

INFORMAL INSURANCE ARRANGEMENTS WITH LIMITED COMMITMENT: THEORY AND EVIDENCE FROM VILLAGE ECONOMIES

ETHAN LIGON, JONATHAN P. THOMAS, AND TIM WORRALL

ABSTRACT. We study efficient insurance arrangements when there is complete information but limited commitment because only limited penalties can be imposed if households renege on their promises. Insurance arrangements must therefore take into account the fact that households will renege if benefits from doing so outweigh costs. Using a general dynamic model with Markovian aggregate and idiosyncratic uncertainty, we show that efficient arrangements are characterised by a simple updating rule, similar to a simple debt contract with occasional forgiveness. We use Indian village data to test the theory against three alternatives: autarky, full insurance, and a static model of limited commitment.

Date: October 9, 2000.

We are very grateful for comments and correspondence with Chris Udry and Stephen Coate at an early stage of the paper, to Andrew Foster and Narayana Kocherlakota for comments on more recent drafts, and to Orazio Attanasio and two thoughtful referees. Thanks are also due to the many useful comments from participants at seminars given at Berkeley, Boston University, Brown, Exeter, Liverpool, the London School of Economics, Manchester, Pennsylvania, Washington, the Federal Reserve Bank of Minneapolis, University College London, the World Bank, and the Econometric Society meetings in Prague and Tokyo. Sections 2 and 3 of the present paper are drawn from an earlier paper titled “Informal insurance arrangements in village economies,” written by the second and third authors (dated November 1994).

1. INTRODUCTION

In his study of risk and insurance in village India, Townsend (1994) tests whether household consumption allocations replicate the Pareto-efficient full risk-pooling outcomes that would result from a complete set of competitive state-contingent markets. He regresses household consumption on aggregate consumption and a vector of other variables including household income.¹ With full risk-pooling only aggregate consumption should enter significantly into the regression. However, he finds that although the null hypothesis of full risk-pooling performs reasonably well, household income is significant in explaining household consumption. The full risk-pooling hypothesis predicts more insurance than is actually observed in the data. Similar conclusions in the context of rural communities in less developed countries have been found by Deaton (1992) for Côte d'Ivoire, Ghana and Thailand, by Udry (1994) for northern Nigeria, by Grimard (1997) for the Côte d'Ivoire, by Lund and Fafchamps (1997) for the Philippines, and by Dubois (2000) for Pakistan. All reject the null hypothesis of full risk-pooling.² This failure of full risk-pooling may be due to either problems of limited information or limited commitment or both. In this paper we consider a model of mutual insurance with limited commitment and test it on the Indian village data used by Townsend.³ We find that the limited commitment model performs substantially better than the full risk-pooling model, and can fully explain the dynamic response of consumption to income, the puzzle originally raised by Townsend's paper. However, while the limited commitment model does a good job of explaining the dynamics of consumption, the model has difficulty explaining the distribution of consumption, in that wealthier households tend to consume less than the model predicts.

With limited commitment a mutual insurance scheme is only feasible if the long-term benefits of making a transfer in terms of future insurance exceed the short-term costs, that is, if promises of future reciprocation are perceived to be credible. That a scheme of mutual insurance with limited commitment may be possible in rural societies was suggested by Posner (1980) and established by Kimball (1988). In an important paper Coate and Ravallion (1993) solved for an efficient mutual insurance arrangement for a symmetric two-household model with a restriction to stationary transfers. The restriction to stationary transfers means that whenever the same state occurs, the same transfers are made and that the past history of transfers is unimportant. We call this the static limited commitment model.

The static limited commitment model can be interpreted as a system of gifts or transfers.⁴ Eswaran and Kotwal (1989), however, emphasise that credit can also be used as a form of mutual insurance — borrow when times are bad and repay when times are good. There is ample evidence that loans are used for mutual insurance purposes (see, e.g., Platteau and Abraham (1987), Udry (1994) and Lund and Fafchamps (1997)). These loans are highly informal and may be better termed as 'quasi-credit' because typically there are no written records, no legal procedures to enforce repayments, no collateral and an understanding that

debts may be delayed or forgiven if circumstances so dictate.⁵ While puzzling at first sight, the ‘quasi-credit’ element of these informal insurance arrangements might be desirable when commitment is limited. It offers a future reward to a household which is being asked to sacrifice current consumption in order to insure a less fortunate household and therefore encourages it to transfer more. On the other hand, it creates an incentive problem for households which have to repay loans previously taken out. We show in Section 3 however, that making current transfers depend upon the past history of transfers is beneficial. We call this the dynamic limited commitment model.

In Section 2 we consider a general bilateral model where income follows a finite state Markov process. This allows for the possibility of both aggregate and idiosyncratic risk and serial correlation. Households are assumed to be infinitely lived, and to consume a single, non-storable consumption good. Given the absence of a formal legal framework, insurance arrangements between households are assumed to be sustained by the joint means of direct penalties against breach⁶ and also the threat of future exclusion from insurance possibilities. The constrained-efficient insurance arrangement is completely characterised in Proposition 1 of Section 3. It can be summarised easily and briefly. There is a simple updating rule which can be expressed in terms of the ratio of the marginal utilities of the two households. Each state of nature (a distribution of endowments across households) is associated with a particular interval of the ratio of marginal utilities. Given the current state and the previous period’s marginal utility ratio, the new ratio lies within the interval associated with the current state, such that the change in the ratio is minimised. As households are risk averse, the ratio of marginal utilities of the two households is monotonic in the transfer between households and knowing the current ratio determines the current transfer. The updating rule then determines how transfers change over time. The updating rule is very intuitive as it implies that the ratio of marginal utilities is kept constant whenever possible (a constant ratio would of course be the outcome of full risk-pooling). However, if full risk-pooling is not attainable, then the ratio must change and one of the households will be constrained (that is, be made just indifferent between adhering to the insurance arrangement and renegeing). Proposition 2 shows how these intervals depend on parameter values and we provide two simple illustrative examples of the interval characterisation and the updating rule.

In Section 4 we show how the bilateral model is extended to $H > 2$ households and to allow for the intertemporal transfer of resources. In particular it is shown how the basic properties of the model are preserved and that a modified updating rule still applies. Changes in welfare due to the introduction of an intertemporal technology are ambiguous; welfare in the limited commitment environment may be reduced as it increases the payoffs to autarky, but on the other hand may have a beneficial effect by making it possible for households to ‘post a bond’ (Attanasio and Rios-Rull 2000; Ligon, Thomas, and Worrall 2000). Finally, we show that, with Gorman aggregable preferences, the H household problem may be regarded as a

sequence of problems between each household and an aggregate of the “rest of the village,” a fact which we are able to exploit in our empirical work.

Section 5 discusses the data used and the estimation procedure. We test the dynamic limited commitment model by using it to predict consumption allocations and comparing the predictions with consumption data from three Indian villages. These predictions are obtained by first assuming that preferences exhibit constant relative risk aversion. We then estimate a household specific income process and use a finite cell approximation for each household and the rest of the village. Next, for a given set of parameter values for the discount factor, the coefficient of relative risk aversion and the direct penalty, we use the results of Section 3 to compute the set of optimal insurance arrangements. Using the data on actual first-period consumption to generate an initial ratio of marginal utilities, we use actual incomes and the updating rule to generate predicted consumptions. Finally we use a search procedure over the set of parameter values to minimise the sum of the squared residual errors between actual and predicted consumptions. A similar procedure is adopted to generate predicted consumption from the static limited commitment model.

Results are reported in Section 5.5. There is strong evidence that limited commitment models fare better than the full risk-pooling model in explaining consumption allocations in all three of the villages studied. Further, the dynamic limited commitment model outperforms the static limited commitment model in each of the several tests we perform. The predictions of the dynamic limited commitment model are most highly correlated with actual consumptions in each of the three villages of all the models estimated. In Section 6, to try to identify the strengths and weakness of the dynamic limited commitment model, we regress residuals from the full insurance model on income (replicating results reported by Townsend (1994)), and show that the consumptions predicted by the dynamic limited commitment model help to explain the errors from the full insurance model. A second set of regressions examines errors from the dynamic limited commitment model. It turns out that the dynamic model can fully explain the dynamic response of consumption to income, but has some difficulty simultaneously explaining the distribution of average consumption across households.

There is a burgeoning literature which develops the dynamic limited commitment model.⁷ In the context of mutual insurance both Fafchamps (1999) and Kocherlakota (1996) consider a limited commitment model. Fafchamps (1999) shows how the dynamic limited commitment solution of mutual insurance can be interpreted as quasi-credit with both a pure transfer and loan element where debt repayments may be rescheduled. The main concern of Kocherlakota (1996), who assumes a symmetric non-autocorrelated endowment process, is establishing the existence of a unique invariant long-run distribution of promised utilities. Neither of these papers provides the explicit characterisation as we do here and which we use for computing predicted consumption values from the model.⁸ Our characterisation is obtained

by generalising the results of Thomas and Worrall (1988), which confirmed Holmström's back-loading principle⁹ and provided a simple interval characterisation and updating rule for the contracted wage in an infinite horizon implicit wage contract model with a risk averse worker and a risk neutral employer in which the worker can quit at any date to work at the random (*i.i.d.*) spot market wage and in which the employer can fire the current worker and hire at the spot market wage.

There has been much less work attempting to test the dynamic limited commitment model. Foster and Rosenzweig (1995) extend the model of Section 2 of the present paper to allow for altruistic links. They use the implied negative relationship between the current transfer and an aggregate of the previous transfer to test the dynamic limited commitment model on data from India and Pakistan. They provide evidence that limited commitment substantially constrains informal transfer arrangements and show that altruism also plays an important role in ameliorating sustainability constraints. Lund and Fafchamps (1997), in a unique study of rice farmers in the Philippines which identifies networks of friends and family, also find evidence consistent with models of dynamic limited commitment. In particular they find that insurance is carried out within networks but find little evidence that transfers are motivated by altruism or by collateral constraints. Beaudry and DiNardo (1995) provide an empirical test in the implicit labour contract context based on the observation that when the wage is decoupled from marginal productivity the only effect of wages on hours is through an income effect, so that an increase in the hourly wage should be associated with a fall in hours worked if leisure is a normal good. In contrast to these papers our characterisation of the limited commitment solution affords the possibility for a more structural empirical approach. Thus model predictions can be computed and compared with actual consumptions, so that the theory can be tested against well specified alternatives. To our knowledge this is the first paper to adopt such an approach.

2. THE MODEL

This section outlines the bilateral risk-sharing model. Suppose that there are two households $i = 1, 2$. Each period $t = 1, 2, \dots$, household i receives an income $y_i(s) > 0$ of a single perishable good, where s is the state of nature drawn from a finite set $s \in \mathcal{S}$, and $\mathcal{S} = \{1, 2, \dots, S\}$. It is assumed that the state of nature follows a Markov process with the probability of transition from state s to state r given by π_{sr} , and we assume that $\pi_{sr} > 0$ for all r and s .¹⁰ We assume that there is some initial distribution over period 1 states $r \in \mathcal{S}$ given by π_r^0 . This formalisation includes as a special case an identical and independent distribution over the possible states of nature (π_{sr} is independent of s). The general specification of the dependence of incomes $y_i(s)$ on the state of nature allows for arbitrary correlation between the two incomes. In agrarian communities incomes are likely to be positively correlated, though idiosyncratic shocks will also be important.

Households 1 and 2 have respective per-period von Neumann-Morgenstern utility of consumption functions $u(c^1)$ and $v(c^2)$, where c^i is consumption of household i . It is assumed that $c^i \geq 0$; this lower bound can be interpreted as subsistence consumption by a suitable translation of the origin. Household 2 is assumed to be risk averse, with $v'(c^2) > 0$, $v''(c^2) < 0$ for all $c^2 > 0$, and household 1 is risk averse or risk neutral, $u'(c^1) > 0$, $u''(c^1) \leq 0$ for all $c^1 > 0$. Households are infinitely lived, discount the future with common discount factor δ , and are expected utility maximisers.¹¹ Define $\xi_s \equiv v'(y_2(s))/u'(y_1(s))$ to be the autarkic ratio of marginal utilities in state s where both households consume their own income. We will assume that there are at least two states s and r such that $\xi_s \neq \xi_r$ (otherwise autarky is first-best).

As at least one of the households is risk averse and there is a pair of distinct autarkic marginal utility ratios, the two households will have an incentive to share risk. We assume that the households enter into a (possibly implicit) risk-sharing contract, and while such a contract is not legally enforceable, there are two consequences for a party which reneges upon the contract. First, it loses future insurance possibilities. We assume that after a contract violation by either party, both households consume at autarky levels thereafter. This can be interpreted as a breakdown of ‘trust’ between the households. Alternatively, viewing the contractual agreement as a non-cooperative equilibrium of a repeated game, since reversion to autarky is the most severe subgame-perfect punishment, not only does a contract which can be supported by reversion to autarky correspond to a subgame-perfect equilibrium when there are no direct penalties for breach, but also there can be no other equilibrium outcomes (see Abreu (1988)).¹² Secondly, it is assumed that contract breaches may meet some direct penalty. While there is no explicit legal enforcement of these credit arrangements, such breaches probably lead to some social stigma and other forms of social punishment. For simplicity we shall assume that an expected discounted utility loss of $P_i(s) \geq 0$ is suffered by household i if it reneges in state s .^{13,14}

Let s_t be the state of the world occurring at date t . A contract $\tau(\cdot)$ will specify for every date t and for each history of states up to and including date t , $h_t = (s_1, s_2, \dots, s_t)$, a transfer $\tau(h_t)$ to be made from household 1 to household 2 (a negative transfer signifying a transfer in the opposite direction). Let us define $U_t(h_t)$ to be the expected utility gain over autarky (or *surplus*) of household 1 from the contract from period t onwards, discounted to period t , if history $h_t = (h_{t-1}, s_t)$ occurs up to period t (i.e., when the current state s_t is known):¹⁵

$$(1) \quad U_t(h_t) = u(y_1(s_t) - \tau(h_t)) - u(y_1(s_t)) + \mathbb{E} \sum_{j=t+1}^{\infty} \delta^{j-t} (u(y_1(s_j) - \tau(h_j)) - u(y_1(s_j))).$$

(\mathbb{E} denotes expectation.) We define $V_t(h_t)$ to be the analogous surplus for household 2. The first term in (1), $u(y_1(s_t) - \tau(h_t)) - u(y_1(s_t))$, is the short-run gain from the contract and the second term is the long-run or continuation gain from the contract. Then household 1

will have no incentive to break the contract if the following *sustainability constraint* holds at each date t after every history h_t ,

$$(2) \quad U_t(h_t) \geq -P_1(s_t),$$

and likewise the constraint for household 2 is

$$(3) \quad V_t(h_t) \geq -P_2(s_t).$$

If both (2) and (3) hold, then we call the contract *sustainable*. Within the class of sustainable contracts, we shall characterise the *constrained-efficient* contracts, those which are not Pareto-dominated by any other sustainable contract.

3. CHARACTERISATION OF CONSTRAINED-EFFICIENT CONTRACTS

This section provides a characterisation of the optimum risk-pooling arrangement with dynamic limited commitment. It outlines the dynamic programming problem used to calculate the predicted consumption allocations under dynamic limited commitment and shows how the solution is completely characterised by a simple and intuitive updating rule for a single parameter.

The dynamic programming procedure used to solve for the (constrained) efficient sustainable contract relies on two key facts. First the Markov structure implies that the problem of designing an efficient contract is the same at any date at which the same state of nature occurs. Secondly, an efficient contract must, after any history, have an efficient continuation contract. The reason is simply that all constraints are (at least weakly) relaxed by moving to a Pareto dominating continuation contract that satisfies the sustainability conditions from an inefficient one, and such a move will make the overall contract Pareto-superior to the original one. This dynamic programming problem is very similar in structure to that analysed by Thomas and Worrall (1988).¹⁶

From the Markov structure, and because each of the sustainability constraints are forward looking, the set of sustainable continuation contracts depends only on the current state. Therefore the Pareto frontier at any date t and given the current state s depends only on s and not on the past history which led to this state. To characterise the efficient contract we shall need to know the shape of the Pareto frontier and its domain of definition. This critically depends upon both the convexity of the set of sustainable contracts and the set of sustainable discounted surpluses for each household (sustainable in the sense that there exists a sustainable contract that delivers each of these surpluses).

Convexity of the set of sustainable contracts is easy to establish. Consider a convex combination of two sustainable contracts, that is, for α satisfying $0 < \alpha < 1$, define the transfer after each history h_t to be $\alpha\tau(h_t) + (1 - \alpha)\hat{\tau}(h_t)$, where $\tau(\cdot)$ and $\hat{\tau}(\cdot)$ are the original two contracts. By the concavity of both $u(\cdot)$ and $v(\cdot)$, this average contract must offer at least the average of the surpluses from the original two contracts for both households and

starting from any history h_t . Consequently the sustainability constraints (2) and (3) must be satisfied by the average contract, which is therefore itself sustainable.

Now for household i consider any pair of sustainable discounted surpluses starting at any date t in state s , and take the convex combination of the corresponding contracts as defined above. Since the average contract is sustainable, and because the discounted surplus corresponding to the average contract is continuous in α , any discounted surplus between the original pair of surpluses must be sustainable. Hence the set of sustainable discounted surpluses for each household must be an interval. For household 1 we denote this interval by $[\underline{U}_s, \bar{U}_s]$, and for household 2 by $[\underline{V}_s, \bar{V}_s]$.¹⁷ By definition the minimum sustainable surpluses for state s , \underline{U}_s and \underline{V}_s , cannot be below $-P_1(s)$ and $-P_2(s)$ respectively. However, it may not be possible to hold household i down to $-P_i(s)$ due to the non-negativity constraint on consumption. It is easily seen¹⁸ that the \underline{U}_s must be the (unique) solutions to

$$(4) \quad \underline{U}_s = \max\{u(0) - u(y_1(s)) + \delta \sum_{r=1}^S \pi_{sr} \underline{U}_r, -P_1(s)\}, \forall s \in \mathcal{S},$$

and the \underline{V}_s solve

$$(5) \quad \underline{V}_s = \max\{v(0) - v(y_2(s)) + \delta \sum_{r=1}^S \pi_{sr} \underline{V}_r, -P_2(s)\}, \forall s \in \mathcal{S}.$$

If $P_1(s) = 0$ then the minimum surplus $\underline{U}_s = 0$ and likewise if $P_2(s) = 0$ then $\underline{V}_s = 0$.

Next we define $V_s(U_s)$ to be the Pareto frontier which solves the problem of maximising, by choice of a sustainable contract commencing at date t , household 2's surplus discounted to date t , subject to giving household 1 at least U_s , given that the current state (at date t) is s .¹⁹ It should be stressed that this is an *ex post* efficiency frontier, calculated once the current state of nature is known. $V_s(U_s)$ is strictly decreasing for all $U_s \in [\underline{U}_s, \bar{U}_s]$ since, starting from any $U_s > \underline{U}_s$, in the corresponding efficient contract there must be some history h_t such that $U_t(h_t) > -P_1(s_t)$ and $y_1(s_t) - \tau(h_t) > 0$ (see equation (4)). A small increase in $\tau(h_t)$ cannot violate the sustainability constraints, but leads to an increase in household 2's utility at the expense of household 1. It follows that the constraint $U_r \leq \bar{U}_r$ can be written equivalently as $V_r(U_r) \geq \underline{V}_r$, where \underline{V}_r is defined as in (5).

The Pareto frontiers must satisfy the following optimality equations:

$$V_s(U_s) = \max_{\tau_s, (U_r)_{r=1}^S} (v(y_2(s) + \tau_s) - v(y_2(s)) + \delta \sum_{r=1}^S \pi_{sr} V_r(U_r))$$

subject to

$$(6) \quad \lambda: \quad u(y_1(s) - \tau_s) - u(y_1(s)) + \delta \sum_{r=1}^S \pi_{sr} U_r \geq U_s,$$

$$(7) \quad \delta \pi_{sr} \phi_r: \quad U_r \geq \underline{U}_r, \quad \forall r \in \mathcal{S}$$

$$(8) \quad \delta \pi_{sr} \mu_r: \quad V_r(U_r) \geq \underline{V}_r, \quad \forall r \in \mathcal{S}$$

$$(9) \quad \psi_1: \quad y_1(s) - \tau_s \geq 0,$$

$$(10) \quad \psi_2: \quad y_2(s) + \tau_s \geq 0.$$

The actual contract can be computed recursively, starting with an initial value for U_s , solving the dynamic program for the current transfer and continuation surpluses, and in each possible state r in the next period, again solving the program with target surplus U_r , and so on (see below for a discussion of the initial values of the U_s). Moreover, take any two distinct sustainable values U_s and \hat{U}_s for household 1's surplus, given that the current state is s . Now applying the same convexity argument used above to the most efficient contracts which deliver these utilities, it follows that any convex combination will offer household 1 more than $\alpha U_s + (1 - \alpha) \hat{U}_s$ and household 2 strictly more than the average of its original surpluses, by the strict concavity of $v(\cdot)$. Consequently each $V_s(\cdot)$ is strictly concave. The dynamic programming problem is thus a concave problem, and the first-order conditions are both necessary and sufficient.²⁰

The first order conditions for this problem yield the following:

$$(11) \quad \frac{v'(y_2(s) + \tau_s)}{u'(y_1(s) - \tau_s)} = \lambda + \frac{\psi_1 - \psi_2}{u'(y_1(s) - \tau_s)},$$

and

$$(12) \quad -V'_r(U_r) = \frac{\lambda + \phi_r}{1 + \mu_r}, \quad \forall r \in \mathcal{S}$$

together with the envelope condition

$$(13) \quad \lambda = -V'_s(U_s).$$

A constrained-efficient contract can be characterised in terms of the evolution over time of λ , which from equation (13) measures the rate at which household 1's surplus can be traded off ex post (once the current state is known) against that of household 2. Once the state of nature r for the following period is known, the new value of λ , which equals $-V'_r(U_r)$, is determined by equation (12). From equation (11), λ also equals the ratio of the marginal utilities of consumption, subject to the non-negativity constraints on consumption being satisfied. Since total resources in each date-state pair are given (i.e. $y_1(s) + y_2(s)$), this ties down the current transfer.²¹ Hence it is sufficient to know the evolution of λ to determine

the contract. Let $\lambda(h_t)$ be the value of λ at date t if the history is h_t . Proposition 1 shows that $\lambda(h_t)$ satisfies a simple updating rule.

Proposition 1. *A constrained-efficient contract can be characterised as follows: There exist S state dependent intervals $[\underline{\lambda}_r, \bar{\lambda}_r]$, $r = 1, 2, \dots, S$, such that $\lambda(h_t)$ evolves according to the following rule. Let h_t be given and let r be the state which occurs at time $t+1$; then*

$$(14) \quad \lambda(h_{t+1}) = \begin{cases} \underline{\lambda}_r & \text{if } \lambda(h_t) < \underline{\lambda}_r \\ \lambda(h_t) & \text{if } \lambda(h_t) \in [\underline{\lambda}_r, \bar{\lambda}_r] \\ \bar{\lambda}_r & \text{if } \lambda(h_t) > \bar{\lambda}_r. \end{cases}$$

This completely characterises the contract once an initial value for λ , λ_0 , is given.

Proof. We define $\underline{\lambda}_r := -V'_r(\underline{U}_r)$ and $\bar{\lambda}_r := -V'_r(\bar{U}_r)$, where \bar{U}_r is the maximum feasible value for U_r ; this satisfies $V_r(\bar{U}_r) = \underline{V}_r$. By the strict concavity of $V_r(\cdot)$, as U_r varies from \underline{U}_r to \bar{U}_r , so $-V'_r(U_r)$ increases from $\underline{\lambda}_r$ to $\bar{\lambda}_r$. Suppose first that $\lambda(h_t) < \underline{\lambda}_r$. Then since $\lambda(h_{t+1}) := -V'_r(U_r) \in [\underline{\lambda}_r, \bar{\lambda}_r]$, we have $\lambda(h_{t+1}) > \lambda(h_t)$, so from equation (12), $\phi_r > 0$. This implies $U_r = \underline{U}_r$, and hence $\lambda(h_{t+1}) = \underline{\lambda}_r$. A symmetric argument holds for the case $\lambda(h_t) > \bar{\lambda}_r$. Suppose finally that $\lambda(h_t) \in [\underline{\lambda}_r, \bar{\lambda}_r]$. Then if $\phi_r > 0$, we have $U_r = \underline{U}_r$ and consequently $\lambda(h_{t+1}) = \underline{\lambda}_r$, and also $\mu_r = 0$. But from equation (12) $\phi_r > 0$ and $\mu_r = 0$ imply $\lambda(h_{t+1}) > \lambda(h_t)$, a contradiction. Hence $\phi_r = 0$. By a symmetric argument $\mu_r = 0$. So by equation (12) $\lambda(h_{t+1}) = \lambda(h_t)$. \square

The idea behind this proposition can be expressed very simply. Suppose for simplicity that the non-negativity constraints on consumption never bind. Consider a first-best risk-sharing contract. This must satisfy the condition that the ratio of the two households' marginal utilities of income is constant across states and over time, and hence this contract satisfies the trivial updating condition that the current transfers are chosen to keep the marginal utility ratio equal to that of the previous period. The rule for constructing a constrained-efficient contract is as follows. If the current state is r , there is an interval of possible marginal utility ratios given by $[\underline{\lambda}_r, \bar{\lambda}_r]$. Given the marginal utility ratio last period, if possible fix the transfer this period so as to keep the ratio constant, i.e. equate the marginal utility growth for the two households. If the previous ratio lies outside the current interval, change the ratio by the minimum possible to get into the new interval. From the proof it can be seen that $\lambda = \underline{\lambda}_r$ corresponds to household 1 being held down to its minimum surplus \underline{U}_r , hence household 1 is constrained and its marginal utility growth will be lower than that of household 2. While $\lambda = \bar{\lambda}_r$ corresponds to household 1 receiving its highest possible sustainable surplus in state r , \bar{U}_r (equivalently, household 2 getting \underline{V}_r) and household 1 has a higher marginal utility growth than household 2.

It should be stressed that these intervals' endpoints, $\underline{\lambda}_r$ and $\bar{\lambda}_r$ are *optimal* values. For example, $\underline{\lambda}_r$ does not generally correspond to the lowest possible marginal utility ratio consistent with a sustainable contract starting in state r , but rather with the optimal ratio given that household 2 will be getting a minimum surplus. Suppose that the previous marginal utility ratio is less than $\underline{\lambda}_r$: it may be possible to reduce the current marginal utility ratio—by cutting c^1 —below $\underline{\lambda}_r$ so that the ratio can be kept constant; this is not however desirable since household 1's future surplus will need to be increased to offset this current loss, and this will lead overall to a worse pattern of consumption from the point of view of risk sharing.

We can think of an initial value of λ , which we denote λ_0 , as determining the distribution of the initial surplus between the two households and the given the rule of Proposition 1, the entire contract. As λ_0 varies from its minimum value of $\min_s \{\underline{\lambda}_s\}$ to its maximum value $\max_s \{\bar{\lambda}_s\}$, all constrained-efficient contracts are traced out, with higher values of λ_0 corresponding to contracts in which household 2 gets more of the potential surplus from trade.²²

Proposition 1 can also be used to demonstrate a number of results from the literature. Central results in Alvarez and Jermann (2000) as in Kocherlakota (1996) concern establishing the existence of a unique invariant long-run distribution of promised utilities. This result actually follows immediately from Proposition 1 as the λ -intervals are time independent and since the transition probabilities are assumed to be strictly positive, the probability of being in a particular state is independent of the initial state and hence of λ_0 .²³ Dixit, Grossman, and Gul (1998) analyse a model of political compromise in which power to divide a 'pie' of fixed size fluctuates between two political parties according to an exogenous stochastic process. As there is no commitment, and the parties are risk-averse and face uncertainty, their model is similar to a mutual insurance model in which the endowment process allocates the entire resource to one or other of the two agents and that agent then has the right to allocate the resource as they choose. Dixit, Grossman, and Gul (1998) show how a party's share reflects not only its current standing but also its historical support and their characterisation of efficient equilibria (Theorem 1) follows straightforwardly from our Proposition 1 and Proposition 2(iv). In the case where household 1 is risk neutral, $\lambda_s = v'(c_s^2)$ so that the updating rule translates into an updating rule for household 2 consumption. A higher value of λ now corresponds directly to a low consumption for household 2 as well as a low surplus. This is formally equivalent to the implicit contract model of Thomas and Worrall (1988). Household 2 is the worker with the contract wage equal to optimal consumption and autarky is determined by the spot market wage. The differences are that in Thomas and Worrall (1988) there are no non-negativity constraints on consumption, no direct penalties and the spot market wage is *i.i.d.*

The static limited commitment model of Coate and Ravallion (1993) too can be expressed using an updating rule which is similar to that given in Proposition 1. As stated in the introduction, the static limited commitment model can be seen as a system of gifts and transfers but without the borrowing/lending element contained in the dynamic limited commitment solution. The modified updating rule for the static limited commitment model is

$$(15) \quad \lambda(h_{t+1}) = \begin{cases} \underline{\lambda}_r & \text{if } \lambda_0 < \underline{\lambda}_r \\ \lambda_0 & \text{if } \lambda_0 \in [\underline{\lambda}_r, \bar{\lambda}_r] \\ \bar{\lambda}_r & \text{if } \lambda_0 > \bar{\lambda}_r. \end{cases}$$

where λ_0 is the initial ratio of the marginal utilities at time zero. The key difference with this modified updating rule is that the ratio of marginal utilities returns to its initial value where possible and promised utilities depend only on the current state and not the previous history. As a consequence, there is much less scope for risk-pooling; households cannot trade future claims to consumption in exchange for consumption today. Since past history is irrelevant the $[\underline{\lambda}_r, \bar{\lambda}_r]$ intervals of the static limited commitment model can be calculated for $\delta = 0$ and appropriate values for the history independent direct penalties $P_i(s)$.²⁴

Returning to the dynamic limited commitment model, Proposition 2 gives some properties of the λ -intervals with respect to the level of the direct penalties $P_i(s)$ and the discount factor δ . These comparative static properties bring out some properties of the solution and show how autarky and full risk-sharing are nested within the model. Recall that ξ_s is the autarkic marginal utility ratio, $v'(y_2(s))/u'(y_1(s))$, in state s .

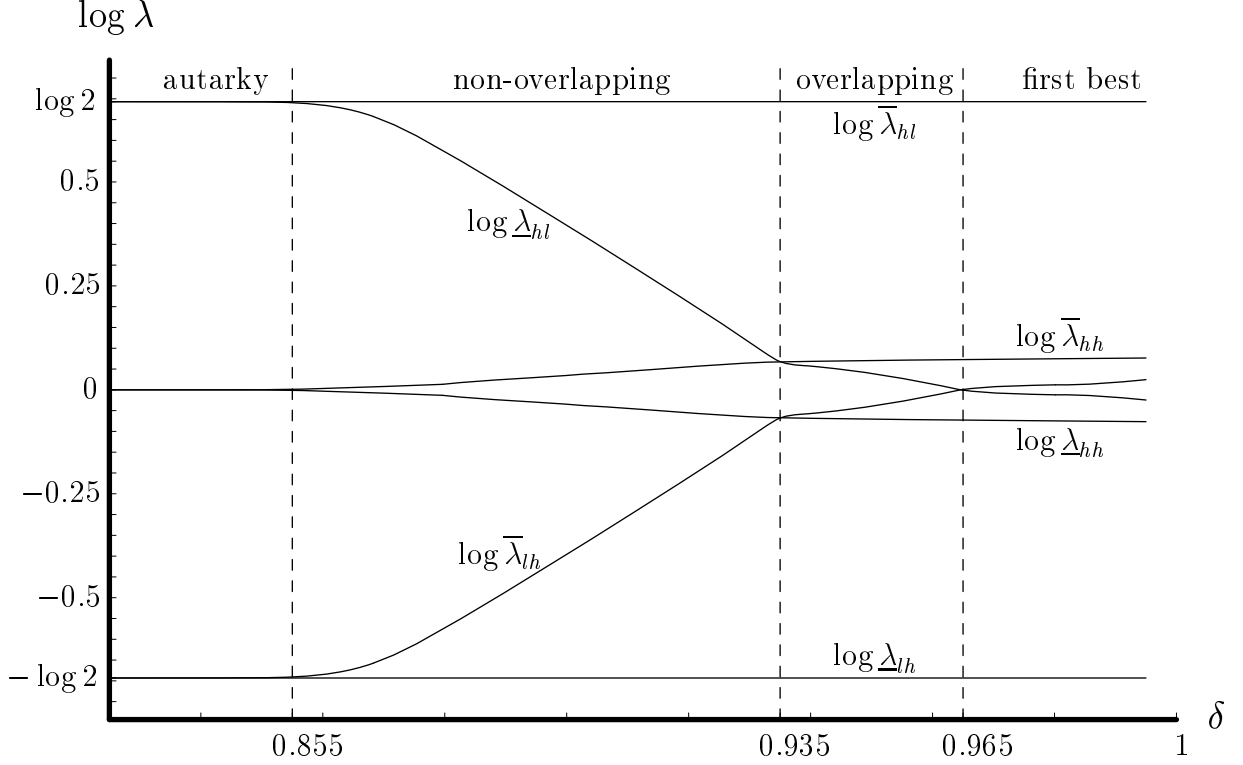
- Proposition 2.** (i). *There exist critical direct penalties, $P_1^*(s) > 0$ and $P_2^*(s) > 0$, for each s , such that if $P_i(s) \geq P_i^*(s)$, $i = 1, 2$ and $\forall s \in \mathcal{S}$, then $\forall 0 \leq \delta < 1$, $\xi_r \in [\underline{\lambda}_s, \bar{\lambda}_s]$, $\forall r, s \in \mathcal{S}$;*
- (ii). *There exists a critical δ^* , $0 \leq \delta^* < 1$, such that the intervals have non-empty intersection if $\delta \geq \delta^*$;*
- (iii). *Given $P_i(s) = P_i > 0$ for all s and $i = 1, 2$, then for each $s \in \mathcal{S}$, ξ_s is contained in the interior of the interval $[\underline{\lambda}_s, \bar{\lambda}_s]$;*
- (iv). *Given $P_i(s) = 0$ for all s and $i = 1, 2$, then for each $s \in \mathcal{S}$, $\xi_s \in [\underline{\lambda}_s, \bar{\lambda}_s]$, and $\min_s \{\underline{\lambda}_s\} = \min_s \{\xi_s\}$, $\max_s \{\bar{\lambda}_s\} = \max_s \{\xi_s\}$; moreover if $\underline{\lambda}_r < \bar{\lambda}_r$ for some r , then $\underline{\lambda}_s < \bar{\lambda}_s$ for all s , and for s such that $\max_r \{\xi_r\} > \xi_s > \min_r \{\xi_r\}$, ξ_s is contained in the interior of the interval $[\underline{\lambda}_s, \bar{\lambda}_s]$;*
- (v). *Given $P_i(s) = 0$ for all s and $i = 1, 2$, then there exists a critical $0 < \delta^{**} < 1$ such that there is no non-autarkic contract for $0 \leq \delta < \delta^{**}$;*
- (vi). *Each $\underline{\lambda}_s, \bar{\lambda}_s$ is continuous in δ , provided that each $P_i(s)$ is also continuous in δ .*

Proof: See Appendix.

Part (i) of Proposition 2 gives the obvious result that if the direct penalties are high enough, then *any* full insurance allocation is sustainable. Part (ii) implies that for a sufficiently high discount factor there is *some* first-best, full insurance contract which is sustainable. As all the λ -intervals overlap there is a λ which simultaneously belongs to each interval, and using this as the initial value to feed into the updating rule, λ remains constant thereafter; also, irrespective of initial value of λ (and hence of the division of the surplus from the contract), the contract converges with probability one to a first-best contract.²⁵ Parts (iii) and (iv) of Proposition 2 relate the λ_s -interval to its own autarkic marginal utility ratio. Specifically, when the penalties are positive but state independent, each interval will contain in its interior the autarkic marginal utility ratio for that state. If all the penalties are zero, but a sustainable non-autarkic contract exists, a similar statement is true except that the lowest (highest) endpoint of all the intervals will be the lowest (highest) autarkic marginal utility ratio. To see this notice that if that $\xi_s \geq \bar{\lambda}_s$, then the contract will never call for a transfer from household 1 to household 2 in state s no matter what the previous history. Household 2 therefore receives a non-negative short run gain from the contract in state s even when it is constrained. Therefore, if household 2 is constrained the long term loss from the continuation contract must be worse than the current penalty (because it is discounted). This can only happen if some of the future penalties are worse than the current penalty. Hence when the penalties are state independent if a household is constrained it is making a net transfer. A similar argument shows that if household 1 is constrained and penalties are state independent, then it too must be making a net transfer. Hence when penalties are state independent the autarkic marginal utility ratio must lie within its associated λ -intervals. The same argument applies if all penalties are zero, except that if, say, $\lambda = \min_s \{\underline{\lambda}_s\}$, then by the updating rule all future surpluses for household 1 are zero and since the current surplus is zero the short-run gain must also be zero and hence there is no net transfer. Thus $\min_s \{\xi_s\} = \min_s \{\underline{\lambda}_s\}$ and the analogous argument for household 2 establishes that $\max_s \{\xi_s\} = \max_s \{\bar{\lambda}_s\}$. Part (v) of Proposition 2 shows that if the households do not discount the future completely but nevertheless sufficiently heavily, then no non-trivial contract exists. Part (vi) shows that the interval endpoints, and hence also the contract, are continuous in the discount factor.

To illustrate some properties of the solution we will consider two simple examples both with *i.i.d.* income distributions. The first with two income levels for each household to show how the solution can be interpreted as a model of reciprocal borrowing and lending and the second with three income levels for each household to illustrate how the solution may involve the relatively unlucky household making a net transfer.

First consider the case where each household has an income of y_h and may suffer a loss of d with probability p , $0 < p < 1$ and there are no direct penalties. Income when there is a loss is $y_l = y_h - d$. The probability p is the same, but independent, for each household

FIGURE 1. Dependence of Intervals on δ .

and constant over time, so the expected income of each household is $y_h - pd$ in each period. There are then four states which we label hl , hh , ll and lh , where hl indicates that household 1 has high income and household 2 has low income, that is, suffers a loss, and so on. We shall consider the example where each household has identical preferences, so that $\xi_{hh} = \xi_{ll} = 1$. Full insurance with equal utilities would then involve a transfer of $d/2$ from household 1 to household 2 in state hl and a transfer of the same value from household 2 to household 1 in state lh .

We assume further that preferences can be represented by the utility function $u(c) = v(c) = \log_e(c)$. The main advantage of the logarithmic utility function for computing the example is that the λ intervals for the states ll and hh defined in Proposition 1 coincide.²⁶ Further, since all penalties are zero, $\underline{\lambda}_{lh} = \xi_{lh}$, in line with Proposition 2(iv). Given that the hh and ll λ -intervals are identical, there are only three intervals to be determined, and since preferences are identical, symmetry dictates that $\underline{\lambda}_{hl} = 1/\bar{\lambda}_{lh}$, $\underline{\lambda}_{lh} = 1/\bar{\lambda}_{hl}$ and $\underline{\lambda}_{hh} = 1/\bar{\lambda}_{hh}$. With this symmetry there are just three possible cases depending on how the intervals overlap. To calculate the interval endpoints we treat each case separately and evaluate the discounted surpluses of each household starting from the interval endpoints, where transfers are determined by equation (11) for the value of λ given by the updating rule of Proposition 1. Using the symmetry of the problem this gives us three equations in three unknowns

which we solve for the interval endpoints. In Figure 1 we plot the logarithm of the interval endpoints against the discount factor;²⁷ the logarithm is taken to preserve symmetry about the equal division of surplus line, $\log(\lambda) = 0$. From the figure, it is easy to see what are the ranges of values for the discount factor for which each of the three cases obtains. For $\delta > 0.965$ all the intervals overlap as discussed in Proposition 2(ii) and the contract converges with probability one to a first-best contract. For $0.935 < \delta < 0.965$, the intervals for states lh and hl overlap with the common interval for states hh and ll but not with each other. For $0.855 < \delta < 0.935$, none of the intervals overlap and for $\delta < 0.855$, there is no non-trivial contract. Consider the case where $0.935 < \delta < 0.965$ where the first-best is not attainable, but the lh and hl intervals overlap the hh and ll intervals.²⁸ Suppose that household 1 is the first to receive a bad shock; λ falls to $\bar{\lambda}_{lh}$, where $1 > \bar{\lambda}_{lh} > \xi_{lh} = 1/2$, and household 2 makes a transfer to household 1 so that the ratio of marginal utilities $v'(c^2)/u'(c^1)$ equals $\bar{\lambda}_{lh}$, where c^i is household i 's consumption in the contract. This is a transfer of less than $d/2$ —less than full insurance. Thereafter, until state hl occurs, $v'(c^2)/u'(c^1)$ is held constant at $\bar{\lambda}_{lh}$, which means that in the symmetric states, hh and ll , household 1 transfers income to household 2. As soon as state hl occurs the situation switches around, with λ taking on the value $\underline{\lambda}_{hl}$. This resembles a debt contract: the household that receives a bad shock receives income from the other household, but thereafter ‘repays’ this ‘loan’ at a constant rate until another bad shock is received by one of the households. When one of the households receives a bad shock²⁹ the resemblance to a standard debt contract ceases. The household suffering the latest bad shock receives a ‘loan’ of the same size as before, and starts repaying the following period. The previous history is *forgotten*, so it doesn’t matter who had previously ‘borrowed’ from whom; all that matters is who was the last to receive a loan. This idea of forgetting the past history is not specific to the two-state example and applies quite generally: when a household is constrained in a particular state at some date, the future course of the contract depends only on that state and not on the previous history up to that date.

Proposition 2(iii) shows that with state independent penalties, the constrained household makes a net transfer. As pointed out by the simulations of Attanasio and Rios-Rull (2000) however, this may involve the relatively unlucky household making a transfer, that is in an opposite direction to the transfer in the first-best outcome. To see this consider a simple example with no penalties, *i.i.d.* incomes and assume that each household has one of three incomes $y_h > y_m > y_l$. Then there are nine states, hh, hm, hl, mh , etc. By Proposition 2 (v) there is some positive discount factor such that there is no non-autarkic contract and hence by the continuity of the contract in δ (Proposition 2 (vi)) it follows that for some discount factor all intervals for states with different ξ_s ’s are disjoint and non-empty (by Proposition 2 (iv)). Suppose then, for such a discount factor, starting from state hl , we move to state ml . By Proposition 2 (iv) the autarkic marginal utility ratio is in the interior of each interval, and as $\xi_{hl} > \xi_{ml}$, this implies that the λ_{hl} -interval lies above the λ_{ml} -interval. Consequently,

by the updating rule of Proposition 1, the ratio of marginal utilities is set to $\bar{\lambda}_{ml}$, which is above the autarkic marginal utility ratio. Hence if state ml follows state hl , there will be a transfer from household 2 to household 1 even though household 2 is the relatively unlucky household, and would in a symmetric first-best contract receive a transfer from household 1. The essential point is that the debt repayment element of the solution may more than offset the static risk-sharing component.

4. EXTENSIONS

In this section we develop three extensions to the model. First, we show how to treat the case in which there is more than two households. This extension is straightforward, but involves a more complicated set of first order conditions; nonetheless, a version of the updating rule of Proposition 1 will continue to hold.

Second, we replace the simple stochastic endowments of Section 2 with a more general intertemporal technology. The new technology can be interpreted as some stochastic endowment plus storage, credit, or some more general stochastic production technology. In particular, let each household i have access to a household-specific, stochastic intertemporal technology, such that if the current state is s and the household invests k , then next period the technology returns some quantity $f_{sr}^i(k)$ in the event that the subsequent state is r . We assume that each of the functions f_{sr}^i is non-decreasing, concave, and continuously differentiable.

Third, the details of what happens to a household's claim to stored assets in the event of default can matter a great deal in determining consumption allocations, even when no default ever occurs.³⁰ Thus, we extend the individual-specific, state-contingent punishments introduced in Section 2 to permit these punishments to depend on additional quantities such as the size of a household's stock of investments. In particular, we let the continuously differentiable function $Z_s^i(k^i)$ denote the autarky utility of household i in state s if its resources available at the beginning of the period is k^i .

At this stage, it will be convenient to change notation slightly. Let there be H households, and let household i have a utility of consumption function given by $u_i(c^i)$. Denote discounted *utilities* (not surpluses) for households i in state s by U_s^i . As before, we set up the programming problem so that the current state is s , and target utilities U_s^i are given for all $i \neq H$. We introduce an additional state variable, z , the *collective* resources available to the village at the beginning of the period, which can be divided into consumption and investment. Let the current state be s . Choice variables in the programming problem will be consumption assignments c^i for $i = 1, \dots, H$, the continuation utilities U_r^i for each possible state r in the next period, and an assignment of investments k^i for each household. The value function for household H can now be written to depend on the current target utilities and collective resources: $U_s^H(U_s^1, \dots, U_s^{H-1}; z)$. Notation is otherwise as before. To simplify somewhat

we assume Inada conditions on the utility functions $u_i(\cdot)$, which allows us to disregard the non-negativity constraint on consumption. The dynamic programming problem becomes

$$U_s^H(U_s^1, \dots, U_s^{H-1}; z) = \max_{((U_r^i)_{r=1}^S)_{i=1}^{H-1}, (c^i, k^i)_{i=1}^H} u_H(c_s^H) + \delta \sum_{r=1}^S \pi_{sr} U_r^H \left(U_r^1, \dots, U_r^{H-1}; \sum_{i=1}^H f_{sr}^i(k^i) \right)$$

subject to an aggregate resource constraint,

$$\mu: \quad \sum_{i=1}^H (k^i + c^i) \leq z$$

and subject to a set of promise-keeping constraints

$$\lambda^i: \quad u_i(c_s^i) + \delta \sum_{r=1}^S \pi_{sr} U_r^i \geq U_s^i,$$

which must hold for all $i \neq H$. The solution must also be sustainable, and so satisfy the sustainability constraints

$$\delta \lambda^i \pi_{sr} \phi_r^i: \quad U_r^i \geq Z_r^i(k^i)$$

for all $r \in \mathcal{S}$, for all households $i \neq H$, and

$$\delta \pi_{sr} \phi_r^H: \quad U_r^H \left(U_r^1, \dots, U_r^{H-1}; \sum_{i=1}^H f_{sr}^i(k^i) \right) \geq Z_r^H(k^H)$$

for all $r \in \mathcal{S}$.

The first-order conditions yield

$$(16) \quad \frac{u'_H(c_s^H)}{u'_i(c_s^i)} = \lambda^i, \quad \forall i \neq H,$$

$$(17) \quad \lambda_r^i = \lambda^i \frac{1 + \phi_r^i}{1 + \phi_r^H}, \quad \forall r \in \mathcal{S}, \forall i \neq H.$$

and

$$(18) \quad u'_i(c_s^i) = \delta \sum_{r=1}^S \pi_{sr} [f_{sr}^{i'}(k_s^i) u'_i(c_r^i)] + \delta \sum_{r=1}^S \pi_{sr} \phi_r^i [f_{sr}^{i'}(k_s^i) u'_i(c_r^i) - Z_r^i(k_s^i)].$$

Note that, unlike our earlier treatment of the two household problem, it's convenient to scale the multipliers associated with the sustainability constraints by the weights λ^i . This scaling issue aside, (16) and (17) together imply exactly the same sort of updating rule for the marginal utility ratio as before, where household H 's marginal utility is treated as a

numeraire. Equation (18) is analogous to the usual Euler equation; the left-hand term is the marginal cost of increased investment associated with foregone contemporaneous consumption, while the first term on the right-hand side is the usual marginal benefit. However, the equation differs from the usual case in that there is a second term. This term reflects both additional marginal benefits measured by the terms $\phi_r^i f_{sr}^i(k^i)u_i'(c_r^i)$, and additional marginal costs measured by the terms $\phi_r^i Z_r^i(k^i)$. The former terms capture the feature that additional resources can help to relax sustainability constraints; the latter terms have to do with the problem that if too many resources are assigned to a household with low surplus, then autarky may become relatively more attractive, and thus make the sustainability constraints more binding, actually reducing welfare (see Ligon, Thomas, and Worrall (2000) for an illustration).

Although the sign of the contribution the additional terms in (18) make is generally ambiguous, in many situations optimal assignment of the k^i to agents who are unlikely to have binding sustainability constraints means that the sign will be positive. In this case, we can interpret the additional terms as a sort of endogenous “liquidity constraint,” since current consumption will be lower relative to future consumption than predicted by the usual Euler equation. Intuitively, we can think of current consumption being reduced due to some households (those who would otherwise be likely to having binding sustainability constraints in the subsequent period) posting “bonds” with other households in the village.³¹

To analyze the extension from a two-household to an H -household economy, it’s convenient to restrict our attention to the case in which utility functions exhibit constant relative risk aversion, so that household preferences are given by

$$(19) \quad \sum_{t=0}^{\infty} \delta^t \left(\frac{(c_t^i)^{1-\gamma} - 1}{1-\gamma} \right),$$

where γ is the coefficient of relative risk aversion. In one important way, the the H -household case is very similar to the two household case discussed earlier. In particular, the updating rule (17) implies that if any two households in an H -household village have non-binding sustainability constraints, then the ratio of their marginal utilities of consumption must remain unchanged. Given CRRA preferences, this further implies that the ratio of the two households’ consumptions must remain constant. Thus, if a third household *does* have a binding sustainability constraint, then other households must finance a greater share of consumption for this constrained household by a common decline in their own consumption shares.

To make the H -household sharing rule precise, we exploit the fact that CRRA utilities are Gorman aggregable to construct a sort of aggregated household representing all but household H , and then consider efficient contracts between household H and this representative household, which we’ll refer to as the “rest of the village.” To this end, let $c_s^H(z)$ denote

household H 's consumption in state s given total resources z , and let $c_s^{-H}(z)$ denote the aggregate consumption assigned to the remaining households $1, \dots, H-1$. Then the collection of individual updating equations (17) together with our assumption of CRRA utilities implies that

$$(20) \quad \left(\frac{c_r^H(z_r)}{c_r^{-H}(z_r)} \right)^\gamma = \frac{1 + \phi_r^H(z_r)}{1 + \phi_r^{-H}(z_r)} \left(\frac{c_s^H(z)}{c_s^{-H}(z)} \right)^\gamma,$$

where

$$\phi_r^{-H}(z_r) = \left[\frac{\sum_{i=1}^{H-1} (1 + \phi_r^i(z_r))^{1/\gamma} (\lambda^i)^{1/\gamma}}{\sum_{i=1}^{H-1} (\lambda^i)^{1/\gamma}} \right]^\gamma - 1.$$

Note then that the term $(1 + \phi_r^{-H}(z_r))^{1/\gamma}$, which determines when the consumption share of the rest of the village increases, is simply a weighted average of similar terms for households one through $H-1$. With CRRA preferences the H household problem collapses to a form very similar to the two household problem presented in Section 2, and yields a sharing rule very much like the two household case, with the important difference that there is now a set of intervals which depend not only on the state, but also on the household and on aggregate resources.

In terms of behaviour, the H -household sharing rule differs from the two household case only in that for the latter, if one household is constrained then the other bears all the cost, whereas this cost is spread over more households in the H -household case. As a consequence, while in the two household case we might often observe that in a given state neither household has a binding sustainability constraint, an H -household village may typically have several households with binding constraints, so there will also typically be a shared decline for unconstrained households. In fact (17) implies that the consumption shares between any two unconstrained households decline at the same rate. It does not follow, however, that a constrained household's consumption share will necessarily increase, but only that it will not fall as much as the shares of unconstrained households. Of course, if *any* household is constrained, then at least one household must receive an increased share, by adding up, and all unconstrained households must see their shares fall. Suppose that household H has a lucky endowment realization and hence a binding constraint at some date t , but that no sustainability constraint binds again for the household until period $t + \ell$. Now as a consequence of its luck at t , the household will receive an increase in its consumption share at that date. However, subsequent to t , household H 's consumption share will decline for ℓ periods, as the household must share in the collective burden of giving increased consumption to any of the *other* households who had binding constraints in the meantime. Finally, at $t + \ell$, household H again receives an increase in consumption share. These simple consumption dynamics imply that in every period the economy is divided into two groups: "winners" with new-found success, and the remaining households who finance increases in consumption for

the winners. There are no “losers” in this economy, since no household is punished for poor endowment realizations—no household’s consumption falls at a rate faster than the rate necessary to finance the winners.

5. TESTING THE MODEL

We wish to estimate the dynamic limited commitment model to see if it can help explain consumption allocations in an actual village. However, while measures of the model’s fit to the data would give us some sense of whether or not the model helps to explain the data, it would be much more satisfactory to test the model against some well-posed alternatives.

Fortunately, our model nests at least two interesting alternatives. As indicated in Sections 2 and 3, even if households’ discount factors are relatively small, Pareto optimality and full consumption insurance will be forthcoming so long as punishments ($P_i(s)$) for renegeing on contracts are sufficiently large. At the opposite pole from Pareto optimal allocations are autarkic allocations. Our model yields autarkic outcomes if the discount factor and punishments are sufficiently small. Finally, we also consider an intermediate case; the static limited commitment model of Coate and Ravallion (1993). Although this model is not nested by the dynamic model, it also nests the full insurance and autarkic allocations.

5.1. The Models. The key parameters discussed above that are required to distinguish these four models (full insurance, autarky, static limited commitment, dynamic limited commitment) were the discount factor (δ) and a state independent punishment for renegeing (P), common to all households. In addition, preferences for each household are assumed to exhibit constant relative risk aversion, and are given by (19), which gives us an additional parameter to estimate (γ). The chief sources of heterogeneity in our estimated model are idiosyncratic household endowment processes,³² and differences in the initial levels of consumption across households.

In a model with some kind of intertemporal technology (as in Section 4), households’ discount factor δ would govern not only the division of consumption, but also savings and investment decisions. Though we are able to numerically solve the model when savings is possible (Ligon, Thomas, and Worrall 2000), this is computationally expensive, and so structural estimation of the model with storage is (presently) ruled out. For pragmatic reasons, then, we will abstract from savings, storage, and investment by scaling household income in each period by a common factor so that aggregate consumption is equal to aggregate income in every period.³³ Assuming away savings also sharpens the distinctions between the different models we wish to test, since the role of discounting in the absence of an intertemporal technology is simply to determine how a fixed quantity of the consumption good ought to be divided among households, not how *much* of the consumption good ought to be allocated.

5.2. **Data.** We use data from three villages in southern India surveyed over the period 1975–1984 by the International Crops Research Institute of the Semi-Arid Tropics. We conservatively discard the first and last three years of data, because of concern over the accuracy of measured consumption in those years (Townsend 1994). Although the design of the survey was such that 40 households were surveyed in any given year, some of the households in later years replaced households lost to attrition. We restrict our attention to households continuously sampled over the entire six year period. This gives us a final sample of 34, 36, and 36 households in the three villages (Aurepalle, Shirapur, and Kanzara). The data on consumption include expenditures on food and clothing, measured at the household level. We follow Townsend (1994) in adjusting this household level measure by converting consumption and income into adult equivalents.

5.3. **Computation.** Although in principle we are able to calculate the efficient contract presented in Section 4 for economies of H households, in practice we are subject to Bellman’s curse of dimensionality. Solving the model for hundreds of households—the magnitude of the population in each village—involves an impractically large computational expense. Instead, we exploit our “rest of the village” characterization of the updating rule (20) for efficient contracts, proceeding as follows. For each household i in our sample, we solve the model *as if* there were only two households in the economy; household i , and the rest of the village (or more accurately, the rest of the sample). This allows us to compute the multipliers on the sustainability constraint for household i (ϕ_r^i) one at a time, and then to use (20) to compute the updating rule for each household’s consumption shares. Because the surplus functions computed for each successive household don’t take into account the costs of financing increases in other households’ consumption shares, this isn’t strictly correct. However, the consequences of this inconsistency are relatively unimportant for computation as long as H is fairly large.³⁴

From the analysis of Section 4, it seems likely that for the case in which the size of our sample (N) is small relative to the size of the population (H), the rest of the village approximation is probably superior to computing the sharing rule *as if* the population only consisted of N households.

5.4. **Estimation.** Suppose that the model is correctly specified up to an unknown vector of parameters θ_0 . In order to fit the model to our data, we need to solve two nested maximisation problems; an inner problem of computing the optimal contract for a given parameter vector, and an outer problem of choosing the parameter vector so as to obtain the best fit to the data. In the inner problem, we iterate on Bellman’s equation to solve the model for a given parameter vector θ . We avoid actually using a hill-climbing algorithm at each step of this iteration by taking advantage of the fact that consumption allocations will be efficient given promised utilities U_i^s ; the only inefficiency has to do with changes in these promised utilities

when sustainability constraints are binding.³⁵ As indicated above, we solve the problem for each household versus the rest of the village.

We estimate the household specific endowment processes separately from the estimation of the other structural parameters, assuming that endowment realizations are independent across both time and households, and identically distributed across time for each household. We then use a finite cell approximation to the distribution of household income, estimated nonparametrically for each household independently of all other households. The endowment process for the rest of the village is represented as an (coarsened) aggregation of each member household's endowment process. In practice we permit three possible levels of income for each household, and five possible levels for the rest of the village, so that there are fifteen possible states for each household. Thus, the state for household i at date t is determined by the cell into which own income falls and the cell into which aggregate income for the rest of the village falls.

Using a finite cell approximation for households' endowment processes raises two issues. First, since there may be extreme outcomes which occur with positive probability but which are not observed in our sample, this procedure may lead us to conclude that autarkic outcomes are more attractive than they are in fact. Second, error due to the approximation of incomes will result in a loss of information regarding small changes in income. If consumption responds to small endowment shocks, as it would if households are autarkic, then we will miss this. However, setting autarky aside, *each* of the models we test tends to predict that consumption should *not* be sensitive to small income shocks, and so the chief consequence of this approximation error will be to bias our model selection against autarky.³⁶

Having solved for the set of efficient contracts for each household, we look at the actual consumption recorded for the household in the first year of our data. We take the coefficient of relative risk aversion to be an element of the parameter vector. Given a guess for γ and observations on household and aggregate consumption, the initial value of the multiplier λ_0^i is given by

$$(21) \quad \lambda_0^i = \left(\frac{c_0^i}{c_0^{-i}} \right)^\gamma,$$

for $i = 1, \dots, H$, where, as in Section 4, c_t^{-i} denotes the rest of the village consumption at time t . We then look at actual income for each household and for the rest of the village. We choose the state closest to this actual income realization, and use this to compute the time series of multipliers associated with i 's sustainability constraints, $\{\phi_t^i\}$. Having computed this sequence of multipliers for each household in our sample, we compute the $\{\phi_t^{-i}\}$ of (20), and use this equation along with data on aggregate consumption to generate a set of predicted consumptions, $\{\hat{c}_t^i(\theta)\}$, which can then be compared with observed consumptions in the data.

We use two different estimating equations to measure the difference between predicted and actual consumptions. The first of these we term our “level” estimator, with resulting estimates denoted by $\bar{\theta}$. Computing these estimates amounts to solving a nonlinear regression problem of the form

$$(22) \quad \log c_t^i = \log \hat{c}_t^i(\lambda_t^i, \bar{c}_t|\theta) + u_t^i,$$

where λ_t^i is determined by λ_{t-1}^i and the realization of endowments according to the updating rule (14). Comparing the logarithms of predicted and actual consumption gives us a set of residuals $\{u_t^i\}$, which we interpret as being due to measurement error in consumption.

Because our main interest is in understanding the dynamics of consumption allocations, and particularly in the dynamic response of consumption to income shocks, we also work with a second estimating equation. The idea is to try and predict, not levels of consumption, but changes in households’ shares of consumption. We begin with some additional notation, to make the role of measurement error precise. For compactness, let $\tilde{\sigma}_t^i$ denote the actual consumption share of household i at time t , noting that consumption shares are related to the λ weights we worked with earlier by $\tilde{\sigma}_t^i = (\lambda_t^i)^{1/\gamma} / \sum_{j=1}^N (\lambda_t^j)^{1/\gamma}$. The econometrician doesn’t observe the actual consumption shares, but instead a quantity contaminated by measurement error, $\sigma_t^i = \tilde{\sigma}_t^i e^{v_t^i}$. Thus, writing out an expression for changes in shares of consumption yields our second estimating equation,

$$(23) \quad \log \frac{\sigma_{t+1}^i}{\sigma_t^i} = \log \frac{\hat{\sigma}_{t+1}^i(\theta)}{\hat{\sigma}_t^i(\theta)} + v_{t+1}^i - v_t^i,$$

which, along with the updating rule (20) for $t \geq 1$ and by (21) for $t = 0$ allows us to make predictions regarding the evolution of shares, conditioned on the parameter vector θ . We call the estimator of θ derived from this estimating equation our “changes-in-shares” estimator, and denote the resulting point estimates by $\hat{\theta}$.

For each of our level and changes-in-shares estimators, we take the sum of squared residuals as our measure of how well the model fits the data, thus solving a nonlinear least squares problem. Both of our estimators will be consistent and asymptotically normal so long as standard regularity conditions are satisfied (White 1984), and so long as there is no measurement error in time zero consumption. In the (more plausible) case in which time zero consumption *is* measured with error, then we need an additional condition to guarantee consistency and asymptotic normality as $T \rightarrow \infty$: for each household at least one of the sustainability constraints must eventually bind with probability one.

Of course, we have a panel with quite a small number of periods, so the property of consistency is of only limited comfort. In this connection we note that in a finite sample, our changes-in-shares estimator is apt to be much less sensitive to time zero measurement error than is our level estimator. To see this, note that (21) uses observed consumption in period zero to compute the initial value of the multipliers $\{\lambda_0^i\}$, and then uses these computed

multipliers along with data on income to predict consumption in the subsequent period. Now, if time zero consumption is measured with error, then using (21) introduces error into the calculation of the time zero multipliers, which in turn leads to error in the prediction of future consumption. However, for our changes-in-shares estimator there's no such difficulty. To see this, we evaluate (23) at $t = 0$, noting that $\hat{\sigma}_0^i(\theta) = \sigma_0^i$, so that

$$\log \frac{\sigma_1^i}{\sigma_0^i} = \log \frac{\tilde{\sigma}_1^i}{\sigma_0^i} + v_1^i.$$

and computed estimates $\hat{\theta}$ will not be directly affected by time zero measurement error.

Of course, there's another equation involved in this estimation procedure, the updating equation (20). Calculation of the updating *rule* at a fixed parameter vector θ is independent of measurement error in consumption; however, using mis-measured multipliers λ_0^i as inputs to the updating rule may produce incorrect predictions of future multipliers. Fortunately, the nature of the updating rule suggests that this may not be a serious problem. Whenever a sustainability constraint is binding, the multiplier λ_t^i will be set to an endpoint of one of the state-dependent intervals described in Proposition 1, and these endpoints are independent of initial consumption. If, on the other hand, no sustainability constraint is binding, then predicted values of the multipliers will inherit measurement error from the previous period, but the error in this case is bounded by the length of the state-dependent interval.

5.5. Comparing Models. Table 1 presents estimates of θ_0 from the level estimating equation (22) for several models: autarky, full insurance, the static limited commitment model, and the dynamic limited commitment model, along with the log quasi-likelihood associated with each model. In addition, the table reports the point estimates for the dynamic limited commitment model from the changes-in-shares estimating equation (23) (note that the parameters of the other models aren't identified by this estimating equation).

We first examine the autarkic model. As remarked above, there are no parameters to estimate for this model, so the value of the criterion function (the log quasi-likelihood) reported in Table 1 simply serves as a benchmark for the other models, each of which must have a better fit by virtue of the nesting of the models. For each of the three villages, the value of the estimating criterion is substantially less than that for any of the other models, suggesting that the other models have a substantially better fit.³⁷ The full insurance (Pareto optimal) model involves estimating the preference parameter γ .³⁸ Estimates vary greatly across villages, ranging from 0.01 in Kanzara to 26.6 in Aurepalle.³⁹ These estimates are implausible, and strongly suggest that the full insurance model with CRRA utility provides a rather poor description of the data. Nonetheless, judging by the value of the estimating criterion, the full insurance model provides a dramatic improvement to the autarkic model in terms of model fit.

Village	Model	γ	$P(1 - \delta)$	δ	Criterion
Aurepalle	Autarky	—	—	—	-207.0320
	Pareto Optimal	26.5659	—	—	-95.0120
	Static LC	1.5576	0.0044	—	-65.0657
	Dynamic LC ($\bar{\theta}$)	1.5439	0.0052	0.7048	-64.1456
	Dynamic LC ($\hat{\theta}$)	0.9501	0.0006	0.8525	—
Shirapur	Autarky	—	—	—	-197.6770
	Pareto Optimal	5.4774	—	—	-76.4276
	Static LC	1.6081	0.0027	—	-70.2262
	Dynamic LC ($\bar{\theta}$)	1.4982	0.0082	0.7093	-63.2487
	Dynamic LC ($\hat{\theta}$)	0.8420	0.0000	0.9487	—
Kanzara	Autarky	—	—	—	-153.5230
	Pareto Optimal	0.0100	—	—	-30.6478
	Static LC	1.5764	0.0043	—	-13.1712
	Dynamic LC ($\bar{\theta}$)	1.4393	0.0166	0.9485	-6.3231
	Dynamic LC ($\hat{\theta}$)	0.8435	0.0000	0.9385	—

TABLE 1. Estimates of Model Parameters. Note that for ease of interpretation, the column labelled $P(1 - \delta)$ reports the estimated punishment scaled by $1 - \delta$, so that it has the interpretation of a per-period punishment, measured in utils. For the static limited commitment model, δ is taken to be zero.

Modifying the updating rule and estimating the parameter P as well as the preference parameter γ gives us the static limited commitment model. Recall that because δ is not identified by this model, we conduct estimation for the static limited commitment model as if households were myopic ($\delta = 0$).⁴⁰ Estimated values of P vary from 0.0027 in Shirapur to 0.0044 in Aurepalle. To give some sense of the magnitude of these punishments, consider, for each village, a single household with the preference parameters reported in Table 1 which consumes the average consumption for that village with certainty. For such households, imposition of the punishment P would amount to, respectively, a 10.9, 10.7, and a 13.5 per cent reduction in certain consumption.

Perhaps more telling than measures of fit, allowing for limited commitment leads to much more sensible estimates of γ , the coefficient of relative risk aversion. Under the static limited commitment model, estimated relative risk aversion is equal to 1.56 in Aurepalle, 1.61 in Shirapur, and 1.58 in Kanzara. Relative to empirical estimates of risk aversion in other settings, these seem perfectly plausible, and seem to indicate a large improvement over the full insurance model, with its utterly implausible estimates of γ .

Estimating a third parameter, the discount factor δ , gives us the dynamic limited commitment model of this paper. We remind the reader that while this model nests the autarkic and full insurance models, it does not nest the static limited commitment model because of the difference in the updating rule for that model. We first consider the level estimates ($\bar{\theta}$) of the dynamic model. Level estimates of δ vary somewhat across villages, with the highest

value of 0.95 in Kanzara. Estimated values of δ in the remaining villages are similar—0.70 in Aurepalle and 0.71 in Shirapur. While the latter two estimates seem low relative to discount factors estimated in developed countries, they are actually somewhat higher than estimates reported by Pender (1996), who uses experimental techniques to estimate rates of discount in Aurepalle.

Level estimates of γ in the dynamic model are similar to estimates in the static limited commitment model. Estimated values of P are scaled by $(1 - \delta)$, to give the interpretation of per period punishments, to facilitate comparison with the values of P reported for the static limited commitment model. The scaled punishments for the dynamic model are uniformly larger than for the static model when measured in utils, but this is a treacherous comparison since estimated preference parameters vary across the two models. For a clearer contrast, in each village we again compute what reduction in certain consumption would correspond to the punishments reported in Table 1 for a household which consumed the village mean with certainty in every period for each village; the reductions amount respectively to an 11.8, 15.9, and a 21.1 per cent reduction in certain consumption in every period, uniformly larger punishments than for the static limited commitment model.

We next turn to the changes-in-shares estimator of the dynamic limited commitment model, comparing the parameter estimates from this model ($\hat{\theta}$) to the level estimates ($\bar{\theta}$). The changes-in-shares estimator yields considerably smaller estimates of the coefficient of relative risk aversion γ , and considerably larger estimates of the discount factor δ in the first two villages (the third is little different). Estimates of punishments (relative to average surplus) are *much* smaller, being very nearly zero in every case. The higher values for P for the level estimates ($\bar{\theta}$) lead to a predicted allocation which is closer to full insurance. As we shall see, this implies a smaller response of consumption to income changes than that observed in the data. The low values for P in the changes-in-shares estimates, on the other hand, predict a consumption response that is of about the right magnitude.

As a consequence, whether one regards the changes-in-shares estimates as more or less plausible than the level estimates, evaluating the dynamic limited commitment model at the change-in-share estimates makes the model do a much better job of capturing the dynamics which we observe in the data. To see this, consider Figure 2. The first column of this figure plots changes in actual consumption shares versus changes in income shares, and gives some sense of how sensitive consumption is to idiosyncratic income shocks. The dotted lines in these figures indicate the least squares prediction of the variable on the vertical axis given the variable on the horizontal axis. Thus, the slopes of these lines measure the average response of consumption shares to changes in income shares. These amount to about 28 per cent in Aurepalle, 12 per cent in Shirapur, and 13 per cent in Kanzara.

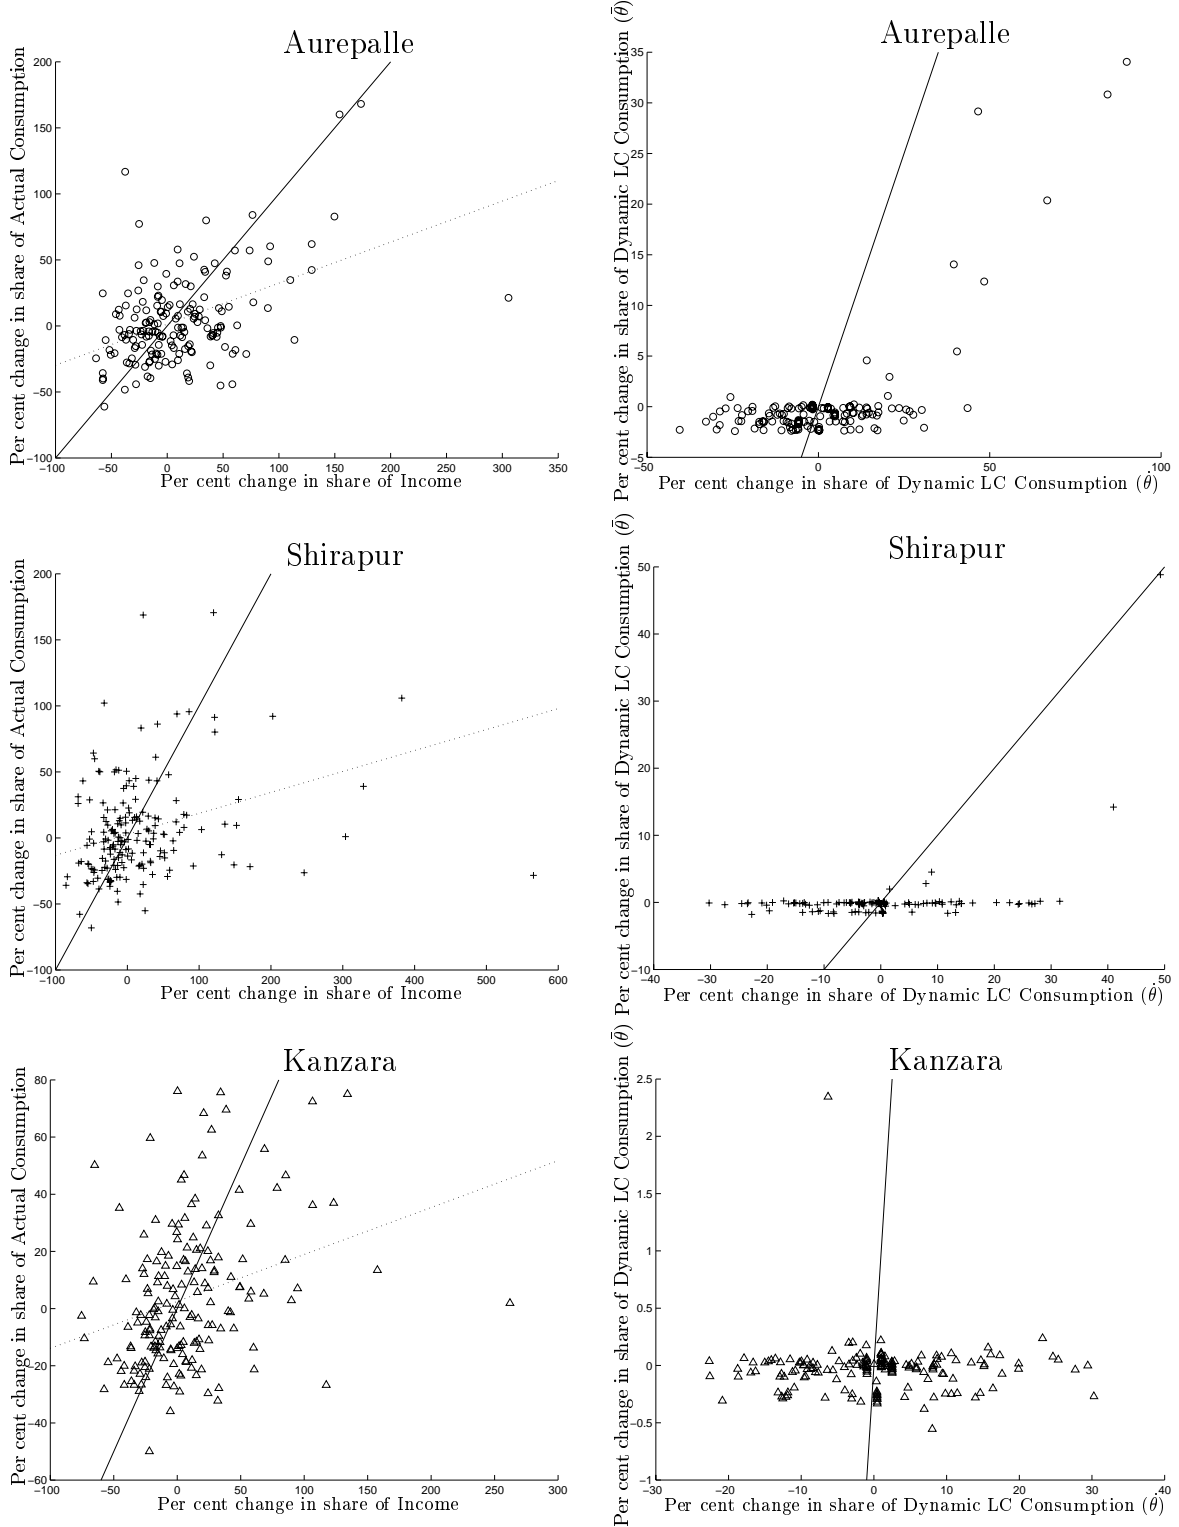


FIGURE 2. Actual Consumption & Income Response; also, Comparison of Dynamic Limited Commitment Model evaluated at $\hat{\theta}$ and θ . The horizontal axis of each plot in the first column measures changes in a household's share of village income, while the vertical axis measures the change in a household's actual share of village consumption. In the second column, we compare predicted consumptions using the level estimates for the dynamic limited commitment model (vertical axis) with predicted consumptions from the dynamic limited commitment model evaluated at the changes-in-shares estimates (horizontal axis). Each row corresponds to a different village. The solid line in each plot passes through the origin, and has a slope of 1; the dotted line is the ordinary least squares fit.

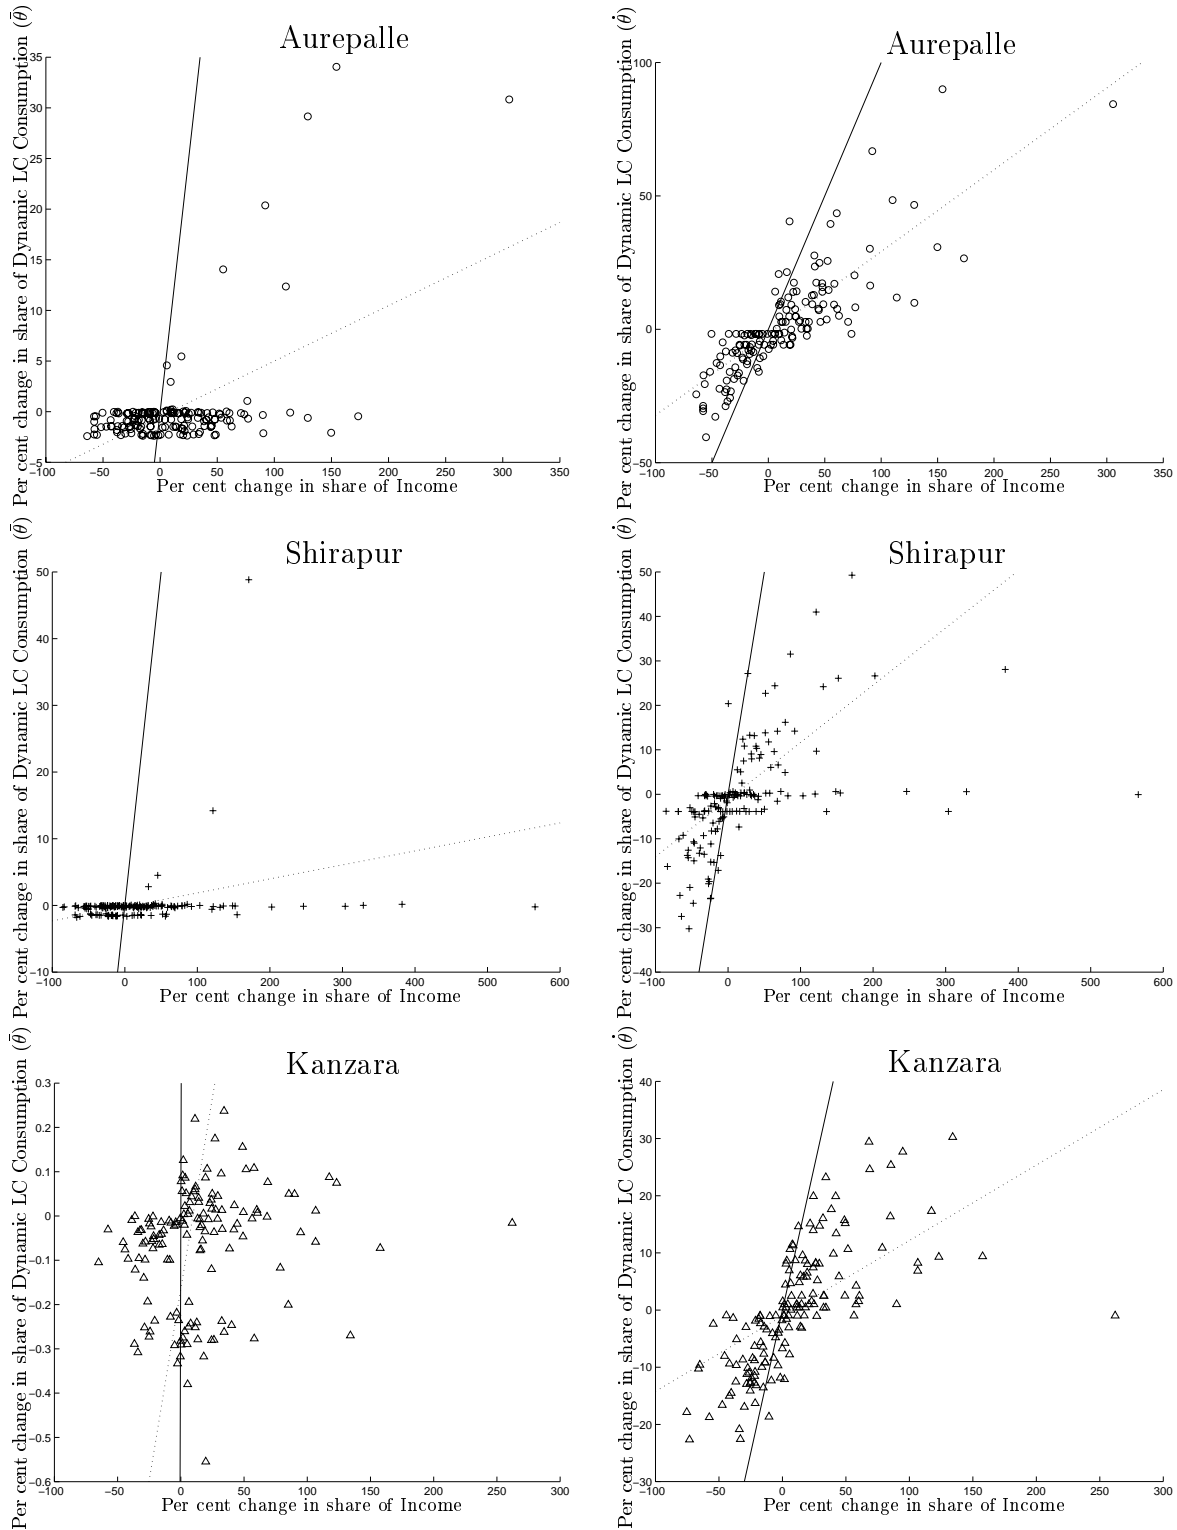


FIGURE 3. Dynamic Limited Commitment & Response to Income. The horizontal axis of each plot measures changes in a household's share of village income. The vertical axis measures the predicted change in a household's share of village consumption. Plots in the first column use the level parameter estimates for the dynamic limited commitment model while plots in the second column use the changes-in-shares estimates. Each row corresponds to a different village. The solid line in each plot passes through the origin, and has a slope of 1; the dotted line is the ordinary least squares fit.

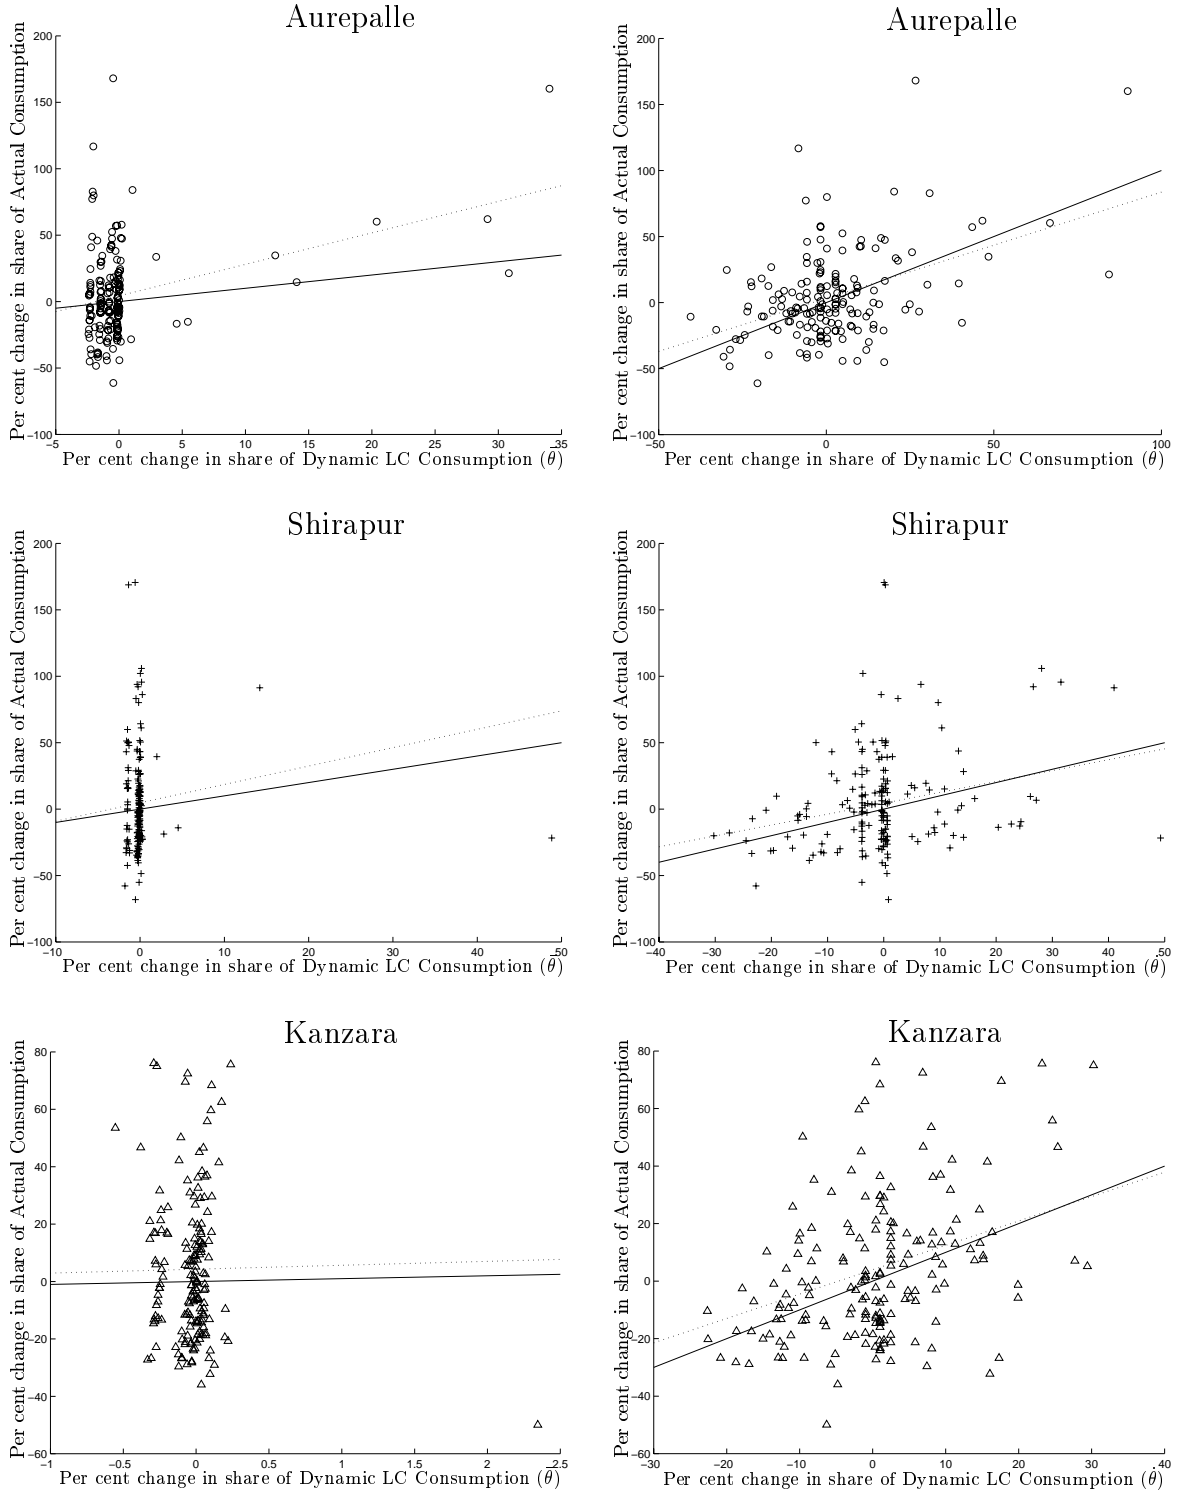


FIGURE 4. Actual & Predicted Consumption. The vertical axis of each plot measures changes in a household's actual share of village consumption. The horizontal axis measures the predicted change in a household's share of village consumption. Plots in the first column use the level parameter estimates for the dynamic limited commitment model ($\bar{\theta}$), while plots in the second column use the changes-in-shares estimates ($\hat{\theta}$). Each row corresponds to a different village. The solid line in each plot passes through the origin, and has a slope of 1; the dotted line is the ordinary least squares fit.

Where Figure 2 shows the response of actual consumption shares to changes in income shares, Figure 3 shows how the consumption shares *predicted* by the dynamic limited commitment model change in response to changes in income shares. The first column of this figure shows the response of the model evaluated at the level estimates $\bar{\theta}$, while the second column reports the response for the model evaluated at the change-in-shares estimates $\dot{\theta}$. As in Figure 2, the solid lines have a slope of one, to help give an indication of relative scale, while the dotted lines correspond to the OLS prediction. The difference between the two sets of estimates is remarkable. In particular, the level model predicts a very small response (the slope of the OLS lines are about 0.03 per cent in each village), while the change-in-shares model predicts average percentage responses of about 12.5 per cent in Aurepalle and Shirapur, and about 13 per cent in Kanzara, very much in line with the the actual average response shown in the first column of Figure 2. This point is reinforced by Figure 4, which plots changes in actual consumption shares versus the predictions of the dynamic limited commitment model, evaluated at $\bar{\theta}$ in the first column, and $\dot{\theta}$ in the second. The level estimates do quite poorly on this comparison, but the changes-in-shares predictions do a fair job of matching the variation observed in actual consumption, in that the estimated OLS fit has a slope slightly less than one in each village. There's greater observed variation in the actual data, but that's exactly what we'd expect given that consumption is measured with error.

Although the changes-in-shares estimates do a much better job of capturing dynamics than do the level estimates, to some extent that's to be expected. The level estimator is more ambitious than the changes-in-shares estimators, since it tries to match not just dynamics but also variation in mean consumption across households. So, as one might expect, the level estimator does a considerably better job of explaining the distribution of consumption than does the changes-in-shares estimator. Figure 5 illustrates this point. In the first column of this figure we plot a set of Lorenz curves, each row corresponding to a different village. Each plot contains the Lorenz curve for each year, and for actual consumption, the changes-in-shares prediction of consumption levels, and the level prediction of consumption levels. Lorenz curves for actual consumption are solid lines, while Lorenz curves using the level prediction are dashed lines, and Lorenz curves using the changes-in-shares predictions are dotted lines. With so many lines, it's not always easy to draw clear comparisons; accordingly, in the second column of the figure we plot the actual Lorenz curves minus the predicted Lorenz curves.

The most striking point in these figures is that the level prediction of distribution is much closer to the actual distribution than is the changes-in-shares prediction. In every case, the changes-in-shares estimates tend to predict much greater inequality in the distribution of consumption than is actually observed. This effect is particularly marked for wealthier (high consumption) households. In two of three villages (Aurepalle and Kanzara), predictions

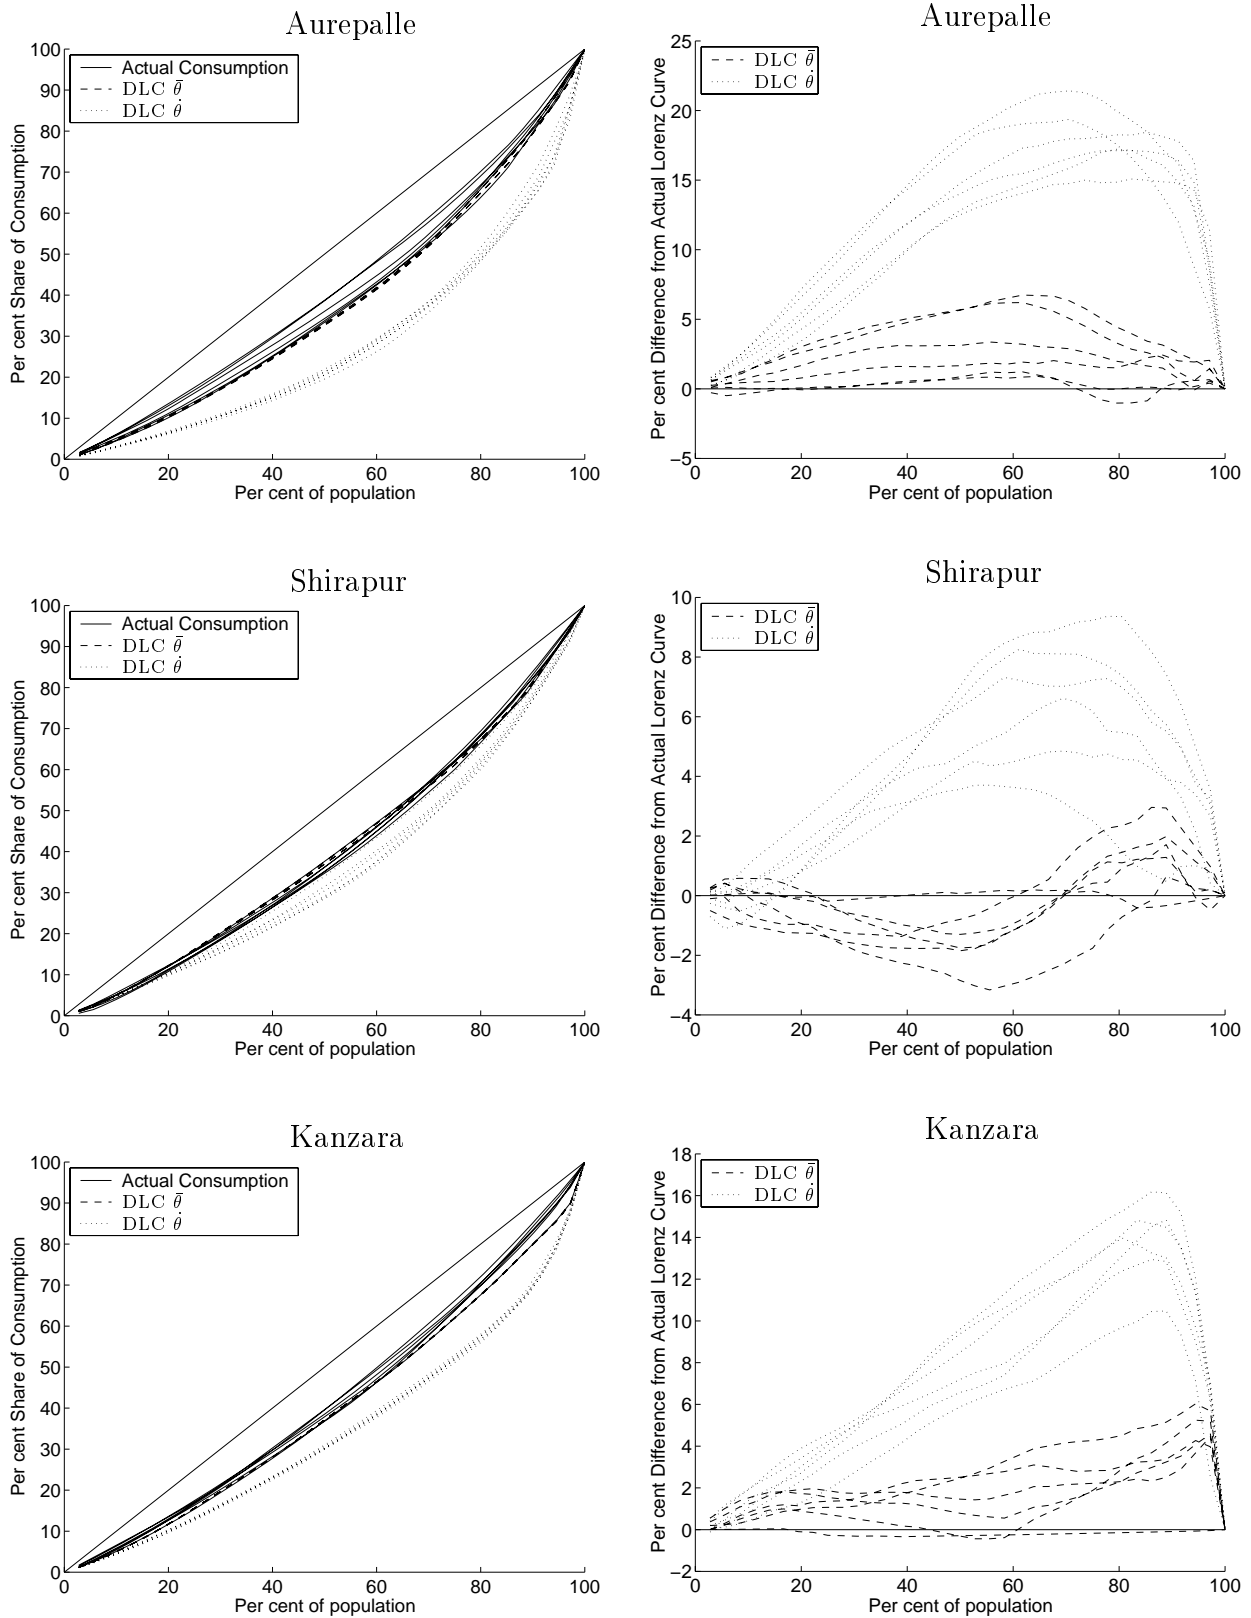


FIGURE 5. Actual and Predicted Lorenz Curves. Curves marked “DLC $\bar{\theta}$ ” use the level parameter estimates for the dynamic limited commitment model while curves marked “DLC $\hat{\theta}$ ” use the change-in-shares estimates. The first column of plots are Lorenz curves of consumption (one curve per year). The second column of plots show the difference between actual and predicted Lorenz curves. Each row corresponds to a different village.

using the level estimates also display too much inequality; however, in Shirapur too little inequality is observed among all but the richest households.

A second important point regarding these figures is that errors in predicted inequality tend to persist across years. For example, in Aurepalle changes-in-shares predictions regarding distribution change rather little across years; in every year outcomes are more egalitarian than predicted, not simply on average.

6. AN APPRAISAL OF THE DYNAMIC LIMITED COMMITMENT MODEL

We have presented a variety of evidence that models incorporating limited commitment seem capable of doing a much better job of explaining actual consumption allocations than are models of full insurance, static limited commitment, or autarky. Nonetheless, given the simplicity of the models we've proposed and the necessarily stylised features of the model economy we've estimated, we would like to have some additional ways to evaluate the performance of the dynamic limited commitment model we've constructed relative to the performance of alternative models.

We begin by examining some informal measures of fit.⁴¹ Table 2 presents the correlation between predicted and actual consumptions. The first column of this table (labelled "Reality") is of the greatest interest, as it gives the correlations between the actual data and consumption in each of the proposed models. Setting aside the "Dynamic LC ($\hat{\theta}$)" row of this table for the moment, the orderings of models according to how highly their predicted consumptions are correlated with actual consumptions is as one would expect from the measures of fit reported in Table 1, with a single exception: income (autarkic consumption) in Aurepalle is actually more highly correlated with consumption than is predicted consumption with full insurance. The highest correlations with actual consumption are given by the dynamic limited commitment model, which has a correlation of 0.73 in Aurepalle, 0.61 in Shirapur, and 0.76 in Kanzara. Since we're looking at correlations between levels of (predicted) consumptions, correlations between the logarithms of actual consumption and the logarithms of the consumptions predicted by the changes-in-shares models tend to be relatively low, reflecting the same shortcoming observed earlier in discussion of Figure 5.

Correlations reported in the first column of Table 2 are uniformly less than one because of some error. This may be innocuous noise in the measurement of actual consumption, or may be error in predicted consumptions. To better understand the error in our model, we investigate whether residuals from the models we estimate are systematically related to income. Thus, in Table 3, we consider a regression in which the dependent variable is the residuals from the limited commitment models, $c_t^i - \hat{c}_t^i$, which we then regress on other variables in an attempt to identify shortcomings. In every case right hand side variables include a set of fixed effects, a set of time effects, and the logarithm of income. The idea is that the predictions of a model which successfully captures consumption's response to income

Village	Model	Reality	Autarky	Full Ins.	Static LC	DLC ($\bar{\theta}$)
Aurepalle	Autarky	0.659	1.000	—	—	—
	Full Ins.	0.600	0.579	1.000	—	—
	Static LC	0.731	0.729	0.883	1.000	—
	Dyn. LC ($\bar{\theta}$)	0.734	0.698	0.908	0.980	1.000
	Dynamic LC ($\dot{\theta}$)	0.712	0.943	0.642	0.806	0.791
Shirapur	Autarky	0.491	1.000	—	—	—
	Full Ins.	0.574	0.347	1.000	—	—
	Static LC	0.590	0.320	0.975	1.000	—
	Dyn. LC ($\bar{\theta}$)	0.616	0.393	0.947	0.971	1.000
	Dynamic LC ($\dot{\theta}$)	0.526	0.832	0.449	0.464	0.537
Kanzara	Autarky	0.592	1.000	—	—	—
	Full Ins.	0.685	0.655	1.000	—	—
	Static LC	0.745	0.616	0.948	1.000	—
	Dyn. LC ($\bar{\theta}$)	0.756	0.630	0.943	0.995	1.000
	Dynamic LC ($\dot{\theta}$)	0.653	0.897	0.699	0.739	0.751

TABLE 2. Simple correlations between the logarithm of consumption from different models.

$$\log c_t^i - \log \hat{c}_t^i = \alpha_i + \beta \log y_t^i + \eta_t + \epsilon_t^i$$

Village	β		
	SLC	DLC ($\bar{\theta}$)	DLC ($\dot{\theta}$)
Aurepalle	0.336* (0.052)	0.274* (0.050)	-0.055 (0.052)
Shirapur	0.186* (0.044)	0.157* (0.043)	0.037 (0.043)
Kanzara	0.212* (0.049)	0.212* (0.049)	-0.006 (0.048)
Pooled	0.236* (0.027)	0.203* (0.027)	0.001 (0.027)

TABLE 3. Consumption regressions to test the null hypothesis of limited commitment. The left hand side variable is the logarithm of actual consumption minus the logarithm of predicted consumption, where consumptions is predicted according to one of the models indicated in the table heading. The right hand side variables in every case are a set of time effects, a set of household fixed effects, and the logarithm of household income. Numbers in parentheses are standard errors.

will differ from actual consumption only in ways which are independent of income. Both the static limited commitment model and the dynamic limited commitment model evaluated at the level estimates do relatively poorly by this measure; after subtracting predicted log consumption from actual consumption, the residual is still significantly related to the log of income. In contrast, the dynamic limited commitment model evaluated at the changes-in-shares estimate does extremely well: differences between predicted log consumption and actual log consumption *are* nearly orthogonal to the logarithm of income. Furthermore, the

changes-in-shares estimates seem fairly plausible, delivering estimates of preference parameters γ and δ well within the range found in empirical studies in other settings. This is an important success, indicating that the dynamic limited commitment model (evaluated at the changes in shares estimates) is capable of explaining the response of consumption to income, the puzzle raised by Townsend (1994). Note, however, that the fixed effects in this last regression are jointly significant; as noted above, the dynamic limited commitment model evaluated at the changes-in-shares estimates does not do particularly well at explaining the distribution of consumption across households.

7. DISCUSSION AND DIRECTIONS FOR FUTURE RESEARCH

Recent work on consumption allocation in village economies has established a model of full risk-pooling as a natural benchmark—such a model provides a much better explanation of consumption than does a model of autarky, for example. Nonetheless, there are well documented shortcomings with the full risk-pooling model in most of the environments in which it’s been tested. In particular, tests of the full risk-pooling model leave us with a puzzle: why is idiosyncratic variation in consumption systematically related to idiosyncratic variation in income?

In this paper we attempt to solve this puzzle by adding limited commitment as a possible impediment to risk-pooling. We provide a complete characterization of consumption allocations in a dynamic environment with limited commitment, and then use data on three Indian villages with which to test the dynamic limited commitment model.

The dynamic limited commitment model does better than any of several alternatives in explaining actual consumption allocations. It provides a better explanation than the benchmark of full risk-pooling; it also performs better than the static limited commitment discussed by Coate and Ravallion (1993). Most importantly, when the dynamic limited commitment model is evaluated at the plausible changes in shares estimates, it predicts an average response of idiosyncratic consumption to idiosyncratic income shocks in line with the empirical response.

Despite these successes, the limited commitment model presented here does not adequately explain consumption. In particular, making the model explain differences in the distribution of mean consumption (as with our level estimator) compromises the ability of the model to simultaneously explain dynamics, particularly the response of consumption to idiosyncratic income shocks. Conversely, if we abandon the goal of explaining distribution, and focus solely on dynamics, then these can be captured very nicely, but at the cost of predicting more consumption inequality than is actually observed in the data.

Why is the dynamic limited commitment model incapable of simultaneously explaining distribution and dynamics? Several possible explanations occur to us. One simple explanation is that rather than having preferences which exhibit constant relative risk aversion, perhaps

households have increasing relative risk aversion—this would be consistent with the argument of Arrow (1965) and empirical work by Binswanger and Rosenzweig (1986), for example. Second, it's possible that using the household as the unit of analysis is a mistake—there may be increasing returns to scale within households or other intra-household allocational issues which would confuse our analysis, in line with the findings of Deaton and Paxson (1998). Because wealthier households also tend to be larger households, this explanation seems consistent with our findings here. Third, our neglect of savings and intertemporal production could possibly explain the result.

Each of these possible explanations amounts to a call for future research. We believe that three different avenues are particularly promising. First, although Section 4 addressed some of the issues regarding the modelling of investment or storage, we've ignored this possibility in our empirical model. The model is considerably more complicated when storage is possible as individual savings can be used to help smooth aggregate consumption when aggregate income is uncertain and households with high savings may find autarky relatively more attractive. Thus it may be desirable to transfer assets across households to improve welfare. Ligon, Thomas, and Worrall (2000) study these issues. What emerges is a theory of dynamic distribution in which shifts in the ownership of storable assets help support an equilibrium. The theory suggests that the existence of an intertemporal technology is likely to lower the cost of smoothing consumption. As a consequence, neglecting such a technology might lead us to incorrectly infer that demand for smoothing is greater than it is in fact. Since demand for smoothing is greater when risk aversion is larger, punishments greater, and households more patient, we would guess that neglect of intertemporal technology would lead to upwardly biased estimates of one or more of these three parameters.

Second, it is important to examine the policy implications of the analysis and the interaction between informal mutual insurance arrangements and government sponsored schemes. Attanasio and Rios-Rull (2000) have considered this issue. They model an economy with many private mutual insurance schemes subject to limited commitment and examine the implications of a compulsory public insurance scheme against aggregate shocks. They find that the welfare effect of such a public insurance scheme is ambiguous as it increases autarky utility and therefore may limit what is sustainable and reduce the extent of private risk-pooling. Third, the role of limited information could be incorporated. Moral hazard and adverse selection are likely to play a role in determining consumption allocations, consistent with the findings of Ligon (1998). Future research might usefully consider both limited commitment and limited information simultaneously.

APPENDIX

7.1. Proof of Proposition 2. Proof. (i) If $\delta = 0$, for sufficiently high penalties in any state s , $\lambda(h_t)$ equal to any ξ_r is sustainable; it follows that with such penalties, for any $\delta > 0$, any efficient allocation with a (constant) marginal utility ratio (and hence constant $\lambda(h_t)$) lying weakly between the maximum and minimum autarkic ratios is also sustainable. By Proposition 1 this is only possible if each interval contains this same range of values. (ii) Let $\delta^* < 1$ be the minimum value of δ such that a first-best contract is sustainable; this exists from usual ‘Folk theorem’ arguments (this requires that $\pi_{sr} > 0$ for all s, r , as we assumed, or at least that all states communicate in the sense that each state is reached with positive probability from each other state). From the definition of $\lambda(h_t)$ a first-best contract requires that $\lambda(h_t)$ is constant for all h_t ; this is possible from Proposition 1 if and only if the intervals have common intersection. The result follows. (iii) Rewrite the sustainability constraint (2) as

$$(24) \quad U_t(h_t) = u(y_1(s_t) - \tau(h_t)) - u(y_1(s_t)) + \delta E(U_{t+1}(h_{t+1})) \geq -P_1.$$

Suppose $s_t = s$ and that the current value of $\lambda(h_t)$ is $\underline{\lambda}_s$ (recall that this means that household 1’s surplus is at its minimum sustainable level of \underline{U}_s). Either household 1’s consumption is zero at time t , in which case $\underline{\lambda}_s < \xi_s$, or the sustainability constraint binds (compare equation (4)). In the latter case, since $U_{t+1}(h_{t+1}) \geq -P_1$ for all s_{t+1} , we have from equation (24), which binds, $u(y_1(s_t) - \tau(h_t)) < u(y_1(s_t))$, which again implies that $\underline{\lambda}_s < \xi_s$. A symmetric argument for household 2 establishes that $\bar{\lambda}_s > \xi_s$. (iv) The first part follows from part (iii) above, replacing P_i by 0, and strict inequalities by weak ones. Next, if state s has the lowest $\underline{\lambda}_s$, suppose that $\lambda(h_t) = \underline{\lambda}_s$. Then the updating rule of Proposition 1 implies that $\lambda((h_t, r)) = \underline{\lambda}_r$ for all states r occurring at date $t + 1$; hence future utilities $U_{t+1}((h_t, r))$ in each state equal \underline{U}_r , which equals zero when $P_1(r) = 0$ for all r . Likewise $U_t(h_t) = 0$, so equation (24) implies that $u(y_1(s_t) - \tau(h_t)) - u(y_1(s_t)) = 0$, and so consumption for household 1 is at the autarkic level at $\underline{\lambda}_s$. Since each $\xi_r \geq \underline{\lambda}_r$, it follows that $\xi_s = \min_r \xi_r$. A symmetric argument for household 2 establishes the result for $\max\{\bar{\lambda}_s\}$. Finally, if at least one interval r is non-degenerate, then this implies that starting in that state, non-autarkic contracts are sustainable. Then by $\pi_{sr} > 0$, the same is true for all other states s since $\delta > 0$, by (v) below (e.g., specify no transfer at the current s , followed by the non-autarkic contract if r occurs next period and autarky otherwise), so they too have non-degenerate intervals. If state s does not have the lowest $\underline{\lambda}_s$, and suppose $s_t = s$ and $\lambda(h_t) = \underline{\lambda}_s$, then if it is followed by state r where $\underline{\lambda}_r < \underline{\lambda}_s$, by the updating rule and the non-degeneracy of each interval, $\lambda((h_t, r)) > \underline{\lambda}_r$. This implies that $U_{t+1}((h_t, r)) > 0$ and given $U_{t+1}(h_{t+1}) \geq 0$, equation (24), which binds, implies that $u(y_1(s_t) - \tau(h_t)) < u(y_1(s_t))$, and so consumption is below the autarkic level for household 1 at $\underline{\lambda}_s$, i.e., $\xi_s > \underline{\lambda}_s$. Since $\min_r \xi_r = \min_r \underline{\lambda}_r$, it follows that s is such that $\xi_s > \min_r \xi_r$. A symmetric argument for household 2 completes the proof. (v) Let $\bar{\tau}_s$ be the transfer in state s if $\lambda = \bar{\lambda}_s$ and let $\underline{\tau}_s$ be the transfer in state s if $\lambda = \underline{\lambda}_s$. From part (iii) $\underline{\tau}_s \geq 0 \geq \bar{\tau}_s$ and from part (iv) $\underline{\tau}_s = 0$ when $\underline{\lambda}_s = \min\{\xi_s\}$ and $\bar{\tau}_s = 0$ when $\bar{\lambda}_s = \max\{\xi_s\}$. For $\delta = 0$ there is clearly no non-autarkic contract. The contract is continuous in δ (see (vi) below), and by part (i) all intervals overlap for large δ , so for δ small the λ -intervals will be disjoint if the the autarkic marginal utility ratios are distinct and approximately coincident if two or more states have the same the autarkic marginal utility ratios. Using Proposition 1, it is then possible to calculate

$$v(y_2(s)) - v(y_2(s) + \bar{\tau}_s) = \delta \sum_{r \text{ s.t. } \xi_r > \xi_s} [\alpha_{sr} \pi_{sr} (v(y_2(r) + \underline{\tau}_r) - v(y_2(r) + \bar{\tau}_r))],$$

$$u(y_1(s)) - u(y_1(s) - \underline{\tau}_s) = \delta \sum_{r \text{ s.t. } \xi_r < \xi_s} [\alpha_{sr} \pi_{sr} (u(y_1(r) - \bar{\tau}_r) - u(y_1(r) - \underline{\tau}_r))],$$

where α_{sr} is some positive parameter. Let $\chi_{rs}^u = u'(y_1(r))/u'(y_1(s))$ and $\chi_{rs}^v = v'(y_2(r))/v'(y_2(s))$. Linearising these two equations about the income levels and adding gives

$$\begin{aligned} (\underline{\tau}_s - \bar{\tau}_s) = & \delta \sum_{r \text{ s.t. } \xi_r > \xi_s} [\alpha_{sr} \pi_{sr} (\underline{\tau}_r - \bar{\tau}_r) \chi_{rs}^v] \\ & + \delta \sum_{r \text{ s.t. } \xi_r < \xi_s} [\alpha_{sr} \pi_{sr} (\underline{\tau}_r - \bar{\tau}_r) \chi_{rs}^u] + o(\underline{\tau}_s - \bar{\tau}_s). \end{aligned}$$

Let $\beta_s = \max\{\alpha_{sr} \pi_{sr} \chi_{rs}^v, \alpha_{sr} \pi_{sr} \chi_{rs}^u\}$ and let $(\underline{\tau}_k - \bar{\tau}_k) = \max\{(\underline{\tau}_s - \bar{\tau}_s)\}$. By choosing δ small enough β_k can be made arbitrarily small, say $\beta_k < 1/S$. Then if $(\underline{\tau}_k - \bar{\tau}_k) > 0$ it follows that $0 < (\underline{\tau}_k - \bar{\tau}_k) \leq (S - n - 1)\beta_k(\underline{\tau}_k - \bar{\tau}_k) + o(\underline{\tau}_k - \bar{\tau}_k)$ where n is the number of states with the same autarkic marginal utility ratio as state k . Hence $\beta_k \geq (1/(S - n - 1)) - O(\underline{\tau}_k - \bar{\tau}_k)$ contradicting the assumption that $(\underline{\tau}_k - \bar{\tau}_k) > 0$ so there can be no non-autarkic contract for δ small. (vi) Fix s , and consider a sequence $\{\delta_v\}$ converging to $\hat{\delta}$, and associated constrained efficient contracts τ_v which give household 1 its minimum surplus, which we write as $\underline{U}(s; \delta_v)$ to allow for possible dependency on δ , starting from initial state s . Note that $\underline{U}(s; \delta)$ is continuous in δ given that the punishments are. Because the space of contracts is sequentially compact (in the product topology), take a convergent subsequence which we also denote $\{\tau_v\}$, and let $\hat{\tau}$ be the limit (i.e., the pointwise limit). Fix any h_t (with $h_1 = s$), and note that $\lim_{v \rightarrow \infty} U_t(h_t; \delta_v) = U_t(h_t; \delta)$, where $U_t(h_t; \delta_v)$ denotes the continuation surplus at δ_v with the contract τ_v ; as $U_t(h_t; \delta_v) \geq \underline{U}(s_t; \delta_v)$ by assumption of sustainability, we have $U_t(h_t; \delta) \geq \underline{U}(s_t; \delta)$, and likewise for household 2, so the limit contract $\hat{\tau}$ is sustainable. To prove that it is constrained efficient, suppose otherwise, so there is some other contract τ^* such that $V_s(\tau^*; \hat{\delta}) > V_s(\hat{\tau}; \hat{\delta})$, where $V_s(\tau; \delta)$ denotes household 2's surplus at δ from τ when the initial state is s . Consider the contract $\mu\tau^*$, where $0 < \mu < 1$, in which all transfers are scaled down by factor μ . Note that it is sustainable by virtue of being a convex combination of the sustainable contracts τ^* and the autarkic contract. Choose μ so that $V_s(\mu\tau^*; \hat{\delta}) > V_s(\hat{\tau}; \hat{\delta})$. Because τ^* obeys the updating rule of Proposition 1, $\mu\tau^*$ can be represented as a finite state Markov chain; moreover by strict concavity of $v(\cdot)$, at each of these finite states, the sustainability constraint for household 2 is slack. The same is true of household 1 if $u(\cdot)$ is strictly concave; if it is linear, a small additional transfer can be made to 1 in each state of the Markov chain to ensure slackness without violating $V_s(\mu\tau^*; \hat{\delta}) > V_s(\hat{\tau}; \hat{\delta})$. Because all constraints are slack, $\mu\tau^*$ is also sustainable for δ_v sufficiently close to $\hat{\delta}$, and $V_s(\mu\tau^*; \delta_v) > V_s(\tau_v; \delta_v)$, contradicting the assumed optimality of τ_v . We conclude that $\hat{\tau}$ is optimal, and thus $\hat{\tau} = \tau^*$ since the optimal contract is unique. Thus given $\{\delta_v\}$ converging to $\hat{\delta}$, there is a subsequence for which τ_v converges to τ^* . Since $\underline{\lambda}_s = v'(y_2(s) + \tau(s))/u'(y_1(s) - \tau(s))$, the $\underline{\lambda}_s$ also converge on the subsequence to the value at $\hat{\delta}$. This is sufficient to establish continuity of $\underline{\lambda}_s$. A symmetric argument establishes the continuity of the $\bar{\lambda}_s$. ■

REFERENCES

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica* 56, 383–396.
- Alvarez, F. and U. J. Jermann (2000). Efficiency, equilibrium, and asset pricing with risk of default. *Econometrica* 68(4), 775–797.
- Arrow, K. J. (1965). *Aspects of the Theory of Risk Aversion*. Helsinki: Yrjö Jahnssonin säätiö.
- Asheim, G. and J. Strand (1991). Long-term union-firm contracts. *Journal of Economics* 53, 161–184.

- Atkeson, A. (1991). International lending with moral hazard and risk of repudiation. *Econometrica* 59, 1069–1089.
- Attanasio, O. and J.-V. Rios-Rull (2000). Consumption smoothing in island economies: Can public insurance reduce welfare? *European Economic Review* 44(7), 1225–1258.
- Beaudry, P. and DiNardo (1995). Is the behaviour of hours worked consistent with implicit contract theory? *Quarterly Journal of Economics* 110, 743–768.
- Binswanger, H. P. and M. R. Rosenzweig (1986). Credit markets, wealth and endowments in rural south India. Presented at the 8th International Economic Association World Congress.
- Bliss, C. and N. Stern (1982). *Palanpur: The Economy of an Indian Village*. Oxford University Press.
- Bulow, J. and K. Rogoff (1989, March). Sovereign debt: Is to forgive to forget? *American Economic Review* 79, 43–50.
- Cashdan, E. (1985). Coping with risk: Reciprocity among the Basarwa of northern Botswana. *Man* 20, 454–474.
- Chari, V. V. and P. Kehoe (1993). Sustainable plans and mutual default. *Review of Economic Studies* 60(1), 175–198.
- Coate, S. and M. Ravallion (1993). Reciprocity without commitment: Characterization and performance of informal insurance arrangements. *Journal of Development Economics* 40, 1–24.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99, 957–976.
- Deaton, A. (1992). Household saving in LDCs: Credit markets, insurance, and welfare. *Scandinavian Journal of Economics* 94, 253–273.
- Deaton, A. and C. Paxson (1998). Economies of scale, household size, and the demand for food. *Journal of Political Economy* 106(5), 896–930.
- Dixit, A., G. M. Grossman, and F. Gul (1998). The dynamics of political compromise. *The Journal of Political Economy* 108(3), 531–568.
- Dubois, P. (2000). Assurance complète, hétérogénéité de préférences et mètayage du Pakistan. *Annales d’Economie et de Statistique* forthcoming.
- Eswaran, E. and A. Kotwal (1989). Credit as insurance in agrarian economies. *Journal of Development Economics* 31, 37–53.
- Fafchamps, M. (1999). Risk sharing and quasi-credit. *Journal of International Trade and Economic Development* 8(3), 257–278.
- Farrell, J. and E. Maskin (1989). Renegotiation in repeated games. *Games and Economic Behavior* 1, 327–360.
- Foster, A. and M. Rosenzweig (1995). Imperfect commitment, altruism and the family: Evidence from transfer behaviour in low-income rural areas. Manuscript.
- Gauthier, C., M. Poitevin, and P. González (1997). Ex ante payments in self-enforcing risk-sharing contracts. *Journal of Economic Theory* 76, 106–144.
- Gobert, K. and M. Poitevin (1998). Non-commitment and savings in dynamic risk-sharing contracts. manuscript.
- Grimard, F. (1997). Household consumption smoothing through ethnic ties: Evidence from Cote d’Ivoire. *Journal of Development Economics* 53, 391–422.
- Hansen, B. E. (1992). The likelihood ratio test under nonstandard conditions: Testing the Markov switching model of GNP. *Journal of Applied Econometrics* 7(Supplement), S61–S82.
- Hayashi, F. (1996). Analysis of household saving: Past, present, and future. *The Japanese Economic Review* 47(1), 21–33.
- Holmström, B. (1983). Equilibrium long-term labor contracts. *Quarterly Journal of Economics* 98(3), 23–54. Supplement.

- Jacoby, H. and E. Skoufias (1997, April). Risk, financial markets and human capital in a developing country. *Review of Economic Studies* 64(2), 311–336.
- Kehoe, P. and F. Perri (1998). International business cycles with endogenous incomplete markets. Manuscript.
- Kimball, M. S. (1988). Farmers' cooperatives as behavior toward risk. *American Economic Review* 78, 224–236.
- Kletzer, K. M. and B. D. Wright (2000). Sovereign debt as intertemporal barter. *American Economic Review* 90(3), 621–639.
- Kochar, A. (1995, May). Explaining household vulnerability to idiosyncratic income shocks. *American Economic Review* 85(2), 159–164.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies* 63(4), 595–610.
- Ligon, E. (1998). Risk-sharing and information in village economies. *Review of Economic Studies* 65, 847–864.
- Ligon, E., J. P. Thomas, and T. Worrall (2000). Mutual insurance, individual savings, and limited commitment. *Review of Economic Dynamics* 3(2), 216–246.
- Lund, S. and M. Fafchamps (1997). Risk-sharing networks in rural Philippines. mimeograph.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99, 928–956.
- McCloskey, D. N. (1976). English open fields as behavior towards risk. In P. Uselding (Ed.), *Research in Economic History*, Volume 1, pp. 124–170. Greenwich, Connecticut: JAI Press.
- Morduch, J. (1995, Summer). Income smoothing and consumption smoothing. *Journal of Economic Perspectives* 9(3), 103–114.
- Pender, J. L. (1996). Discount rates and credit markets: Theory and evidence from rural India. *Journal of Development Economics* 50(2), 257–296.
- Platteau, J.-P. (1997, August). Mutual insurance as an elusive concept in traditional rural societies. *Journal of Development Studies* 33(6), 764–796.
- Platteau, J.-P. and A. Abraham (1987). An inquiry into quasi-credit contracts: The role of reciprocal credit and interlinked deals in small-scale fishing communities. *Journal of Development Studies* 23(4), 461–490.
- Posner, R. (1980, April). A theory of primitive society, with special reference to law. *Journal of Law and Economics* 23, 1–53.
- Ravallion, M. and S. Chaudhuri (1997). Risk and insurance in village India: Comment. *Econometrica* 65(1), 171–184.
- Rosenzweig, M. and O. Stark (1989). Consumption smoothing, migration and marriage: Evidence from rural India. *Journal of Political Economy* 97(4), 905–927.
- Rosenzweig, M. R. and K. I. Wolpin (1993). Credit market constraints, consumption smoothing, and the accumulation of durable assets in low-income countries: Investments in bullocks in India. *Journal of Political Economy* 101, 223–244.
- Sahlins, M. (1974). *Stone Age Economics*. London: Tavistock.
- Thomas, J. P. and T. Worrall (1988). Self-enforcing wage contracts. *Review of Economic Studies* 55, 541–554.
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.
- Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria. *Review of Economic Studies* 63, 495–526.
- White, H. (1984). *Asymptotic Theory for Econometricians*. Orlando: Academic Press.

NOTES

¹A similar procedure was first adopted by Mace (1991) and Cochrane (1991) who studied U.S. consumption data.

²Ravallion and Chaudhuri (1997) re-examine Townsend's results using an alternative procedure to measure consumption but find a more definite rejection of full risk-pooling. Deaton, and Lund and Fafchamps, argue that the rejection cannot be accounted for by the ad hoc introduction of borrowing constraints.

³Mutual insurance is only one method of pooling risk. Townsend (1994) identifies many possibilities. These may be usefully divided into income smoothing and consumption smoothing strategies (see Morduch (1995)). Income smoothing might include plot diversification (McCloskey (1976)), planting lower yielding but hardier varieties (Morduch (1995)), delaying planting until more accurate weather forecasts are available (Bliss and Stern (1982)), having a family member working in fixed wage employment or outside the village (Rosenzweig and Stark (1989)) or intertemporal labour substitution (Kochar (1995) and Jacoby and Skoufias (1997)). Consumption smoothing may take the form of mutual insurance through gifts and loans as we will stress in the next two sections, but might also include intertemporal transfers or asset accumulation (e.g. the sale and purchase of animals, Rosenzweig and Wolpin (1993).) We will examine the extension to allow for intertemporal transfers in Section 4.

⁴That gifts can play a role as insurance is recognised in the anthropological literature (see, e.g., Cashdan (1985)).

⁵Platteau (1997) goes further and argues that mutual insurance is a concept alien to those in traditional agrarian societies and that informal insurance arrangements in these societies are guided by the principle of balanced reciprocity. Balanced reciprocity (see also Sahlins) means that for any 'gift' there is a strong assumption that at some, as yet unknown, time in the future there will be a 'counter-gift'. Or as Platteau puts it, for any payment there must be a "tangible *quid pro quo*."

⁶Direct penalties might include peer group pressure or being brought before a village council for admonishment.

⁷Examples include applications to sovereign debt by Kletzer and Wright (2000) and Atkeson (1991) (who also considers asymmetric information), Chari and Kehoe (1993) who consider a model in which both the government and private citizens can default on their debt to each other, a two-country international business cycle model with both production and capital accumulation by Kehoe and Perri (1998), a two-party model of political bargaining by Dixit, Grossman, and Gul (1998) and a model of asset pricing by Alvarez and Jermann (2000).

⁸Alvarez and Jermann (2000) obtain a similar characterisation to that here in a two-sided risk aversion case but assume symmetry, no aggregate uncertainty and a monotonicity property.

⁹Holmström's back-loading principle (Holmström (1983)) established in a two-period model can be summarised as saying that in order to deter an agent from renegeing, their rewards should be back-loaded into the future.

¹⁰Nothing important depends on this assumption that all transition probabilities are positive. Income is assumed to be positive so marginal utility is bounded at autarky.

¹¹The assumption of an infinite horizon can be justified by appealing to the continuity of households through their offspring. In fact all that is needed is the belief that the insurance game defined below will continue to be played with some positive probability, this probability being reflected in the discount rate that the households use. See Coate and Ravallion (1993) for more discussion of the

dynastic interpretation of this assumption in the rural village context. The fact that we allow for exogenous penalties consequent upon contract violation also implies that in a finite horizon model backwards unravelling does not occur, and we conjecture that our results would be approximately valid if this time horizon were sufficiently long.

¹²Hence this assumption allows us to characterise the most efficient non-cooperative (subgame perfect) equilibria (see also footnote 22 below for further discussion of this point). If the reversion to autarky assumption seems too extreme, then replacing it with the assumption of an eventual return to risk-sharing will not substantially change the contract characterisation that we obtain. This will increase the utility from renegeing, changing the right-hand sides of the incentive constraints (2) and (3) below. In the case of *i.i.d.* shocks each period, with say an n -period exclusion from risk-sharing, and some fixed division of the gains from risk-sharing thereafter, this will simply add a constant, and our general characterisation is unchanged.

¹³Obviously if $P_i(s)$ were large enough, there would be no enforceability problems and full insurance would be possible. Equally Proposition 2(iv) below shows that if $P_i(s) = 0$ for each state and each household and if the discount factor δ is small enough then only autarkic consumptions will be feasible. We shall be mainly interested in intermediate cases where some but not full risk-sharing is possible.

¹⁴Since income is non-storable and consumption opportunities are limited to joint income in each period we have implicitly ruled out credit market transactions with parties outside the village. Different assumptions about access to credit markets might make a substantial difference to the results. For example in the absence of direct penalties, the possibility of saving in a ‘cash-in-advance’ account which offers an average return of $(1/\delta) - 1$, if this can be made state contingent in a suitable fashion, will undo any sustainable risk-sharing contract (see Bulow and Rogoff (1989)). Nevertheless, we do not consider this type of credit transaction to be realistic in most rural village contexts.

¹⁵For period 1, h_{t-1} is the empty set.

¹⁶Thomas and Worrall (1988) analysed a long-term wage contract between a risk-averse worker and a risk-neutral firm in which the worker can at any date quit the firm and work at the random (*i.i.d.*) spot-market wage. This would be formally equivalent in the current context to assuming that one of the households is risk-neutral and has no non-negativity constraint on consumption and that there are no direct penalties.

¹⁷Technical details of the dynamic programming procedure, *e.g.*, that these intervals are closed, can be found in Thomas and Worrall (1988) and the same proofs carry over *mutatis mutandis* to the current context.

¹⁸Clearly \underline{U}_s cannot be smaller than either term in the max operator; if \underline{U}_s is strictly larger than both, then it is possible to cut either household 1’s current consumption or one of its future surpluses without violating the sustainability constraints.

¹⁹The actual contract starts at date 1, but, as argued above, continuation contracts must be efficient.

²⁰The objective function and constraints are easily seen to be concave and the Slater condition is satisfied whenever the constraint set is more than a singleton.

²¹That is, there is a unique solution for τ_s to equation (11) given a value for λ and taking into account the complementary slackness conditions on the non-negative consumption constraints. Hence either there is a unique interior solution with the ratio of the marginal utilities equal to λ , or λ lies outside the set of marginal utility ratios which can be generated by feasible transfers in state

s , namely $[v'(y_1(s) + y_2(s))/u'(0), v'(0)/u'(y_1(s) + y_2(s))]$ in which case there is a corner solution with all income going to one of the households.

²²As stated above, provided that there are no penalties other than the return to autarky for breach of contract, there is a one-to-one relationship between our sustainable contracts and subgame perfect equilibria. The constrained-efficient contracts which we characterise then correspond precisely to the Pareto frontier of the equilibrium payoff set. The Pareto frontier can also be shown to be renegotiation proof in the sense that a contract can be devised for each point on the frontier which involves continuation payoffs lying exclusively on the frontier; the idea is to replace the return to autarky punishment by the point on the Pareto frontier for the current state which gives the lowest surplus to the deviant household as defined by (4) or (5). The other household will not agree to a renegotiation of this equilibrium since it is receiving its maximum surplus. This corresponds to the weak renegotiation proof concept of Farrell and Maskin (1989). Renegotiation proofness (including stronger concepts) for models very close to that of Thomas and Worrall (1988) has been established in Asheim and Strand (1991) and in Kletzer and Wright (2000), and a similar argument is applicable here.

²³Alvarez and Jermann (2000) relax the assumption that all transition probabilities are positive.

²⁴This is our procedure for calculating the static limited commitment model. In Coate and Ravallion (1993) the direct penalties are zero and the discount factor positive, but this alternative procedure gives equivalent results since the state-dependent direct penalty can be used to replicate the promised utilities which are history independent in Coate-Ravallion.

²⁵The long-run value of λ will be at the bottom of the common intersection of all intervals if λ_0 lies below the intersection (the long-run value is attained as soon as the state with the highest $\underline{\lambda}_s$ occurs), and at the top if initially it lies above; if the initial value of λ belongs to the common intersection then λ will remain constant and the contract will be first best. For some distributions of the potential surplus from the relationship the contract will not be a first-best allocation; nevertheless if (and only if) some first-best allocation is sustainable, the contract must end up (with probability one) having a first-best continuation contract.

²⁶A proof is available upon request. The logarithmic utility function also implies that all the calculations presented below are independent of the absolute size of income levels, and only depend on the percentage loss in the bad state.

²⁷With the assumptions $p = 0.1$ and $d/y_h = 0.5$.

²⁸The case where the intervals are non-overlapping is very similar.

²⁹If both households simultaneously receive bad shocks then the repayments continue, except they are reduced for that period, proportionately to the fall in aggregate income (50%).

³⁰For example, contrast the treatment of storage in Gobert and Poitevin (1998) in which defaulting households forfeit stored assets with the treatment in Ligon, Thomas, and Worrall (2000), where holding a large store of assets can provide an incentive for a household to renege on existing arrangements. It's important to note that the latter treatment (which is similar to our approach here) may introduce nonconvexities, since the value associated with autarky now depends on choice variables. We ignore this possible difficulty here, but refer the reader to Ligon, Thomas, and Worrall (2000) for a more satisfactory treatment.

³¹Save for timing, this mechanism is similar to one considered by Gauthier, Poitevin, and González (1997) which they term "ex ante payments."

³²In fact, households employ labour in production. However, by assuming that labour and other input decisions are efficient, and that utility from leisure is additively separable from utility from consumption, we can abstract from production with no further loss of generality.

³³In Aurepalle the ratio of aggregate income to aggregate consumption over the period 1976–81 is 2.31, 2.89, 2.47, 1.97, 1.61, and 1.93, while in Shirapur the corresponding series is 1.18 1.39 1.48, 1.72, 1.20, 1.49, and in Kanzara is 1.87, 2.24, 2.17, 1.64, 1.60, and 2.12. It's not clear what accounts for this apparently large surplus; either there's mismeasurement, or a great deal of aggregate savings. Some evidence in favor of the latter point of view is that although these villages are sometimes subject to severe drought, no such drought occurred during the sample period.

³⁴We have used this approach to approximate the solution to a simple example with many (500) households. In each period households draw an endowment from an identical two point income distribution and have log utility. This allowed us to compute sample consumption paths for each household, using only the last 100 periods of a 500 period simulation in order to approximate a steady state. We then compare predicted consumptions from this approximate solution to the known steady-state solution in the continuum case. Solutions to the problem with a continuum of households have the property that each constrained household with high income has the same consumption, each constrained household with low income has consumption equal to income, and all unconstrained households have the same (positive) growth rate in marginal utility. As a consequence, the solution is determined by two numbers, the consumption of the high income constrained and the growth rate of marginal utility for unconstrained households.

We expect that the correct solution to the 500 household example will be quite close to the continuum case. Thus, using the “rest of village” algorithm outlined here gives us an approximation to the 500 household case, which in turn is an approximation to the continuum case. Comparing our approximate solution to what would obtain in the continuum case may thus give us some sense of how important approximation error is in our empirical work.

The comparison is revealing. This particular example points out a shortcoming of our algorithm, which does not correctly reproduce the feature of the continuum solution that all households with high-income will have the same level of consumption; rather, a household with repeated high realizations may (incorrectly) appear to be unconstrained after the first high realization, so that the marginal utility of this household will increase at the same rate as all other unconstrained households. This problem is due to the fact that our approximation doesn't properly take into account the fall in consumption associated with being unconstrained in an economy with many households.

The example we've described is designed to highlight shortcomings of our approximation, and does a nice job of illustrating the qualitative difference we've just described. Nonetheless, the average value of the correlation coefficient between the two consumption paths is 0.972, which suggests that the approximation is fairly good, at least in this example. (Details available on request.)

³⁵This simplification of the problem is due to Fumio Hayashi (1996).

³⁶The reason we have chosen a small number of cells is due more to the paucity of data we have for any given household's income realizations than it is to the computational expense; it's not difficult to quickly solve models with dozens or even hundreds of different states.

³⁷If one were to assume that disturbances in (22) were i.i.d. normal, or alternatively if one's sample was sufficiently large, then one could interpret the criterion statistics in Table 1 as log-likelihood statistics. Because of our nested specification, the difference between the log-likelihoods of the different models would then have the interpretation of a likelihood ratio, and under some

simple regularity conditions twice this ratio would be distributed χ^2 , with degrees of freedom for each test equal to the differences in the number of parameters estimated. Here we do not do so, for several reasons. First, one of our estimates (the estimate of $\gamma = 0.01$ in Aurepalle for the full insurance model) lies on the edge of our parameter space, violating one of the usual regularity conditions. Second, the assumption of *i.i.d.* normal disturbances seems very strong, and our sense is that our sample is too small to rely on the usual asymptotic arguments which might replace this distributional assumption. Third, since some of the parameters we estimate aren't identified in all of the models (e.g., δ when there's full insurance), the appropriate limiting distribution won't actually be χ^2 (but note that the methods of Hansen (1992) could be used to compute the correct distribution). Nonetheless, in a larger sample one could compare the fit of these models more formally.

³⁸Since, for CRRA preferences the full insurance model simply predicts that consumption shares ought to remain constant over time, it may seem that γ wouldn't be identified by the full insurance model. We actually test a slightly weaker restriction than constant shares. In particular, after solving for the initial λ -weights in terms of time zero consumption, we actually use the restriction

$$\sigma_t^i = \sigma_0^i \left(\frac{\sum_{j=1}^N (\sigma_t^j)^\gamma}{\sum_{j=1}^N (\sigma_0^j)^\gamma} \right)^{1/\gamma}.$$

Of course, if full insurance prevails throughout the village and there's no measurement error in consumption, then the last factor in this expression disappears, and we recover the prediction that consumption shares ought to be constant. This weaker form of the restriction allows us to consider the case in which some (but not all) households are fully insured. Further, this equation makes some allowance for time zero measurement error, by using information from consumption in all periods to estimate shares.

³⁹It's possible to compute estimates of average relative risk aversion from the consumption-smoothing regressions of Townsend (1994) (see his Tables IX and A.I) under the null hypothesis of full insurance. These estimates are similarly wild and mostly negative, ranging from -32 in Shirapur to 2.8 in Aurepalle.

⁴⁰If households are *not* myopic, then the parameter P in the static limited commitment model may be construed as the sum of punishments for default plus the utility cost of exclusion from future risk-pooling, since (unlike the dynamic limited commitment model) this latter quantity is simply a number, not a function.

⁴¹A warning here seems in order. Although we hope that the correlations we report are instructive, the reader should bear in mind that the correlations measure only linear relationships between different consumptions; since we're interested in a highly nonlinear relationship, the correlation coefficient may be quite misleading when regarded as a measure of fit.

UNIVERSITY OF CALIFORNIA, BERKELEY AND GIANNINI FOUNDATION

UNIVERSITY OF ST. ANDREWS

KEELE UNIVERSITY