearing within the village. Then, so long as preferences are Gorman aggregable,

$$(34) \qquad U'(c_{0t}) = \beta(1 + r_t)U'(c_{0,t+1}).$$

Note that $c_{0,t+1}$ is assumed to be known at t; this assumption may not innocuous, but can be tested (and will be below). Now, if $b_{0t} = 0$ for all t, then the village economy is closed, and we can use (34) to solve for r_t as a function of the measured endowment. If, on the other hand, $b_{0t} \neq 0$ for some t, then the village economy is open, and it seems reasonable to suppose that the interest rates are taken as given by the agents within the village. In either case, we can use (34) to solve for $\beta(1 + r_t)$, and substitute the result into (33), yielding

$$(35) \qquad \mathrm{E}_t \left\{ \frac{U'(c_{i,t+1})}{U'(c_{it})} \frac{U'(c_{0t})}{U'(c_{0,t+1})} - 1 \right\} = 0.$$

Thus, as before, the private information and permanent income hypotheses may be nested by

$$(36) \qquad \mathrm{E}_t \left[\left(\frac{c_{i,t+1}/c_{it}}{c_{0t+1}/c_{0t}} \right)^{b_0} - 1 \right] = 0.$$

Once again, if we infer that b_0 is positive, then we will prefer the hypothesis that the village belongs to the private infomation regime; if negative, to the permanent income regime. If our ability to draw inferences about the sign of b_0 is very poor, then we may be unable to reject the full insurance hypothesis.

Estimates using equation (36) are presented in Table XI. Qualitatively, estimates of b_0 in Shirapur and Kanzara (and model selection) are similar to estimates made under the strict version of these

Instrument Vector	Village			
	Aurepalle b $(\sigma(b))$ J	Shirapur b $(\sigma(b))$ J	Kanzara b $(\sigma(b))$ J	All b $(\sigma(b))$ J
—	0.4960 (0.2612) 0	0.8348* (0.1737) 0	1.4398* (0.2668) 0	0.6229* (0.0969) 0.4599
Income	0.5761 (0.3030) 0.0161	0.8714* (0.1541) 0.4886	1.5160* (0.2846) 0.0440	1.3567* (0.1338) 8.9090
Income, Land	0.4191 (0.3427) 0.0113	0.9151* (0.2016) 0.0989	1.3757* (0.2947) 0.2067	2.5033* (0.2140) 0.8743
Family Size, Income, Land	0.3260 (0.1777) 5.0265	0.7852* (0.1884) 0.1349	1.2968* (0.3207) 0.0916	1.3841* (0.1123) 11.1001
Family Size, Income, Land, Rain, Avg. Consumption	0.4691* (0.0751) 0.5453	0.8967* (0.1107) 1.7583	1.2571* (0.1335) 1.4543	0.8360* (0.0883) 1.9911

TABLE XI. Estimates of relative risk aversion under private information regime with a risk averse principal.

regimes, and reported in Table III. However, the case is very different in Aurepalle, where the magnitude of the estimates is similar to the magnitudes estimated in Table XI, but the signs are different, implying that if the principal is risk averse and has a highly variable consumption stream, then in Aurepalle consumption is best explained by the private information model. In part this switch in sign may be uninteresting; our model selection in Aurepalle was suspect earlier since most of our estimates of γ in Table III were not significantly different from zero, and this lack of significance tends to occur in Table III, as well. We would also do well to recall the estimation of heterogeneous regimes in Section 6.1. Nonetheless, the fact that different assumptions regarding the behavior of unobserved aggregates can lead to such starkly different conclusions in Aurepalle is a cautionary tale.

6.3. **Aggregate shocks.** The estimation to this point has relied on the assumption that aggregate shocks are small relative to the shocks that affect individual agents' consumption. It is certainly not the case that aggregate shocks matter not at all in the villages we've examined. This section first examines the data to see how important aggregate shocks are. We first present evidence demonstrating that they are large. Fortunately, despite this fact, evidence from a calibrated Monte Carlo experiment assures us that our estimation procedure remains consistent (and nearly unbiased) and that the estimated standard errors of our point estimates are not too badly estimated. Next, we propose a more general way to discriminate between the opposing hypotheses of the permanent income regime and the private information regime which permits any degree of correlation between individual and aggregate outcomes.

Importance of aggregate shocks. A casual look at the disturbance terms[17] suggests that they are approximately distributed log-normally, with a threshold parameter of -1 (this follows from moment condition (equation (36)) and the positivity of consumption). Following this logic, assume that the logarithm of the ratio of marginal rates of substitution, under whatever regime, can be additively decomposed into an aggregate and an idiosyncratic component, or

$$\log(\xi_{it}^{\gamma}) = \log(h(\xi_{it}, \gamma) + 1) = \eta_t + \epsilon_{it}.$$

Each component is normally distributed, with variances σ_{η}^2, σ_{ϵ}^2 and means μ_{η}, μ_{ϵ}, respectively. The mean of the disturbance term $h(\xi_{it}, \gamma)$ is given by

$$\mathrm{E}h(\xi_{it}, \gamma) = \exp\left(\mu_{\eta} + \mu_{\epsilon} + \frac{\sigma_{\eta}^2 + \sigma_{\epsilon}^2}{2}\right) - 1.$$

Accordingly, since $\mathrm{E}h(\xi_{it}, \gamma) = 0$, our distributional assumptions imply that

$$\mu_{\eta} + \mu_{\epsilon} = -\frac{\sigma_{\eta}^2 + \sigma_{\epsilon}^2}{2}.$$

Just how much aggregate variability is there? If we let $\xi_{it} = \frac{c_{0t} c_{i,t+1}}{c_{0,t+1} c_{it}}$ as above, and perform the decomposition indicated above at our preferred estimate of γ, 0.836, then the proportion of the variance of the disturbance term attributable to the aggregate component is 0.27 in Aurepalle, 0.32 in Shirapur, and 0.48 in Kanzara. These proportions seem large enough to cause concern regarding the estimates of standard errors presented in Table III.

We might also ask what proportion of total variance is explained by the aggregate component when we include the principal's consumption in the moment condition, as in equation (36). Evidently this should increase cross-sectional correlation, since this amounts to explicitly adding a common component to the disturbance terms. This turns out to be correct, with the proportion of variance due to the aggregate component increasing to 0.66 in Aurepalle, 0.84 in Shirapur, and 0.77 in Kanzara.

In order to determine the effect of the cross-sectional correlation of forecast errors indicated by the large proportion of variance attributable to the aggregate component we can perform a Monte Carlo exercise. We first select a value for the variance of the aggregate component, and then generate 1000 artificial samples of the same size as the actual sample from the log-normal distribution, matching the mean and the variance of the idiosyncratic components to those of the actual sample. By computing the standard error of the estimates thus derived, and repeating this exercise for a sequence of values of aggregate variances, we can trace out the effect of cross-sectional correlation on the standard errors estimates of relative risk aversion. Figure 2 provides a graph of

---

[17] Disturbances were computed from equation (equation (36)), evaluated at $\gamma = 0.8360$.

Instrument Vector	Village		
	Aurepalle b $(\sigma(b))$	Shirapur b $(\sigma(b))$	Kanzara b $(\sigma(b))$
Strict Estimates	-0.5554* (0.1961)	1.5267* (0.5132)	0.6575* (0.3131)
Principal Estimates	0.4443 (0.2496)	0.7899 (0.4975)	1.1445* (0.5255)

TABLE XII. Estimates of risk aversion under the private information regime. 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.

these effects, with the standard error of the estimate on the y axis of the graph, and the variance of the aggregate component on the x axis. It can be seen from this picture that, while a small amount of cross-sectional correlation may have dramatic effects on the magnitude of the true standard error of estimates of risk aversion, beyond a certain point increases in the magnitude of the variance of the aggregate component matter little. The points on the graph which correspond to the variance decomposition for each village are indicated with small circles. Note that there are two points for each village; the lower grouping (marked with asterisks) for the 'strict' hypotheses, and the higher grouping (marked with circles) for the case with a risk-averse principal.

Not much is known about the small sample properties of GMM estimators, and less is known about small sample properties in the presence of cross-sectional correlation. Given the degree of cross-sectional correlation present in these data, it is not even obvious that our estimator has desirable asymptotic properties.[18] Our Monte Carlo experiment is quite reassuring on this point. While we have not investigated the asymptotic properties of our estimator, there appears to be little or no bias, even in the small samples of our experiment, and regardless of the degree of cross-sectional correlation. The average estimate of b_0 for our 1000 artificial samples is invariably quite close to the 'true' value of 0.836. There is some evidence of an upward bias as aggregate variance reaches quite large values, with the mean of b equal to 0.8367 when $\sigma_\eta = 0$, and increasing to 0.9676 for $\sigma_\eta = 0.2$. Correcting for this bias and also correcting our estimated standard errors yields the estimates presented in Table XII, corresponding to the final row of uncorrected tables III and XI.

The first row of Table XII presents corrected 'strict' estimates, using equation (24) for a moment restriction. The second row presents corresponding estimates for a non-fictive principal, and

---

[18]Our estimator will be consistent and asymptotically normal as $N \to \infty$ if cross-sectional correlation is sufficiently small, or as $T \to \infty$ if cross-sectional correlation is large. Since $T = 8$ in our dataset, this provides an important reason to prefer the 'strict' specification of our test, in which cross-sectional correlation is relatively small.

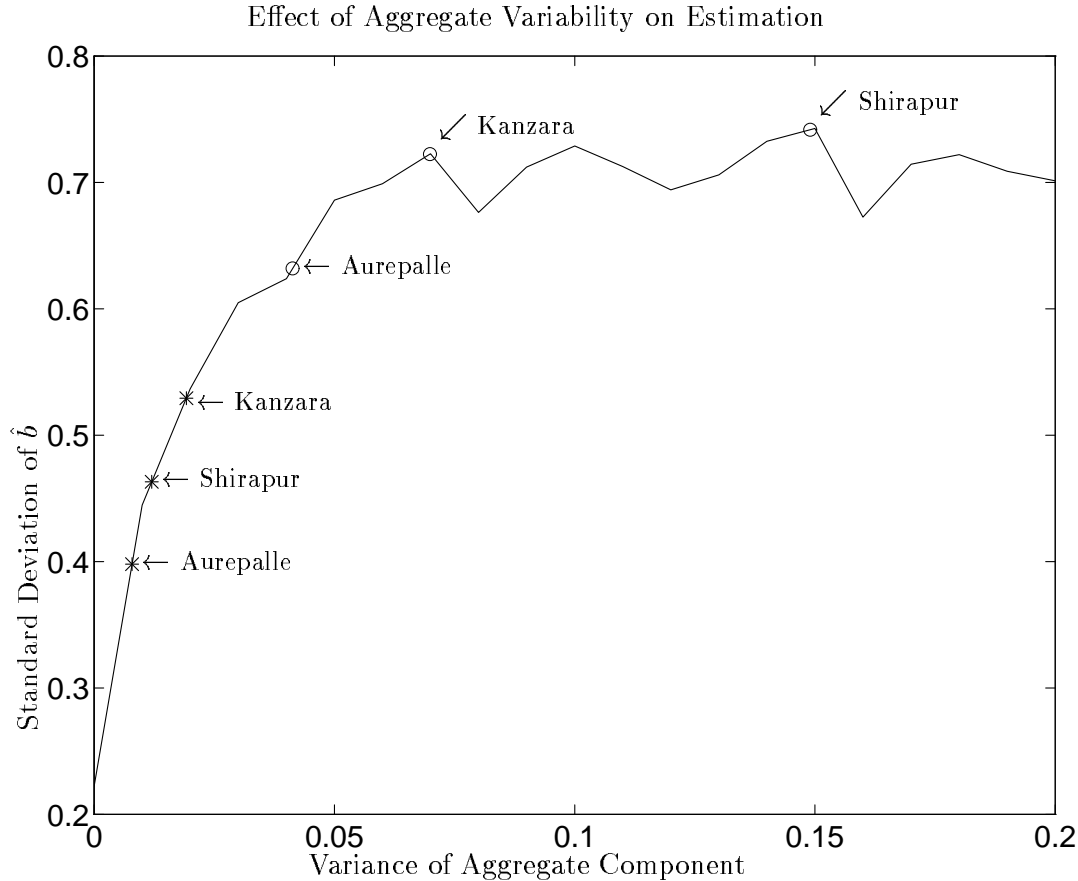Effect of Aggregate Variability on Estimation

FIGURE 2. Monte Carlo investigation of the effect of aggregate uncertainty on estimation

employs equation (36). After correcting for small sample bias and underestimated standard errors, all of the strict estimates remain significant. The 'principal' estimates fare less well, with only the point estimate for Kanzara remaining significant.

These results suggest that we treat the standard errors reported in Tables XI and III with caution; it may make sense to think of the standard errors there as lower bounds on the the true standard errors. The reported J statistics will, as a consequence, probably be lower than they ought to be. Fortunately, none of our results hinges on the values of these standard errors or J statistics; what is important for our purposes is the sign of the estimated relative risk aversion. The fact that true standard errors may be underestimated reduces our confidence that we have correctly selected the regime, but does not change which regime we think more likely.

6.4. **Distinguishing regimes under aggregate uncertainty.** Although the discovery that there is an important aggregate component to the forecast errors in our sample has important effects on estimation, it has a more profound effect on our formulation of the moment condition associated

with the permanent income regime. In particular, in deriving equation (21) we relied on the assumption that there is no aggregate uncertainty. Our results above seem to indicate that this is not a reasonable assumption. Here we formulate a test which does not rely on this assumption.

We know from equation (19) that under the private information regime,

$$(37) \qquad \qquad \mathrm{E}_t \left\{ \frac{U'(c_{0,t+1})}{U'(c_{i,t+1})} \right\} = \frac{U'(c_{0t})}{U'(c_{it})}.$$

Similarly, from equation (33) it is clear that

$$(38) \qquad \qquad \frac{\mathrm{E}_t U'(c_{0,t+1})}{\mathrm{E}_t U'(c_{i,t+1})} = \frac{U'(c_{0t})}{U'(c_{it})}.$$

The sole difference between these two expressions is the arrangement of the expectations operator(s): the difference between the two is the result of the partial insurance provided to both principal and agent under the private information regime. Section 4 explained a technique for characterizing (37), and one could use the techniques of, e.g., Gallant, Hansen, and Tauchen (1990) or Chamberlain (1991) in order to compute the left hand side of (38). Alternatively, under some more or less restrictive distributional assumptions, one could parameterize the expectation operators in equations (37) and (38) and produce another nested specification. This is left as a topic for future research.

6.5. **Measurement Error.** The VLS dataset is unusual in that it provides consumption data for a panel over a period of ten years. However, Gautam (1991) has done a careful study of the consumption data, and alleges that there may be problems with measurement error.

The VLS consumption data is constructed from expenditure data on a variety of foodstuffs and sundries, and from data on initial levels of stocks of various goods. There are two main concerns with this procedure. The first is that, in the first year of the survey, inaccurate data on initial stocks may have led to poorly estimated final consumptions. We have obviated this concern in this study by simply dropping the first year of data.

The second concern is that households may sometimes under-report consumption in a particular expenditure category, perhaps because of poor recollection or perhaps because the interviewer may simply neglect some category altogether. The source of this possible error that most concerns Gautam is that households may fail to report consumption out of their own production, possibly of several different goods. Another reason for concern over missing expenditures is that there appears to be a dramatic drop-off in the number of expenditures reported in several expenditure categories in the final years of the survey. This drop-off appears to be unrelated to reporting of home production. There is no apparent drop-off in the number of households purchasing any of

the cereals (rice, wheat, bajra, jowar) or in pulses, milk, or clothing; there is an apparent drop-off in recorded households purchasing oil, sugar, vegetables and meat in the last three years of the survey.

Fortunately, from Gautam's work we can tell a good deal about the structure of this potential measurement error. By taking the household's total reported income and subtracting all reported transactions *not* including consumption goods, he constructs an upper bound on the measurement error due to underreporting of expenditures.

Let c_{it}^* denote the true level of consumption, and let us take Gautam's upper bound as the true measurement error, and express it as a share of total consumption. Hence observed consumption is related to true consumption by

$$c_{it} = c_{it}^*(1 - \epsilon_{it})$$

where ϵ_{it} is the share of household i's time t consumption which is not reported. Households with zero reported consumption in any given year would have been dropped from the sample,[19] so let $[0, 1)$ be the support of ϵ_{it}.

If it were the case that zero expenditures in some categories were due to interviewer error, we might suppose that the interviewer made this mistake for some households in some years, and not for others. This suggests that there may be many household-years for which there is no measurement error. We wish to capture not only underreporting of consumption from own production, but also the drop-offs in the final years of the survey. Accordingly, let the proportion of household-years for which there *is* measurement error be ρ; hence the distribution of ϵ_{it} has a mass of $1 - \rho$ at zero. Assume that, for households and periods for which consumption is measured with error, the conditional density of the error is log-normal and truncated at one.

This gives us a mixed distribution for ϵ_{it} which is completely characterized by three parameters (ρ, μ, σ^2) where $\mu = \text{E} \log \epsilon_{it}$ and $\sigma^2 = \text{Var} \log \epsilon_{it}$. From Gautam's Table 1, we know that the mean of the measurement error is 0.2414 and that the variance is 0.3884 (assuming equal numbers of households in each village, which is almost correct). Denote these statistics by μ_ϵ and σ_ϵ^2, respectively. Then we can relate the parameters of the error distribution to Gautam's estimates by

$$(39) \qquad \rho e^{\mu + \sigma^2/2} = \mu_\epsilon$$

$$(40) \qquad (1 - \rho)\left(\rho^2 e^{2\mu + \sigma^2}\right) = \sigma_\epsilon^2.$$

Fix $\rho \leq 1$. There is a unique (μ, σ^2) satisfying these equations (and the restriction that σ^2 be non-negative) for values of ρ greater than about 0.75 (no solution exists for values of ρ less than this). Index these solutions by ρ, and denote them by $(\mu(\rho), \sigma^2(\rho))$.

---

[19] Such households would not have well-defined marginal rates of substitution given iso-elastic preferences.

Now consider the following Monte Carlo exercise. Draw a sequence of consumptions c_{it}^* for N households over T periods in such a way that the consumptions satisfy equation (23) when the coefficient of relative risk aversion, γ, is equal to our earlier estimate of 0.836. Assume that consumption realizations are independent across agents, and draw the sample so that its variance matches the empirical variance (see Section 6.3, *infra*). Now fix $\rho \in [0.75, 1]$, compute $(\mu(\rho), \sigma^2(\rho))$, and generate some random measurement errors from the distribution described above. Calculate the observed consumption from the c_{it}^* and the measurement errors, and use the resulting observations to estimate γ under the hypothesis that no measurement error is present.

Figure 3 gives estimates of γ when the 'true' value used to generate the consumption data is 0.836, but under the incorrect hypothesis that there is no measurement error. The mean and variance of the measurement error match the estimates given by Gautam, but the incidence of measurement error, given by ρ, is allowed to vary. Note that the bias is greater the greater is ρ, and that the bias is toward zero.

The line denoting estimates of γ is accompanied by a 95 per cent confidence interval, which shrinks slightly as ρ goes to one. Estimated values of γ remain significant and positive, however, even at the extreme of $\rho = 1$.

We've been treating 0.836 as the true value of γ, but of course it's merely an estimate. Figure 4 suggests that, in the presence of the sort of measurement error postulated above, it may be a very poor estimate. This figure traces out estimates of γ for a rather large range of true values of γ when $\rho = 0.75$ and the mean and variance of the postulated measurement error match Gautam's estimates. Estimates appear to be nearly unbiased for γ less than about 0.5 in absolute value, but to be increasingly biased for values of γ greater than this. In fact, it appears that estimating $0.7 < b < 5.0$ provides little information about the true γ at all.[20]

However, even if we're quite uncertain about the true value of γ, we can remain confident about our selection of regime; Figure 4 seems to indicate that there is little danger of mis-estimating the sign of gamma. Furthermore, because the bias (for values of γ for which there is one) is toward zero, we might actually feel more confident of our selection of regime that we would otherwise be.

7. Conclusion

We've developed a set of techniques 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 techniques in three villages in India are not fully satisfactory,

---

[20] The reader who suspects (after looking at Figure 4) that $\hat{\gamma}$ converges to zero for large (absolute) values of γ is correct. If it were the case that the true value of γ in an economy were extremely large, then estimated values of relative risk aversion might be indistinguishable from zero. This would be a poor estimate of the true γ, of course, and would also preclude distinguishing among regimes.

FIGURE 3. Bias versus ρ

but indicate nonetheless that in these environments, the private information, constrained-efficient allocation rule seems to provide the best explanation of consumption in two of these three villages. In the third village, allocations are broadly consistent with the permanent income hypothesis, 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 class of private information models used here possess some unusual features, some of which are appealing, while others are more troublesome. 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.

FIGURE 4. Bias versus γ

This paper has been fairly agnostic on the important issue of what contractual devices might be used to support private information allocations. We know that a contract specifying utility at every state and date will work, but it is not clear that we can observe such contracts. More reasonably, we might ask whether combinations of simpler contracts might suffice to deliver private information allocations. Besley (1988) shows that interlinking contracts specifying goods and credit offer potential welfare gains; Udry (1990) suggests that sequence of contingent loans he observes in Nigeria suffice to support the full insurance optimum. In a particular environment, precise specification of the primitives of the economy allow us to predict the characteristics of efficient institutions beyond the general characterization given in this paper; this is left for future research.[21]

---

[21] Ligon (1993) makes a preliminary attempt to provide a more direct characterization of efficient institutions in the ICRISAT villages. Fafchamps (1992) argues that many of the stylized institutional features of village economies might be predicted by some repeated game with limited commitment.

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've presented 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.

University of California, Berkeley

REFERENCES

Allen, F. (1985). On the fixed nature of sharecropping contracts. *The Economic Journal 95*, 30–48.

Altonji, J., F. Hayashi, and L. Kotlikoff (1989). Is the extended family altruistically linked? NBER Working Paper No. 3046.

Altug, S. and R. Miller (1990). Household choices in equilibrium. *Econometrica 58*, 543–570.

Atkeson, A. (1991). International lending with moral hazard and risk of repudiation. *Econometrica 59*, 1069–1089.

Atkeson, A. and R. E. Lucas, Jr. (1992). On efficient distribution with private information. *The Review of Economic Studies 59*.

Bardhan, P. K. (1980). Interlocking factor markets and agrarian development: A review of the issues. *Oxford Economic Papers 32*, 82–98.

Basu, K. (1989). Technological stagnation, tenurial laws, and adverse selection. *American Economic Review 79*(1), 251–255.

Bell, C. (1988). Credit markets and interlinked transactions. In H. Chenery and T. Srinivasan (Eds.), *Handbook of Development Economics, Volume I*. Elsevier Science Publishers.

Besley, T. J. (1988). Tied-in credit with a monopoly credit market. *Economics Letters 28*, 105–108.

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory 16*, 252–92.

Bhaduri, A. (1973). Agricultural backwardness under semi-feudalism. *Economic Journal 83*, 120–137.

Bhende, M. J. (1986). Credit markets in rural south india. *Economic and Political Weekly 21*(38–39), A119–A124.

Binswanger, H. P. et al. (1984). Common features and contrasts in labor relations inthe semiarid tropics of India. In H. P. Binswanger and M. R. Rosenzweig (Eds.), *Contractual Arrangements, Employment, and Wages in Rural Labor Markets in Asia*. New Haven: Yale University Press.

Binswanger, H. P. et al. (1985). Credit markets in rural South India: Theoretical issues and empirical analysis. World Bank Discussion Paper Report ARU 45, World Bank.

Binswanger, H. P. and M. R. Rosenzweig (1986). Credit markets, welath and endowments in rural south India. Presented at the 8th International Economic Association World Congress.

Braverman, A. and J. Stiglitz (1982). Sharecropping and the interlinking of agrarian markets. *American Economic Review 72*, 695–715.

Browning, M. and C. Meghir (1981). The effects of male and female labor supply on commodity demands. *Econometrica 59*(4), 925–951.

Chamberlain, G. (1991). Sequential moment restrictions in panel data. Manuscript.

Cheung, S. N. S. (1969). *The Theory of Share Tenancy; with special application to asian agriculture and the first phase of Taiwan land reform*. Chicago: University of Chicago Press.

Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy 99*, 957–976.

Deaton, A. (1990). On risk, insurance, and intra-village consumption smoothing. Manuscript.

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.

Eichenbaum, M. S. and L. P. Hansen (1990). Estimating models with intertemporal substitution using aggregate time series data. *Journal of Business and Economic Statistics 8*, 53–69.

Fafchamps, M. (1992). Solidarity networks in preindustrial societies: Rational peasants with a moral economy. *Economic Development and Cultural Change 41*(1), 147–174.

Frees, E. W. (1995). Assessing cross-sectional correlation in panel data. *Journal of Econometrics 69*, 393–414.

Gallant, A. R., L. P. Hansen, and G. Tauchen (1990). Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics 45*, 141–179.

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.

Green, E. J. (1987). Lending and the smoothing of uninsurable income. In E. C. Prescott and N. Wallace (Eds.), *Contractual Arrangements for Intertemporal Trade*. Minneapolis: University of Minnesota Press.

Green, E. J. and S.-N. Oh (1991). Contracts, constraints, and consumption. *Review of Economic Studies 58*, 883–899.

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.

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

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

Hayashi, F. (1987). Tests for liquidity constraints: A survey and some new observations. In T. F. Bewley (Ed.), *Advances in Economic Theory, Fifth World Congress*. Cambridge: Cambridge University Press.

Hayashi, F., J. Altonji, and L. J. Kotlikoff (1996). Risk-sharing between and within families. *Econometrica 64*(2), 261–294.

Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics 10*, 74–91.

Jodha, N. S. (1981). Agricultural tenancy: Fresh evidence from dryland areas in India. *Economic and Political Weekly 16*(52), A118–A128.

Leme, P. (1984). Integration of international capital markets. Manuscript.

Ligon, E. (1993). Notes on simulata. Manuscript.

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.

Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy 99*, 928–956.

Morduch, J. (1991). Consumption smoothing across space. Manuscript.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association 55*.

Newey, W. K. and K. D. West (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review 28*(3), 777–787.

Pakes, A. (1991). Dynamic structural models: Problems and prospects. mixed continuous discrete controls and market interactions. Prepared for the Sixth World Congress of the Econometric Society.

Phelan, C. (1995, August). On the long run implications of repeated moral hazard. Manuscript.

Phelan, C. and R. M. Townsend (1991). Computing multi-period, information-constrained optima. *Review of Economic Studies 58*, 853–882.

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.

Prescott, E. and R. Townsend (1984). General competitive analysis in an economy with private information. *International Economic Review 25*, 1–20.

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

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

Rosenzweig, M. R. and H. P. Binswanger (1993). Wealth, weather risk and the composition and profitability of agricultural investments. *Economic Journal 103*, 56–78.

Runkle, D. E. (1991). Liquidity constraints and the permanent-income hypothesis: Evidence from panel data. *Journal of Monetary Economics 27*, 73–98.

Sharma, K. C. and R. P. Singh (undated). Rainfall and crop production in three villages of India's SAT. Undated manuscript.

Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *Review of Economic Studies 54*, 599–618.

Stiglitz, J. E. (1974). Incentives and risk sharing in sharecropping. *Review of Economic Studies 41*, 219–255.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *The American Economic Review 71*, 393–410.

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*, 1166–1185.

Townsend, R. M. (1987). Microfoundations of macroeconomics. In T. F. Bewley (Ed.), *Advances in Economic Theory, Fifth World Congress*. Cambridge: Cambridge University Press.

Townsend, R. M. (1992). Risk and insurance in village India. Manuscript.

Udry, C. (1990). Credit markets in northern Nigeria: Credit as insurance in a rural economy. *The World Bank Economic Review 4*, 251–269.

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.

Zeldes, S. (1989). Consumption and liquidity constraints. *Journal of Political Economy 97*, 305–346.