

INFORMAL INSURANCE ARRANGEMENTS IN VILLAGE ECONOMIES

ETHAN LIGON, JONATHAN P. THOMAS, AND TIM WORRALL

ABSTRACT. This paper studies efficient insurance arrangements in village economies when there is complete information but limited commitment. Commitment is limited because only limited penalties can be imposed on households which renege on their promises. Any efficient insurance arrangement must therefore take into account the fact that households will renege if the benefits from doing so outweigh the costs. We study a general model which admits aggregate and idiosyncratic risk as well as serial correlation of incomes. It is shown that in the case of two households and no storage the efficient insurance arrangement is characterized by a simple updating rule. An example illustrates the similarity of the efficient arrangement to a simple debt contract with occasional debt forgiveness. The model is then extended to multiple households and a simple storage technology. We use data from the ICRISAT survey of three villages in southern India to test the theory against three alternative models: autarky, full insurance, and a static model of limited commitment due to Coate and Ravallion (1993). Overall, the model we develop does a significantly better job of explaining the data than does any of these alternatives.

1. INTRODUCTION

A village economy is a closed, cohesive, agrarian economy consisting of a group of mainly subsistence household-farmers. Probably the most important fact of village life is the risk to crop yields caused by climatic conditions (poor rainfall and wind damage), flooding, crop disease, insect infestation and the farmer's health. Some of these risks will be

Date: February 1997.

We are very grateful for comments and correspondence with Chris Udry and Stephen Coate at an early stage of the paper and to Andrew Foster and Narayana Kocherlakota for comments on the current draft. Thanks are also due to the many useful comments from participants at seminars given at Brown, Berkeley, Exeter, Liverpool, the London School of Economics, Manchester, the Federal Reserve Bank of Minneapolis and University College London, and the Econometric Society meetings in Prague and Tokyo. Sections 2 to 4 of the present paper are drawn from an earlier paper of the same title, written by the second and third authors (dated November 1994).

common to the village (like general rainfall) but some will be localized or idiosyncratic affecting only one (like illness) or a small sub-group (like hailstone damage) of farmers. Furthermore, since climatic conditions are positively serially correlated over time, crop yields are likely to be positively serially correlated too.

As risk is prevalent and critical to households near to subsistence, much effort will be expended on risk mitigation. In the absence of perfect capital markets, one response in traditional agrarian communities is the development of informal insurance arrangements.¹ The importance of such arrangements has been widely recognized by social anthropologists, sociologists and economists.² Informal insurance arrangements are essentially a form of mutual insurance which provide for those in need based on an understanding of reciprocity. Platteau (1996) argues that mutual insurance is a concept alien to those in traditional agrarian societies and that such arrangements although providing insurance are more importantly guided by a principle of balanced reciprocity. Balanced reciprocity 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.' That is for any payment there is a "tangible *quid pro quo*". Thus informal insurance arrangements might appear similar to credit or quasi-credit arrangements.³ These insurance arrangements are informal because there is no formal legal framework within traditional agrarian societies to make binding commitments or enforce promises of reciprocity.⁴

The purpose of this paper is to develop a testable theory of informal insurance arrangements in village economies and to confront it with data from three villages in southern India from the Village Level Studies

¹Townsend (1994) provides a taxonomy of risk mitigating strategies. Risk may be mitigated through the adoption of less risky technologies or crop diversification (McCloskey 1976) or through storage (Deaton 1992) or asset accumulation (Rosenzweig and Wolpin 1993). These are *ex ante* measures taken before realizations are known. *Ex post* measures include gifts and transfers, credit market transactions and migration. Informal insurance arrangements fall in the *ex post* category, arrangements being made *ex ante* but exchanges taking place *ex post*.

²See, e.g. Fafchamps (1992), Nash (1966), Platteau and Abraham (1987).

³See Platteau and Abraham (1987) and Fafchamps (1995). Udry (1990, 1994) provides evidence from northern Nigeria that repayments on loans are state-contingent. On average a borrower with a good realization repays 20.4% more than he borrowed but a borrower with a bad realization repays 0.6% less than he borrowed. Moreover, repayments are contingent on the lender's realization. A lender with a good realization receives on average 5% less than he lent, but a lender with a bad realization receives 11.8% more than he lent.

⁴Udry (1994) documents that loans are made without witnesses or written record and in only 3% of cases were loans backed by collateral.

survey of the International Crops Research Institute of the Semi-Arid Tropics over the period 1975–1984.

In Section 2 we consider a general model where yields within the village follow 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, risk averse, and consume a single consumption good which may be stored using a non-random storage technology and may enter into a long-term insurance contract with other households. Given the absence of a formal legal framework, any insurance contract is assumed to be sustained by the joint means of direct penalties against breach⁵ and also the threat of future exclusion from insurance possibilities. We characterize the constrained-efficient insurance arrangements when commitment is limited.

In this approach we follow Kimball (1988), Coate and Ravallion (1993) and Kocherlakota (1996). Kimball (1988) examines only whether such schemes might exist⁶ and Coate and Ravallion (1993) consider only two-household, symmetric environments with a restriction to stationary arrangements which turn out not to be optimal when the first-best is not attainable. Kocherlakota (1996) examines the multi-household case but only for symmetric non-autocorrelated income processes, no saving, and no direct penalties. His main concern is also the long run properties of the contract and the efficient contract is not completely characterized. His analysis is therefore less amenable to empirical investigation for two reasons. First, the VLS data on incomes is non-symmetric and savings are non-trivial. Secondly and more importantly, we derive a simple, complete characterization of the constrained-efficient contract in the no-storage, bilateral exchange case which we use extensively to predict consumptions in the model, which we then compare with the actual data. The direct penalties also play an important role in our estimation procedure.

The basic characterization of the bilateral exchange case given in Section 3 can be easily and briefly summarized. Each state of nature determines the income of each household. Absent storage, the constrained-efficient insurance arrangement specifies an amount to be transferred from one household to the other. Because households are risk averse, the ratio of the marginal utilities of the two households is

⁵Direct penalties might include peer group pressure such as embarrassment in front of one's father or another village authority figure, or more colorfully as something along the lines of having the mob break your leg (Udry 1994).

⁶Fafchamps (1992) also discusses their possible existence and emphasizes the game theoretic underpinnings of how future exclusion may encourage repayments even in the absence of legal sanctions.

monotonic in the size of the transfer. Each state of nature is associated with a certain interval of marginal utility ratios, and the insurance arrangement satisfies a simple updating rule. Given the previous period's marginal utility ratio, and the current state, the new ratio lies within the interval associated with the current state, such that the change in the ratio is minimized. This implies that the marginal utility ratio is kept constant whenever possible. This is very intuitive since a constant ratio would of course be the outcome of a first-best insurance arrangement. If the first-best is not attainable, however, then the ratio must change at some point, with some household constrained (that is, just indifferent between adhering to the contract and reneging). In this case the updating rule specifies that the change is as small as possible given that the new interval must be attained. This simple rule allows the entire insurance arrangement to be determined.

In Section 4 we analyze an example which allows us to interpret the insurance arrangement in terms of credit or quasi-credit. Credit contracts do have a desirable property when commitment is limited. They offer a future reward to a household with a good realization which is being asked to sacrifice current consumption to help insure a household with a bad current realization. At the same time, however, they create incentive problems for households which have to repay loans previously taken out. In this example we show how the constrained efficient contract can be interpreted as a form of debt contract with debt forgiveness in certain states of nature. The possibility of forgiveness mitigates the debtor's incentive problem when they in turn are required to be insurance provider. Although specific to our two household example, this idea of forgetfulness generalizes to more complex environments with several households. If a household is constrained at some date, then the future course of the contract depends on the state at that date and not the previous history. Section 5 extends the model to allow for any finite number of households and allow for a simple storage technology.

The second part of the paper considers a test of the theory. Using the VLS data we test the theory by predicting consumption allocations from the model and comparing the predictions with the actual data. For simplicity we abstract from savings and solve the bilateral exchange model for each household against the rest of the sample. This is done for a particular set of parameter values: the discount factor, the coefficient of relative risk aversion and a state and household independent default penalty. Using the actual data on consumption in the first period the vector of marginal utilities can be determined. Income data is then used to determine the evolution of states over time so that

predicted consumptions can be generated for each household. Assuming normally distributed disturbances, we estimate the parameters of the model using a maximum likelihood estimator. This procedure is described in Section 6. A similar procedure is adopted under the alternative assumptions of full insurance, no insurance and a static limited commitment variant of the Coate-Ravallion model. Several of these alternatives are nested within our own so that a likelihood ratio test can be used to make comparisons.

Results are reported in Section 7. The models of limited commitment do better than the full insurance model in each of the three villages. The dynamic limited commitment model outperforms the static model in two out of the three villages but produces an unreasonably low estimate of the discount factor in one of these villages. At a more informal level a similar ranking is obtained by examining simple correlation coefficients between the consumptions predicted by the models and actual consumption. Although the dynamic limited commitment model performs well it clearly does not explain everything. We therefore regress actual consumption on actual household income, aggregate consumption and predicted consumption. Adding fixed effects suggests that aggregate consumption is unimportant but that individual income still has explanatory power. The dynamic limited commitment model predicts too much insurance.

The VLS data has also been used by Townsend (1994) to test the full insurance hypothesis. He adopts a similar procedure to that of Mace (1991) and Cochrane (1991) who studied U.S. consumption data. Townsend (1994) regresses household consumption on aggregate consumption and a vector of other variables including household income. Under full insurance these other variables should not enter as significant variables in the regression. Although the null of full insurance is rejected it does not perform too badly and can be considered a benchmark case. A very similar conclusion is drawn by Udry (1994) in his study of a district in northern Nigeria and our results here are also consistent in suggesting that full insurance performs reasonably well. Although supportive of the full insurance hypothesis these results suggest that there may be significant reasons why consumption allocations do not fully replicate the Pareto-efficient risk-pooling outcomes of a complete set of competitive state-contingent markets. Potentially the lack of full risk-pooling may be due to either problems of limited information or problems of limited commitment or both. Although we

concentrate here on the limited commitment problem,⁷ the importance of limited information also needs to be assessed⁸ and future research might consider both aspects simultaneously.⁹

As far as we know there are only a few studies which have attempted to explicitly test the limited commitment model. Foster and Rosenzweig (1995) provide a test of the limited commitment also using the VLS data as well as other data from India and Pakistan. They extend the model of this paper to allow for altruistic links between households. Their test is based on an implied negative relationship between the current transfer and an aggregate of previous transfers. They provide evidence that limited commitment substantially constrains informal transfer arrangements and further show that altruism also plays an important role in ameliorating sustainability constraints. Beaudry and DiNardo (1995) also provide a test of a limited commitment model but in a different context. They consider a market in implicit labor contracts. Their test is 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 hourly wage should be associated with a fall in hours if leisure is a normal good. Again their results are supportive of the theory.

The paper is organized as follows. The basic bilateral model is presented in Section 2 and the constrained-efficient contract characterized in Section 3. Section 4 presents a simple example and Section 5 extends the model to accommodate multiple households and storage. The empirical analysis is carried out in Sections 6 and 7.

2. THE 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 $s \in S$, and $S = \{1, 2, \dots, S\}$. It is assumed that the state of nature follows a Markov process with the probability of transition from state s to state

⁷Udry (1990) argues that informational asymmetries do not play an important role in loan transactions in northern Nigeria. He argues that village economies tend to be cohesive with a ready flow of information within the village.

⁸This is done in Ligon (1996b) who uses the same data set to conclude that private information also plays an important role in these villages.

⁹In a wider context there may be a trade-off between trading risk with nearby villages which may suffer different shocks but where information problems between villages might be more important and trading risk within the village.

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 given by π_r^0 . This formalization 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.¹¹

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 maximizers. The 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.¹²

Because of the risk aversion of at least one of the households, the two households will generally have an incentive to share risk. We assume that the households enter into a 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

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.

¹¹As income is non-storable we are also implicitly ruling out outside credit market transactions so that in each period consumption opportunities are limited to joint income. In Section 5 we suppose that households have access to a simple storage technology. Different assumptions about access to credit markets might make a substantial difference to the results. Absent any direct penalties, for example, 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.

¹²The 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.

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 sustainable contract correspond to a subgame-perfect equilibrium, but also there can be no other equilibrium outcomes (see Abreu 1988). Hence this assumption allows us to characterize the most efficient non-cooperative (subgame perfect) equilibria (see also footnote 22 below for further discussion of this point). If reversion to autarky seems too extreme an assumption, then replacing it with the assumption of an eventual return to risk-sharing will not substantially change the contract characterization that we obtain.¹³ Secondly, it is assumed that contract breaches 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, as discussed in the introduction. 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 . Given $0 \leq \delta < 1$ and the finite gains from risk sharing, it is obvious that 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 interested in intermediate cases where some but not full risk-sharing is possible.

Let s_t be the state of the world occurring at date t . A contract 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 $h_t = (h_{t-1}, s_t)$

¹³Formally this will increase the utility from reneging, 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 characterization is unchanged.

occurs up to period t (i.e. when the current state s_t is known).¹⁴

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

where 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 characterize the *constrained-efficient* contracts, those which are not Pareto-dominated by any other sustainable contract.

3. CHARACTERIZATION OF CONSTRAINED-EFFICIENT CONTRACTS

To solve for the (constrained) efficient set of sustainable contracts a straightforward dynamic programming procedure can be followed. This 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 why all continuation contracts should be efficient 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—such a move will make the overall contract Pareto superior to the original one. This dynamic programming problem is similar in structure to that analyzed by Thomas and Worrall (1988). Necessary technical details have been established there, and the same proofs carry over *mutatis mutandis* to the current context.¹⁵

¹⁴For period 1, h_{t-1} is the empty set.

¹⁵Thomas and Worrall (1988) analyzed 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. Kletzer and Wright

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 characterize 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 $[\underline{V}_s, \bar{V}_s]$.¹⁶ 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 seen¹⁷ that the \underline{U}_s must be the (unique) solutions to

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

(1996) look at the same model in a sovereign debt market; they derive a simpler proof for the main characterization of Thomas and Worrall.

¹⁶The proof that the intervals are closed is as in Thomas and Worrall (1988).

¹⁷Clearly \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.

and the \underline{V}_s solve

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

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 maximizing, 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 .¹⁸ 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

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

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

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

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

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

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,

¹⁸The actual contract starts at date 1, but, as argued above, continuation contracts must be efficient.

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.¹⁹

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

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

together with the envelope condition

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

A constrained-efficient contract can be characterized 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.²⁰ 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 characterized as follows: There exist S state dependent intervals $[\underline{\lambda}_r, \bar{\lambda}_r]$, $r = 1, 2, \dots, S$,*

¹⁹The 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.

²⁰That 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.

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

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

This completely characterizes 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.²¹ 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,

²¹ This resembles the characterization found in Thomas and Worrall (1988) where the contract *wage* is held constant where possible.

household 2 getting 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.

Given the rule of Proposition 1, we can think of an initial value of λ , which we denote λ_0 , as determining the entire contract. As λ_0 varies from $\min_s \{\underline{\lambda}_s\}$ to $\max_s \{\bar{\lambda}_s\}$, 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

It is possible to say more about the λ -intervals; this is done in the next proposition.

- Proposition 2.** (i) *There exists a critical δ^* , $0 \leq \delta^* < 1$, such that the intervals have non-empty intersection if and only if $\delta \geq \delta^*$;*
- (ii) *Assume that $P_i(s) = P_i$ for all s and $i = 1, 2$. Then for each state $s \in S$, the interval $[\underline{\lambda}_s, \bar{\lambda}_s]$ contains the autarkic marginal utility ratio $v'(y_2(s))/u'(y_1(s))$;*
- (iii) *Assume that $P_i(s) = 0$ for all s and $i = 1, 2$. Then $\min\{\underline{\lambda}_s\} = \min\{v'(y_2(s))/u'(y_1(s))\}$ and $\max\{\bar{\lambda}_s\} = \max\{v'(y_2(s))/u'(y_1(s))\}$;*

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 characterize 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 (1996), and a similar argument is applicable here.

- (iv) Assume that $P_i(s) = 0$ for all s and $i = 1, 2$. Then there exists a critical $0 < \delta^{**} < 1$ such that there is no non-autarkic contract for $0 \leq \delta < \delta^{**}$.

Proof. See Appendix. □

Part (i) of Proposition 2 does not imply that *any* full insurance allocation is sustainable (remember this is implied if the penalties are large enough) but that for a large enough discount factor there is *some* first-best, full insurance contract which is sustainable. As all λ -intervals overlap, there is a λ which simultaneously belongs to each interval. Hence once a state occurs such that λ belongs to the common intersection, it remains constant thereafter; the contract therefore converges (with probability one) to a first-best contract. (The long-run value of λ will be at the bottom of the common intersection of all intervals if λ_0 lies below the intersection 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 is not a full insurance allocation; nevertheless if (and only if) some full insurance allocation is sustainable, the contract must end up (with probability one) having a first-best continuation contract.

Parts (ii) and (iii) of Proposition 2 relate the λ -interval to the autarkic marginal utility ratio. Specifically, when the penalties are state independent, each interval will enclose the autarkic marginal utility ratio for that state, and if the penalties are all zero then the lowest (highest) point of all the intervals will be the lowest (highest) autarkic marginal utility ratio. To see this, suppose the autarkic marginal utility ratio in state s , lay above the λ_s -interval. Then the contract will always call for a transfer from household 1 to household 2 in state s no matter what the previous history. Household 2 therefore receives a positive 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 contracts must be worse than the current penalty. 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 must be making a net transfer. Similarly $\min\{v'(y_2(s))/u'(y_1(s))\}$ is relatively the worst state for household 1. If there are no penalties there cannot be a sustainable contract which calls for household 1 to make a transfer in this state.

With no penalties and with $\delta = 0$, it is clear that the only sustainable contract is autarkic. Part (iv) of Proposition 2 shows that even

if the households do not discount the future completely but nevertheless sufficiently heavily, then the only sustainable contract is autarkic. Basically since the gains from risk-sharing are finite for some small discount factor the future gains from risk sharing can never offset the short run loss of marginal utility from a current transfer.

4. AN EXAMPLE

In this section we consider a special case of the general model presented in Section 2. This special case is the problem of mutual insurance of households where each of the two households may have either a high or a low income. We shall derive the constrained efficient limited commitment (LC) contract and compare it with the static constrained-efficient contract studied by Coate and Ravallion (1993) and with a one period debt contract with occasional debt forgiveness.

In the example each household has an income of y_h with probability $(1 - p)$ and income y_l with probability p , $0 < p < 1$. Thus we may consider this as the situation in which in each period a household may suffer a loss of $d = y_h - y_l$ with probability p . The probability p is the same, but independent, for each household and constant over time. Hence the expected income of each household is $y_h - pd$ in each period. In the insurance context it is natural to think of p as being a small probability and as d as a relatively large loss. In the calculations we present below we consider the case where both households have a 10% chance of a 50% loss ($p = 0.1$ and $d/y_h = 0.5$).²³ 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 $v'(y_h)/u'(y_h) = v'(y_l)/u'(y_l) = 1$. 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 that preferences can be represented by the utility function $u(c) = v(c) = \log_e(c)$.²⁴ The main advantage of the logarithmic utility function for computing the example is that the λ intervals for

²³Such drastic income fluctuations are not uncommon in the Indian village data which we examine below: of the 104 households which were sampled continuously over a nine year period, 32 experienced at least one year in which income was less than 50 percent of the median year's income (Walker and Ryan 1990).

²⁴Logarithmic utility is a special case of preferences exhibiting constant relative risk aversion. We use this more general class in the empirical work of later sections.

the states ll and hh defined in Proposition 1 coincide.²⁵ 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 which we have assumed to be 50%. In addition we shall assume that penalties are either zero, or state independent.

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; each is illustrated in Figure 1. The figure is drawn for the case of no penalties, so that in each case $\underline{\lambda}_{lh}$ and $\bar{\lambda}_{hl}$ equals the autarky ratio of marginal utilities as stated in Proposition 2(iii). 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 2 we plot the logarithm of the interval endpoints against the discount factor; 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. As δ converges to one, the (logarithm of the) common intersection of the intervals converges to $[-0.0717, 0.0717]$, which corresponds precisely to the (logarithm of the) set of marginal utility ratios of the first-best insurance arrangements which give a non-negative *average* surplus overall, that is to say which are ex ante individually rational (in accordance with the 'folk theorem' for repeated games).

Case 1 where all the intervals overlap has been discussed above. No matter what the initial distribution of the surplus the constrained efficient contract ends up, with probability one, with λ between $\underline{\lambda}_{hl} < 1$ and $\bar{\lambda}_{lh} > 1$. In particular, in the long term full risk pooling results. Perhaps the more interesting cases are 2 and 3.

When no first-best contract is sustainable, the constrained-efficient contract is easy to interpret. Consider Case 2 where the first-best is not attainable, but the lh and hl intervals overlap the hh and ll intervals. Suppose that household 1 is the first to receive a bad shock; λ falls to $\bar{\lambda}_{lh}$, where $1 > \bar{\lambda}_{lh} > v'(y_h)/u'(y_l) = 1/2$, and household

²⁵A proof is available upon request.

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. At this point 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.²⁶ If 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%).

In Case 3 this story is essentially the same; the only difference is that the ratio of marginal utilities differs between the borrowing state and the repayment states (thus if state lh is followed by hh or ll , then λ rises from $\bar{\lambda}_{lh}$ to $\underline{\lambda}_{hh}$, but this still involves 'repayments' by household 1).²⁷ In either case the promise of future repayments induces the household with a good realisation to lend more to a household with a bad realisation than would be the case if no such repayments were anticipated, as under the static contract characterised by Coate and Ravallion (1993). The drawback to such repayments is that while they achieve significant insurance at a particular date, it is at the cost of variable consumption over time, as the level of consumption will be higher when a household is in a 'creditor' position than in a 'debtor' position in the symmetric states. The problem with a more conventional debt contract (or sequence of debt contracts) is that it 'remembers' all previous loans: if a household which already has built up debt is supposed to lend to the other household when the latter has a bad shock, then it will not anticipate future repayments if its overall debt is still positive, and so

²⁶This idea of "forgetfulness" generalises to more complex environments as follows: if a household is constrained in a particular state at some date then the future course of the contract depends only on the state and not the previous history.

²⁷Indeed, this general interpretation of the contract is not dependent upon the logarithmic utility function assumed in this section; the only additional complication is when the ll interval lies within the hh interval, which might imply that the occurrence of the ll state will affect the repayments in the hh state.

the default option may be preferable to sacrificing current income. Our solution says that a contract which forgets the previous debt altogether allows a larger transfer to be made for insurance purposes.²⁸

As a comparison to the constrained-efficient, limited commitment contract, we shall consider the optimal static contract which has been studied by Coate and Ravallion (1993). In the context of the example, this amounts to choosing a single transfer (b) to be made in states hl and lh from the household with a good realisation to that with a bad one. This can be thought of as a loan with a rate of interest of -100%. We choose this transfer so that the sustainability constraints are not violated. This will deliver the first-best utilities when the first best is sustainable; that is for $\delta > 0.964$ (or equivalently for discount rates below 3.7%), and in this case the transfer is $d/2$. For discount factors below this, the household making the transfer would be better off under autarky than offering first-best insurance, and so it is necessary to reduce the transfer, until the sustainability constraint of the household making the transfer is just satisfied.

Now consider a sustainable contract which has the feature of the static contract that in the states hl and lh the better-off household makes a fixed transfer (\hat{b}) to the other household, but it resembles more closely a debt contract in that now the contract also specifies a repayment (r) due in the following period. Intuitively a standard debt contract with a fixed, state independent repayment, might be superior to the static contract in that it relieves some of the binding constraints. In the static contract, below $\delta = 0.964$ the binding sustainability constraint is for the household with a good realisation which has to make a transfer to the other one. If the lending household also expects some future return, this relaxes the constraint (see the interpretation of the constrained-efficient contract above). But in the state where a household owes from the previous period and is supposed to lend this period, the extra current commitments will more than offset any beneficial future effects, so a standard debt contract does not (at least in our example) improve upon the static contract. If however the debtor household is *forgiven* any repayment in the state where it is due to lend again, then only the beneficial effect of anticipated repayments remains. When the first-best is not attainable, this allows a larger loan

²⁸It is clear that this contract is not incentive compatible when income shocks are not observed by the other household; claiming to have a bad shock is attractive not only because of receiving a current positive transfer, but also because previous debts are forgotten, and consequently the household would make this claim each period. See Wang (1995) for an analysis of two-sided asymmetric information when contracts *are* enforceable.

δ	Static contract		Debt contract			LC contract
	$2b/d$	Surplus	$2\hat{b}/d$	Interest rate	Surplus	Surplus
> 0.964	100	100	100	-100	100	100
0.95	63	89	84	-76	95	98
0.94	43	69	75	-63	89	97
0.93	23	43	58	-64	80	95
0.92	4	9	42	-64	65	89
0.91	0	0	27	-64	47	79
0.90	0	0	13	-64	25	66
0.89	0	0	0	—	0	50

TABLE 1. Lending, interest rate and surpluses in three different contracts. Surpluses are measured as a percentage of the surplus under the first best contract.

which helps risk-sharing (of course the repayment element of the loan does not, in itself, help risk sharing).

In our example, this type of debt contract does considerably better than the static contract. For example, for a value of the discount factor equal to 0.92, the static contract transfer is reduced to 4% of that of the first-best contract (i.e. $2b/d = 0.04$), and the surplus from this contract is 9% of the potential first-best surplus. In the debt contract with forgiveness the corresponding amounts are 42% and 65%.²⁹ The rate of interest on the loan $((r/\hat{b} - 1) \times 100)$ is -64%. See Table 1 for further details. This contract shares some features of the constrained-efficient limited commitment contract, and does correspondingly better than the static contract.

In Figure 3 we compare how well the three types of contract do, where for each contract type we plot contract surplus as a percentage of the first-best surplus against the discount factor. For the constrained-efficient contract it is assumed that the ex ante surplus is shared equally ($\lambda_0 = 1$). We have also calculated the surplus from a constrained-efficient contract when there is a state-independent penalty from contract violation. We have set this at 5% of the discounted first-best surplus. With a positive penalty, some risk sharing is possible at all discount factors. What is interesting is that over a certain range of discount factors, close to where the no-penalty constrained-efficient contract converges to the autarkic contract, the surplus from this new

²⁹When $\delta = 0.92$, the constrained-efficient contract that shares surplus evenly has values of 69% and 89% respectively.

contract is larger by up to 30% of the first-best surplus, an amount far greater than the size of the penalty itself.

5. EXTENSIONS OF THE MODEL

In principle there is little difficulty in extending the model to a situation where there are more than two households, and where some intertemporal transfer of resources is feasible.³⁰ Unravelling the first-order conditions will, however, be less straightforward. Suppose that there are H households. In addition suppose that each household i has access to a linear storage technology which allows it to transform A_t^i units of the good stored at t into ρA_t^i units at $t + 1$, where $\rho \geq 0$ and $A_0^i \geq 0$ is given. We shall impose the condition that $A_t^i \geq 0$ for all $t = 1, 2, \dots$. Depending on the value of ρ , this technology may also be interpreted as access to a simple credit market, where borrowing is excluded. As before, we assume that contract violation will lead to a breakdown of trust and exclusion from future risk-sharing arrangements together with a direct penalty.³¹ A household which defaults will no longer consume its income each period however, but will be able to self-insure by using the storage technology.³² A slightly different model is possible in which storage, instead of being at the household level, is concentrated at the village level in a common store which is controlled jointly by the village, or by a responsible individual such as the local priest. In this case it would be natural to assume that consequent upon a default the household would lose access to the village store, and autarky would imply consumption equal to income.

At this stage, it will be convenient to change notation slightly. Let household i have a utility of consumption function given by $u_i(c^i)$, and we shall 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 $i \neq H$. Additional state variables will be the end of period storage or inventories, A_s^i , for each household i in state s . We shall use r to index the

³⁰As in the two household case only net transfers are determined and we assume that there is a single point in the period at which all net transfers both to other households and into storage are taken simultaneously.

³¹We do not consider coalitional deviations. An analysis of coalitional deviations is considered by Fafchamps (1995).

³²One important caveat should be made about the manner in which storage has been modelled here. We have assumed that consumption and transfers are simultaneous. Conceptually, the transfer decision could be made prior to the consumption decision, and so an additional sustainability constraint would need to be introduced at the point of consumption, where the household has to choose whether to abide by the contractual stipulation of the consumption/saving division.

state in the following period. Let $f_r^i(A_s^i)$ denote the autarky utility of household i in state r if its storage at the end of the previous period is A_s^i ; this is the utility from self-insurance only. Choice variables in the programming problem will be transfers τ_s^i for $i \neq H$, consumptions c_s^i for each household, the continuation utilities U_r^i for each possible state in the next period and the end of period levels of storage A_s^i for each household. The value function for household H can now be written to depend on the current target utilities and the storage levels from the previous period: $U_s^H(U_s^1, \dots, U_s^{H-1}; A^1, \dots, A^H)$. Notation is otherwise as before.³³ 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

$$U_s^H(U_s^1, \dots, U_s^{H-1}; A^1, \dots, A^H) = \max_{((\tau_s^i), (U_r^i)_{r=1}^S)_{i=1}^{H-1}, ((c_s^i), (A_s^i))_{i=1}^H} u_H(c_s^H) + \delta \sum_{r=1}^S \pi_{sr} U_r^H(U_r^1, \dots, U_r^{H-1}; A_s^1, \dots, A_s^H)$$

subject to a set of updating rules for storage for each household,

$$(14) \quad \rho^i: \quad A_s^i = \varrho A^i + y_s^i - \tau_s^i - c_s^i$$

for all $i \neq H$, and

$$(15) \quad \rho^H: \quad A_s^H = \varrho A^H + y_s^H + \sum_{i=1}^{H-1} \tau_s^i - c_s^H,$$

and subject also to a set of promise keeping constraints

$$(16) \quad \lambda^i: \quad u_i(c_s^i) + \delta \sum_{r=1}^S \pi_{sr} U_r^i \geq U_s^i,$$

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

$$(17) \quad \delta \pi_{sr} \phi_r^i: \quad U_r^i \geq f_r^i(A_s^i),$$

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

$$(18) \quad \delta \pi_{sr} \phi_r^H: \quad V_r(U_r^1, \dots, U_r^{H-1}; A_s^1, \dots, A_s^H) \geq f_r^H(A_s^H),$$

³³The constraint set is no longer convex because of the constraints (17) and (18) below.

for all $r \in S$. Finally, maximization is subject to a set of non-negativity constraints on storage,

$$(19) \quad \omega^i: \quad A_s^i \geq 0$$

for all $i = 1, \dots, H$.

The first-order conditions yield (see Appendix)

$$(20) \quad \frac{u'_H(c_s^H)}{u'_i(c_s^i)} = \lambda^i, \quad \forall i \neq H,$$

$$(21) \quad \lambda_r^i = \frac{\lambda^i + \phi_r^i}{1 + \phi_r^H}, \quad \forall r \in S, \forall i \neq H.$$

and

$$(22) \quad u'_i(c_s^i) = \delta \varrho E_r [u'_i(c_r^i)] + \omega^i / \lambda^i + \delta E_r [\phi_r^i (\varrho u'_r(c_r^i) - f_r^i)].$$

Together, (20) and (21) imply exactly the same updating rule for the marginal utility ratio as before, where household H 's marginal utility is treated as a numeraire. Equation (22) is analogous to the usual Euler equation for an agent able to borrow or save at a gross rate of return of ϱ (see e.g. Hall (1978)). The first term on the right hand side is the expectation of future marginal utility of consumption multiplied by $\delta \varrho$; the second term is the multiplier on the "liquidity constraint" that inventories must have nonnegative value; and the third term reflects the possibility of binding sustainability constraints next period. The sign of this third term is ambiguous and depends upon the sign of the expressions $\{\varrho u'_r(c_r^i) - f_r^i\}$. To interpret this, suppose that an extra unit of the good is stored today by household i , and the extra income this provides next period, ϱ , is consumed by the household. Its utility next period rises by $\varrho u'_r(c_r^i)$, but its autarky utility rises by f_r^i . If $\{\varrho u'_r(c_r^i) - f_r^i\}$ is positive, the extra storage relaxes the sustainability constraint in state r next period. If this holds for each state, then the third term is strictly positive whenever at least one sustainability constraint binds. Whenever the third term is positive, current consumption will be lower relative to future consumption than predicted by the usual Euler equation. Intuitively there is an additional return to saving due to the relaxation of the sustainability constraints. The third term can also be signed if all storage is held communally. In this case the non-negativity constraint on inventories should apply to the sum of inventories instead of applying household by household. Hence, the f_r^i terms drops out of equation (22) as storage does not improve autarky utility and the third term is strictly positive provided only that some sustainability constraint binds next period.

Two questions which arise when there are storage possibilities are whether their existence is welfare improving in the limited commitment environment and whether it is desirable to introduce *ex ante* transfers at the beginning of each period before the current state is known. In fact the addition of storage possibilities need not enhance welfare if there is individual storage as storage may both widen what is technologically feasible but may also increase the payoff in autarky restricting what is sustainable.³⁴ Consider the following two situations. Suppose that there are only two households as before, no direct penalties, and first, let the discount factor be sufficiently low that in the absence of storage no non-autarkic sustainable contract exists (Proposition 2(iv)). Suppose that ρ is greater than the expected marginal rate of substitution between state s in period t and period $t + 1$, that is

$$(23) \quad \rho > \frac{u'_i(y_s^i)}{\delta \sum_{r=1}^S \pi_{sr} u'_i(y_r^i)}$$

for some household i and some state s . Then trivially storage provides some self-insurance against random income, and must be welfare improving. Secondly, suppose for simplicity that aggregate income (though not individual income) is constant and assume that δ is sufficiently high that the first-best contract is just sustainable in the absence of storage. Then each household's consumption will be completely stabilized. If storage is now possible, with $\rho < 1/\delta$, the first-best allocation is unchanged (storage is not utilized), but if ρ satisfies (23) for some household with a binding sustainability constraint then the autarky utility will be increased and the first-best contract is rendered unsustainable. In this case storage reduces potential welfare. In contrast to this ambiguous situation, in the communal storage case the no-storage allocation remains a sustainable contract; hence storage can only ever push out the potential welfare frontier.

The possibility of *ex ante* transfers has been made in an interesting recent paper by Gauthier and Poitevin (1994).³⁵ They do not model storage possibilities but assume that agents have resources which they

³⁴A similar ambiguity may arise if altruism is introduced into the model. The more altruistic are households the more they are willing to transfer to the other household but it also renders the threat to return to autarky incredible so making sustainability more difficult. For an extension to our model incorporating altruism, see Foster and Rosenzweig (1995).

³⁵Their results became available after we had obtained our own results. Their model is more specific than ours in a number of respects, for example, in that only one of the two agents has a random (i.i.d.) income. Nevertheless it seems clear that their basic point will extend to our context.

can transfer *ex ante*. They show how *ex ante* transfers may be used to alleviate the sustainability constraints and in certain circumstances improve welfare. To see why these *ex ante* transfers might be advantageous, consider the two household model and suppose that household 1's sustainability constraint will bind in some states in a particular period, but household 2's constraints never bind so that household 1 may be transferring less than full risk-sharing dictates. Suppose household 1 makes an *ex ante* transfer at the beginning of the period, offset by an equivalent reduction in the non-binding states but not in the binding states so as not to violate sustainability at the point of the *ex ante* transfer. This effectively relaxes the *ex post* sustainability constraints by allowing a larger transfer in binding states but at the expense of smaller transfers in non-binding states. The size of the *ex ante* transfer should be small enough that household 2's *ex post* sustainability constraints are not violated. In our storage model, *ex ante* transfers have an even more direct effect on relaxing the *ex post* constraints by reducing the autarky payoff. However, their use, except possibly in the initial period, cannot lead to a welfare improvement. Consider the above example where household 1 makes an *ex ante* transfer. Reducing this *ex ante* transfer by one unit and increasing the *ex post* transfer in the previous period by $1/\rho$ units³⁶ and stipulating that it is stored provides the same benefits and at the same time relaxes the *ex ante* constraint that household 1 make the *ex ante* payment. Since consumption is unaffected all other constraints will be unaffected. Hence the *ex ante* transfers can be reduced to zero without diminishing welfare.

6. TESTING THE MODEL

We would like to estimate a version of the model described above to see whether limited commitment indeed plays a role in determining consumption allocations. 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 optimal behavior will be forthcoming so long as punishments ($P_i(s)$) for reneging 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

³⁶This argument of course does not work for the initial period where there is no previous period.

punishments are sufficiently small. Finally, we also test 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 Pareto optimal and autarkic allocations.

6.1. 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 reneging (P_i). In addition to these parameters, we will estimate a preference parameter measuring risk aversion. Unbelievable estimates for this preference parameter would provide evidence against our specification.

The discount factor δ would govern not only the division of consumption, but also savings and investment decisions. Although we are able to numerically solve the model when saving is possible (Ligon 1996a), structural estimation of the model with storage is (presently) ruled out on computational grounds. For pragmatic reasons, then, we will abstract from savings (and storage) by scaling household income in each period by a common factor so that aggregate income is equal to aggregate consumption in every period. 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 will be divided among households, not how much of the consumption good ought to be allocated.

Preferences for each household are given by

$$\sum_{t=0}^{\infty} \delta^t \left(\frac{c_{it}^{1-\gamma} - 1}{1-\gamma} \right),$$

where c_{it} is household i 's consumption at time t and γ is the coefficient of relative risk aversion. Preferences are presumed to be common to all households. The chief source of heterogeneity in our estimated model is an idiosyncratic household endowment process. In fact, households employ labor in production. However, by assuming that labor 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.

Although in principle we are able to calculate the efficient contract for economies of H households, in practice we are subject to Bellman's curse of dimensionality; solving the model for 34 households—roughly the size of our sample in each village—involves an impractically large computational expense. Accordingly, we proceed 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). Given assumed preferences, aggregating the rest of the village in this manner is reasonable so long as consumption allocations *within* the rest of the village are fully efficient. Assuming this seems inconsistent, since after all our model is a model of potential inefficiencies in consumption allocation. However, we suspect that the consequences of this inconsistency are relatively unimportant. Even if a few households in the rest of the village have binding sustainability constraints in any given period, the remaining households are likely to have a fully efficient allocation.

6.2. 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. Some summary statistics for this sample are presented in Table 2.

6.3. Estimation. In order to fit the model to our data, we need to solve two nested maximization problems. 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.³⁷ As indicated above, we solve the problem for each household versus the rest of the village.

Estimation of household specific endowment processes is done separately from the estimation of the other structural parameters. We assume that endowment realizations are independent across both time

³⁷This simplification of the problem is due to Fumio ?).

	Village			All
	Aurepalle	Shirapur	Kanzara	
Consumption (’75 Rs)	371.35 (173.03)	495.71 (186.88)	481.59 (174.63)	451.03 (186.43)
Income (’75 Rs)	787.44 (810.06)	698.14 (514.31)	927.72 (773.75)	804.75 (714.95)
Landholdings (Acres)	0.7362 (0.8102)	0.7707 (1.0167)	0.7113 (1.0321)	0.7379 (0.9598)
Household Size	5.8409 (2.5101)	6.3145 (2.8394)	6.6042 (3.5399)	6.2625 (3.0317)

TABLE 2. Summary Statistics. Numbers reported in parentheses are standard deviations; others are means.

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. Since there may be very good or very bad outcomes which are seldom seen in the data, this procedure may lead us to conclude that autarkic outcomes are more attractive than they in fact are. The 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.

Having solved the efficient contract 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, we can deduce the initial

value of the multiplier λ_0^i by solving the system of equations³⁸

$$\lambda_0^i = \left(\frac{c_0^i}{\bar{c}_0} \right)^\gamma \left(\sum_{j=1}^H (\lambda_0^j)^{1/\gamma} \right)^\gamma,$$

where \bar{c}_0 denotes aggregate consumption at time zero. 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 predict the time series $\{\lambda_t^i\}$. Given knowledge of this sequence and of aggregate consumption in each period, we are able to generate a set of predicted consumptions, $\{c_t^i(\theta)\}$.

Comparing predicted and actual consumptions gives us a set of residuals, which we interpret as being due either to measurement error in consumption, or to the influence of some unobserved (by the econometrician) state variables (Rust 1994). We take the sum of squared residuals as our measure of how well the model fits the data. Under the assumption that these residuals are iid normal, minimizing this measure of fit is equivalent to maximizing the likelihood, and the likelihood ratio tests we conduct below are exact. If the assumption of normality is incorrect, then under a weak set of regularity conditions the resulting quasi-maximum likelihood estimator is consistent and asymptotically normal, and the tests below are valid asymptotically.

6.4. Calculation of Models. We begin by fitting each of the different models to the data. Since the autarkic and Pareto optimal models are special cases of each of the limited commitment models, we attempt to estimate each of these special cases first, and then compare the different models.

Autarky. Autarkic allocations will result when the discount factor and punishments for renegeing are both sufficiently small. By Proposition 2(iv), setting δ and P_i all to zero certainly yields autarkic outcomes. Plainly this rules out any transfers between households in any state.

Full insurance. Full insurance allocations result when no household has a binding sustainability constraint and we calculate full insurance allocations by fixing $\delta = 0$ and $P = \infty$. We then estimate the risk aversion parameter, γ , as outlined above.³⁹

³⁸This system follows directly from our parameterization of utility and equation (20).

³⁹If in fact consumption allocations are autarkic, then we will not, of course, arrive at consistent estimates of γ . Since the full insurance and autarkic outcomes correspond when households are risk neutral (there is no need for insurance in this case), we should expect this misspecification to yield estimates of γ biased toward

Coate-Ravallion. Although the Coate-Ravallion (or static limited commitment) model is not nested by the dynamic model, as are the autarkic and Pareto optimal models, it is quite closely related.

In order to solve the Coate-Ravallion model, we begin with the dynamic limited commitment model. We fix $\delta = 0$, and estimate γ and P . With this parameterization, the λ -intervals described above can vary according to the state, just as in the dynamic model, but the shares of consumption evolve rather differently than in the dynamic model. For example, consider an economy with only two households. Let λ_0 be the initial ratio of marginal utilities between the two households; let h_t be some given history, and let the state which occurs at time $t + 1$ be r . Then the sharing rule for the static limited economy will be given by some rule of the form

$$\lambda(h_{t+1}) = \begin{cases} \bar{\lambda}_r & \text{if } \lambda_0 > \bar{\lambda}_r \\ \lambda_0 & \text{if } \underline{\lambda}_r \leq \lambda_0 \leq \bar{\lambda}_r \\ \underline{\lambda}_r & \text{if } \lambda_0 < \underline{\lambda}_r. \end{cases}$$

To see this, consider adapting the model of this paper to the problem addressed by Coate and Ravallion (1993).⁴⁰ The difference between our models is that shares of consumption (alternatively, promised utilities) are permitted to vary in our model, while in the static model of Coate-Ravallion, consumption shares are determined at date zero, and do not vary thereafter save when a sustainability constraint is binding. As a consequence, there is much less scope for risk-sharing in their model; households cannot trade future claims to consumption in exchange for consumption today. Thus, there is no real mobility in the Coate-Ravallion model; a household with a low initial future expected utility can never improve its expected future lot.

Assuming (as we do in our estimation) that $P_i = P$ for all $i = 1, 2, \dots, H$ amounts to assuming that all surpluses are equitably distributed in the Coate-Ravallion model. Note that this is much weaker than assuming that expected utility levels are equal; there is still room for considerably heterogeneity across households under our assumption, since endowment processes and initial promised utilities vary by household.

zero. If one of the limited commitment models is correct, then it is not clear what the bias in our estimate of γ will be.

⁴⁰Coate and Ravallion (1993) assume that households have identical preferences and endowment processes; that is, that they are *ex ante* identical, and hence that outcomes are symmetric. The restrictions we impose in this section actually generalize their model to the extent that we do not assume identical endowment processes, and permit asymmetric outcomes.

Finally, to see that the static limited commitment model nests both Pareto optimal and autarkic outcomes, one can simply compute the λ -intervals as suggested above for $P_i = 0$ (the autarkic case) and for $P_i = \infty$ (the Pareto optimal case).

7. RESULTS

Table 3 presents the structural estimates for each of the nested models described above, along with the log-likelihood associated with each model.

Village	Model	γ	P	δ	Log likelihood
Aurepalle	Autarky	—	—	—	-1621.4723
	Pareto Optimal	1.7415	—	—	-1293.8627
	Static LC	1.6418	0.8231	—	-1284.0782
	Dynamic LC	1.3390	0.8667	0.0433	-1279.6161
Shirapur	Autarky	—	—	—	-1608.7583
	Pareto Optimal	1.8881	—	—	-1391.3818
	Static LC	2.1906	0.9556	—	-1382.5174
	Dynamic LC	2.2916	0.9233	0.5036	-1380.0622
Kanzara	Autarky	—	—	—	-1706.8602
	Pareto Optimal	1.9009	—	—	-1384.8700
	Static LC	2.2198	0.9511	—	-1379.1900
	Dynamic LC	2.3194	0.7926	0.6945	-1365.6128

TABLE 3. Estimates of Model Parameters

We first examine the autarkic model. As remarked above, there are no parameters to estimate for this model, so the log-likelihood reported in Table 3 is simply a measure of the probability that consumption is equal to (scaled) income. For each of the three villages, the log-likelihood is substantially smaller than the log-likelihoods for any of the other models, suggesting that we can firmly reject the autarkic model. In particular, because of our nested specification, twice the difference between the log-likelihoods of the different models has the interpretation of a likelihood ratio. Table 4 presents χ^2 statistics testing the pairwise difference between models. The degrees of freedom for each test are equal to the differences in the number of parameters estimated.

The full insurance (Pareto optimal) model involves estimating the preference parameter γ . Estimates are fairly consistent across villages, ranging from 1.7415 in Aurepalle to 1.9009 in Kanzara. These estimates seem quite reasonable, falling squarely in the range of estimates of relative risk aversion from other empirical studies. Judging by the log

Village	Model	Autarky	Pareto Opt.	Static LC
Aurepalle	Pareto Optimal	655.2193	—	—
	Static LC	674.7882	19.5689	—
	Dynamic LC	683.7124	28.4931	8.9242
Shirapur	Pareto Optimal	434.7532	—	—
	Static LC	452.4818	17.7287	—
	Dynamic LC	457.3923	22.6391	4.9104
Kanzara	Pareto Optimal	643.9805	—	—
	Static LC	655.3403	11.3598	—
	Dynamic LC	682.4947	38.5142	27.1544

TABLE 4. LR tests of differences between models. The relevant critical values for the χ^2_1 and χ^2_2 statistics are 3.84 and 5.99, respectively, so that every statistic in the table is significant at a 95 per cent level of confidence.

likelihoods, the full insurance model provides a dramatic improvement over the autarkic model in terms of model fit, an observation which is confirmed by an examination of Table 4. The full insurance model provides a significantly better fit to the data than does the autarkic model in each village.

Modifying the updating rule and estimating the punishment parameter P as well as the preference parameter γ gives us a version of the static limited commitment model of Coate and Ravallion.⁴¹ P is measured in units of utils; in order to make estimates of P easier to interpret, we calculate the ratio of estimates of P to the average surplus utility for each village. So, for example, the estimated punishment for reneging on the Coate-Ravallion contract in Aurepalle amounts to 82.31 per cent of the average surplus generated by this contract, relative to autarky. The sizes of the estimated punishments are similar in Shirapur and Kanzara; 95.56 per cent and 95.11 per cent, respectively. Estimated punishments in Shirapur and Kanzara are also close when expressed in terms of utils; 9.698×10^{-4} and 1.006×10^{-3} , respectively.⁴²

⁴¹If allocations are in fact Pareto optimal, then P will be unidentified; although there will be some minimum level of P which delivers the Pareto optimal outcome, any larger punishment will deliver the same allocation. In this case, we'd like to recover the smallest value of P consistent with the allocation; accordingly, we add to the objective function a small penalty function which is increasing in P at a linear rate for positive P .

⁴²By this point, the reader has no doubt noticed that no standard errors are reported for the coefficient estimates presented in Table 3. The presence of small 'flat' areas in our likelihood in the neighborhood of our estimates, combined with

Although the absolute punishment is much larger in Aurepalle, with a value of 0.0121, it is actually slightly smaller when expressed as a proportion of average surplus, a consequence of the fact that average surpluses appear to be considerably larger in Aurepalle than in the other two villages. Although the absolute values of punishments may appear to be small, their introduction leads in every case to a significant improvement in the likelihood, indicating that the punishments are in fact significantly different from zero, and that the introduction of static limited commitment does in fact improve our ability to explain the data.

Estimating a third parameter, the discount factor δ , gives us a version of the dynamic limited commitment model of this paper. We remind the reader that while this model nests the autarkic and Pareto optimal models, it does not nest the static limited commitment model, due to the difference in the updating rule for that model. Estimates of δ vary dramatically across villages, with the lowest value of 0.0433 in Aurepalle. Estimated values of δ in the remaining villages are more similar—0.5036 in Shirapur and 0.6945 in Kanzara. While even the latter two estimates seem low relative to rates of discount estimated in developed countries, they are consistent with estimates reported by Pender (1996), who uses experimental techniques to estimate rates of discount in Aurepalle. Though this experimental evidence is reassuring in the case of Shirapur and Kanzara, it makes the estimated value of δ in Aurepalle quite unsatisfactory. However, low estimates of δ in Aurepalle can be interpreted as a sign that savings or storage (which we neglect in our estimation) is very important in this village.

Another possibility, of course, is that the dynamic limited commitment model is simply incorrect. Of course, the model is bound to be literally incorrect—a better question is how well the model performs relative to some well-specified alternatives. On this question, there is only weak evidence for the dynamic model in Shirapur, where dynamic limited commitment improves only moderately (but significantly) on the alternatives; however, the dynamic model easily beats the alternatives in both Aurepalle and Kanzara (Table 4). Estimated values of γ under the dynamic model are higher than in the Coate-Ravallion

the presence of many local maxima makes computing either asymptotic approximations to the standard errors or bootstrapping the standard errors extraordinarily problematical (this latter alternative would also be extremely expensive). Because nothing in our model selection procedure hinges on the standard errors, we have chosen not to attempt to calculate them. Nonetheless, a simple likelihood ratio test leaves little doubt that punishments in Shirapur and Kanzara are, indeed, quite close.

or full insurance models in Shirapur and Kanzara. Estimated punishments are not too different from the static model in Shirapur and Kanzara, but are much larger in Aurepalle.

From the tests of Table 3 and Table 4, then, we conclude that for each of the three villages, the dynamic limited commitment provides a better explanation of consumption allocations than does either the Autarkic or Pareto optimal models. The static limited commitment model improves over autarky in every village, but provides a significantly better fit to the data than the Pareto optimal model only in Aurepalle and Kanzara. Finally, estimated values of the discount factor in Aurepalle are unrealistically low, suggesting misspecification; we suspect that privately controlled savings and storage may be of particular importance in this village.

It may be instructive to examine some less formal measures of fit as well.⁴³ Table 5 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. The orderings of models according to how highly their predicted consumptions are correlated with actual consumptions is the same as the ordering provided by the likelihood ratios of Table 4, with a single exception: the static model in Shirapur predicts consumptions which have a correlation coefficient of 0.6770 with actual consumption, while the dynamic model has a similar coefficient which is slightly smaller, at 0.6674. This is the lowest correlation between the dynamic model and actual consumption: Aurepalle and Kanzara record more respectable coefficients of 0.7473 and 0.7695.

Looking at the fourth column, we can examine the correlation between the consumptions of the Coate-Ravallion model and the dynamic limited commitment model. These are generally quite high, exceeding 97 per cent in each of Aurepalle and Kanzara. Curiously, this correlation is lowest in the one village in which the static and dynamic models are least statistically distinguishable: in Shirapur the correlation is only 96 per cent. On the other hand, nearly all of the correlations for Shirapur tend to be somewhat lower; it seems that none of these models provides as good a fit in Shirapur as they do in the other two villages,

⁴³ A 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 clearly interested in a highly nonlinear relationship, the correlation coefficient may be quite misleading when regarded as a measure of fit.

a sense which is confirmed by the relatively low likelihoods reported for Shirapur in Table 4.

Table 5 also adds yet another model of consumption allocation, titled 'Ad hoc' in the table.⁴⁴ Consumption allocations for this rule are simple but, as the name suggests, *ad hoc*. In particular, at time t household i 's consumption is determined by a weighted average of own income and aggregate village income, or

$$c_t^i = \alpha_i y_t^i + (1 - \alpha_i) \frac{1}{H} \sum_{j=1}^H y_t^j.$$

The parameter α_i was estimated for each household using restricted least squares; incomes were scaled to have a sample mean identical to the sample mean of consumption. Despite the fact that this model has many more parameters than does the structural model, it does not perform particularly well; in no village does it outperform either of the limited commitment models, and in only one (Aurepalle) does it outperform the Pareto optimal model.⁴⁵

We have presented very strong evidence that models incorporating limited commitment are capable of doing a much better job of explaining actual consumption allocations than are models of full insurance or autarky. We have also seen that a dynamic model of limited commitment outperforms a static model in each of the three villages we consider. Nonetheless, given the simplicity of the models we've proposed, and the necessarily stylized features of the model economy we've estimated, we would like to have some way to evaluate ways in which the dynamic limited commitment model fails to capture some aspect of consumption allocation in these villages.

Table 6 takes a simple approach to the task of identifying strengths and weaknesses of our model. In its first panel, it presents coefficient estimates from a regression of actual consumption on a constant, household income, aggregate consumption, and finally the predicted consumptions from the dynamic limited commitment model. The results are quite striking. With a single exception (aggregate consumption in Shirapur), in each of the three villages, as well as for all three villages pooled, each of the explanatory variables is highly significant. If the limited commitment model actually captured all of the systematic variation in consumption, we would expect to observe a coefficient

⁴⁴We thank Naryana Kocherlakota for suggesting that we add this allocation rule.

⁴⁵If we were to add a set of fixed effects to the estimation of the *ad hoc* model, and take logs of consumptions and incomes, then the allocation rule of the full insurance model would be a special case of the *ad hoc* rule, with $\alpha_i = 0$.

Village	Model	Reality	Autarky	Pareto Opt.	Static LC	Dynamic LC
Aurepalle	Autarky	0.6705	1.0000	—	—	—
	Pareto Optimal	0.6826	0.6071	1.0000	—	—
	Static LC	0.7093	0.6618	0.9824	1.0000	—
	Dynamic LC	0.7473	0.6931	0.9782	0.9844	1.0000
	Ad hoc	0.7067	0.7857	0.5293	0.5756	0.5783
Shirapur	Autarky	0.4972	1.0000	—	—	—
	Pareto Optimal	0.6512	0.3655	1.0000	—	—
	Static LC	0.6770	0.4045	0.9834	1.0000	—
	Dynamic LC	0.6674	0.4493	0.9576	0.9617	1.0000
	Ad hoc	0.6160	0.6725	0.4901	0.5036	0.5047
Kanzara	Autarky	0.6330	1.0000	—	—	—
	Pareto Optimal	0.7539	0.7649	1.0000	—	—
	Static LC	0.7584	0.7732	0.9883	1.0000	—
	Dynamic LC	0.7695	0.7744	0.9709	0.9768	1.0000
	Ad hoc	0.6854	0.6779	0.6225	0.6260	0.6327
All	Autarky	0.5648	1.0000	—	—	—
	Pareto Optimal	0.7178	0.5836	1.0000	—	—
	Static LC	0.7336	0.6103	0.9863	1.0000	—
	Dynamic LC	0.7456	0.6386	0.9695	0.9748	1.0000
	Ad hoc	0.6846	0.6225	0.5731	0.5884	0.5883

TABLE 5. Simple correlations between consumption from different models.

of one on the LC variable, and for the remaining coefficients to be insignificant; instead we observe coefficient estimates ranging from 0.2901 in Aurepalle to 0.6293 in Shirapur, and nearly each of the remaining coefficients is significant. The limited commitment model seems to contribute something quite important to explaining consumption, but clearly does not explain all of the variation in consumption.

One possible explanation for this failure is that our model has failed to capture some important sources of heterogeneity across households. Accordingly, in the second panel of Table 6, we add a set of fixed effects to control for this possible heterogeneity. Although this does not much affect our estimates of the income coefficients, which remain right around 0.10 and are all significant, it has a dramatic effect on our estimates for the LC and aggregate consumption variables. The estimated coefficients for the LC variable all increase substantially, ranging from 0.6627 in Aurepalle to 1.5905 in Shirapur. The standard errors also rise, presumably because of the relatively high correlation between the LC variable and the fixed effects; however, the LC variable remains significant in both Aurepalle and Kanzara, as well as for all three villages

Village	Income	Agg. Cons.	LC Cons.
Aurepalle	0.0944* (0.0129)	0.6376* (0.1078)	0.2901* (0.0627)
Shirapur	0.0956* (0.0193)	0.3564 (0.3707)	0.6293* (0.0615)
Kanzara	0.0420* (0.0152)	0.4995* (0.1756)	0.4989* (0.0569)
All	0.0476* (0.0087)	0.3731* (0.1448)	0.5930* (0.0324)
Aurepalle	0.0926* (0.0198)	0.2427 (0.1917)	0.6627* (0.1668)
Shirapur	0.0976* (0.0241)	-0.6118 (0.9928)	1.5905 (0.9385)
Kanzara	0.1199* (0.0211)	0.2214 (0.3527)	0.7655* (0.3290)
All	0.1020* (0.0126)	0.0850 (0.1371)	0.8583* (0.0775)

TABLE 6. Consumption Regressions. Figures in parentheses are standard errors. The first set of estimates regresses consumption on a constant, household income, aggregate consumption, and estimated consumption from the dynamic limited commitment model. The second set of estimates the constant with a set of fixed effects, one for each household.

pooled. What is most striking, however, is that *none* of the aggregate consumption coefficients is significant under this specification. The magnitude of these coefficients falls and standard errors increase for each village.

In interpreting these results, it is useful to note that we can get a version of the consumption allocation rule for each of autarky, full insurance, and the *ad hoc* rule introduced above as linear combinations of the fixed effects, incomes, and aggregate consumptions which appear on the right hand side variables of the estimating equation for the second panel of Table 6. The fact that the LC variable continues to help to explain consumption provides yet more evidence that the dynamic limited commitment model dominates each of these other three models as providing an explanation for observed patterns in consumption.

The other thing we should note is that, despite our best efforts, the limited commitment model predicts too much insurance. This is

reflected in the fact that the income coefficients continue to help explain consumption. One possible explanation for this has to do with our assumption (for tractability) regarding the efficiency of allocation in the rest of the village (see Section 6.1). However, we regard it as more likely that consumptions sometimes respond to income even if no sustainability constraints are binding. Such a response would be generated by a model with both limited commitment and private information about individual levels of storage, for example. Ligon (1996b) presents direct evidence that private information plays a role in determining allocations in these villages, so constructing and testing models with both limited commitment and private information seems a promising direction for future research.

APPENDIX

7.1. Proof of Proposition 2.

Proof. (i) 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 overlap. The result follows. (ii) 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 $\underline{\lambda}_s$ is smaller than the autarkic marginal utility ratio, 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) $u(y_1(s_t) - \tau(h_t)) \leq u(y_1(s_t))$, which again implies that $\underline{\lambda}_s$ is at least as small as the autarkic marginal utility ratio. A symmetric argument for household 2 establishes that $\bar{\lambda}_s$ is at least as large as the autarky marginal utility ratio. (iii) 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 $U_{t+1}(h_{t+1})$ 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 is at the autarkic level at $\underline{\lambda}_s$. A symmetric argument for household 2 establishes the result for $\max\{\bar{\lambda}_s\}$. (iv) 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 (ii) $\underline{\tau}_s \geq 0 \geq \bar{\tau}_s$ and from part (iii) $\underline{\tau}_s = 0$ when $\underline{\lambda}_s = \min\{v'(y_2(s))/u'(y_1(s))\}$ and $\bar{\tau}_s = 0$ when $\bar{\lambda}_s = \max\{v'(y_2(s))/u'(y_1(s))\}$. Let $\chi_{rs}^u = u'(y_1(r))/u'(y_1(s))$ and $\chi_{rs}^v = v'(y_2(r))/v'(y_2(s))$. For $\delta = 0$ there is clearly no non-autarkic contract. By the maximum theorem the contract is continuous in δ , and by part (i) all intervals overlap for large δ , so for δ small the λ -intervals will be disjoint if the autarkic marginal utility ratios are distinct and approximately coincident if two or more states have the same autarkic marginal utility ratios. Using Proposition 1, it is then possible to calculate

$$v(y_2(s)) - v(y_2(s) + \bar{\tau}_s) = \delta \sum_r [\alpha_{sr} \pi_{sr} (v(y_2(r) + \underline{\tau}_r) - v(y_2(r) + \bar{\tau}_r)) \mid \chi_{rs}^v > \chi_{rs}^u]$$

$$u(y_1(s)) - u(y_1(s) - \underline{\tau}_s) = \delta \sum_r [\alpha_{sr} \pi_{sr} (u(y_1(r) - \bar{\tau}_r) - u(y_1(r) - \underline{\tau}_r)) \mid \chi_{rs}^v < \chi_{rs}^u]$$

where α_{sr} is some positive parameter. Linearizing these two equations about the income levels and adding gives

$$\begin{aligned} (\underline{\tau}_s - \bar{\tau}_s) &= \delta \sum_r [\alpha_{sr} \pi_{sr} (\underline{\tau}_r - \bar{\tau}_r) \chi_{rs}^v \mid \chi_{rs}^v > \chi_{rs}^u] \\ &\quad + \delta \sum_r [\alpha_{sr} \pi_{sr} (\underline{\tau}_r - \bar{\tau}_r) \chi_{rs}^u \mid \chi_{rs}^v < \chi_{rs}^u] + o(\underline{\tau}_s - \bar{\tau}_s) \end{aligned}$$

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. \square

7.2. Derivation of First Order Conditions in Section 5. The first order conditions yield

$$(25) \quad \rho^i = \rho^H, \forall i,$$

$$(26) \quad u'_i(c_s^i) = \rho^i / \lambda^i, \quad i \neq H,$$

which implies from (25), (20) in the text

$$(27) \quad -\frac{\partial U_r^H(U_r^1, \dots, U_r^{H-1}, A_s^1, \dots, A_s^H)}{\partial U_r^i} = \frac{\lambda^i + \phi_r^i}{1 + \phi_r^H}, \forall r \in S, \forall i \neq H,$$

where

$$(28) \quad \rho^i = \delta \sum_r \pi_{sr} \frac{\partial U_r^H(U_r^1, \dots, U_r^{H-1}, A_s^1, \dots, A_s^H)}{\partial A_s^i} (1 + \phi_r^H) - \delta \sum_r \pi_{sr} \phi_r^i f_r'^i + \omega^i,$$

for all $i \neq H$; together with the envelope conditions

$$(29) \quad \lambda^i = -\frac{\partial U_s^H(U_s^1, \dots, U_s^{H-1}, A^1, \dots, A^H)}{\partial U_s^i}, \quad \forall i \neq H,$$

and

$$(30) \quad \frac{\partial U_s^H(U_s^1, \dots, U_s^{H-1}, A^1, \dots, A^H)}{\partial A^i} = \gamma \rho^i, \forall i \neq H.$$

From (26), substituting ρ^i from (28), and using (30) moved forward one period to substitute for $\frac{\partial U_r^H}{\partial A_s^i} = \gamma \rho_r^i$ (where a subscript r denotes a future variable):

$$(31) \quad u'_i(c_s^i) = (\lambda^i)^{-1}(\omega^i + \delta \sum_r \pi_{sr} [(1 + \phi_r^H) \gamma \rho_r^i - \phi_r^i f_r^i]).$$

Next, from (27) and (29) we get the updating equation (21) in the text. In (31), substitute for $(1 + \phi_r^H)$ from (21), and for ρ_r^i from (26) moved forward one period, which yields (22).

REFERENCES

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica* 56, 383–396.
- Asheim, G. and J. Strand (1991). Long-term union-firm contracts. *Journal of Economics* 53, 161–184.
- Beaudry, P. and DiNardo (1995). Is the behaviour of hours worked consistent with implicit contract theory? *Quarterly Journal of Economics* 110, 743–768.
- Bulow, J. and K. Rogoff (1989, March). Sovereign debt: Is to forgive to forget? *American Economic Review* 79, 43–50.
- 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 Idcs: Credit markets, insurance, and welfare. *Scandinavian Journal of Economics* 94, 253–273.
- Fafchamps, M. (1992). Solidarity networks in preindustrial societies: Rational peasants with a moral economy. *Economic Development and Cultural Change* 41(1), 147–174.
- Fafchamps, M. (1995). Risk sharing, quasi-credit, and the enforcement of informal contracts. Manuscript.
- 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. and M. Poitevin (1994). Using ex ante payments in self-enforcing risk-sharing contracts. Discussion Paper 0394, Universite de Montreal.
- Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86, 971–987.
- Hayashi, F. (1996). Analysis of household saving: Past, present, and future. *The Japanese Economic Review* 47(1), 21–33.
- Kimball, M. S. (1988). Farmers' cooperatives as behavior toward risk. *American Economic Review* 78, 224–236.
- Kletzer, K. and B. Wright (1996). Sovereign debt as intertemporal barter. Manuscript.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies* 63(4), 595–610.

- Ligon, E. (1996a). Notes on storage, informal insurance arrangements, and limited commitment. <http://are.berkeley.edu/~ligon/Papers/storage.ps>.
- Ligon, E. (1996b). Risk-sharing and information: Theory and measurement in village economies. <http://are.berkeley.edu/~ligon/Papers/reid.ps>.
- 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.
- Nash, M. (1966). *Primitive and Peasant Economic Systems*. San Francisco: Chandler Publishing.
- 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. (1996). Mutual insurance as an elusive concept in traditional rural societies. Forthcoming in *Journal of Development Studies*.
- 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.
- 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.
- Rust, J. (1994). Estimation of dynamic structural models, problems and prospects: discrete decision processes. In C. A. Sims (Ed.), *Advances in Econometrics*, Volume II, Cambridge, pp. 119–170. Sixth World Congress: Cambridge University Press.
- 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. (1990). Credit markets in northern Nigeria: Credit as insurance in a rural economy. *The World Bank Economic Review* 4, 251–269.
- Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria. *Review of Economic Studies* 63, 495–526.
- Walker, T. S. and J. G. Ryan (1990). *Village and Household Economies in India's Semi-arid Tropics*. John Hopkins.
- Wang, C. (1995, October). Dynamic insurance with private information and balanced budgets. *The Review of Economic Studies* 62(4), 577–596.

UNIVERSITY OF CALIFORNIA, BERKELEY AND GIANNINI FOUNDATION

UNIVERSITY OF WARWICK

KEELE UNIVERSITY

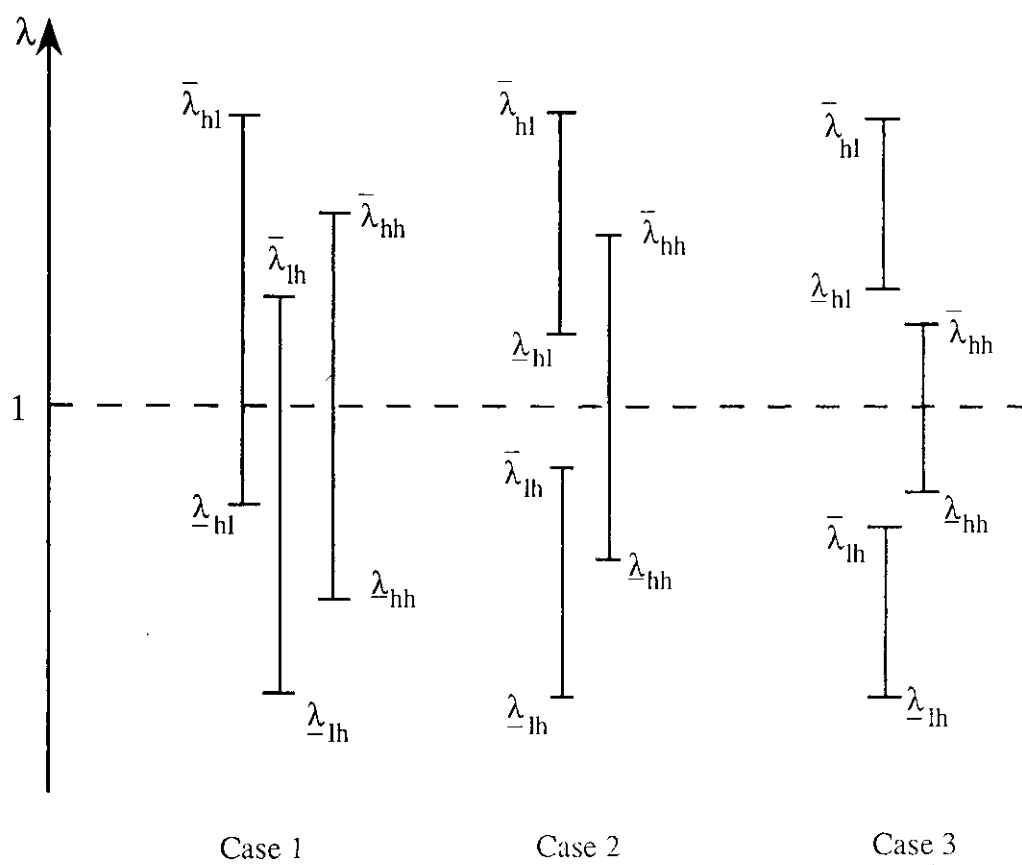


FIGURE 1. The relationship between the intervals

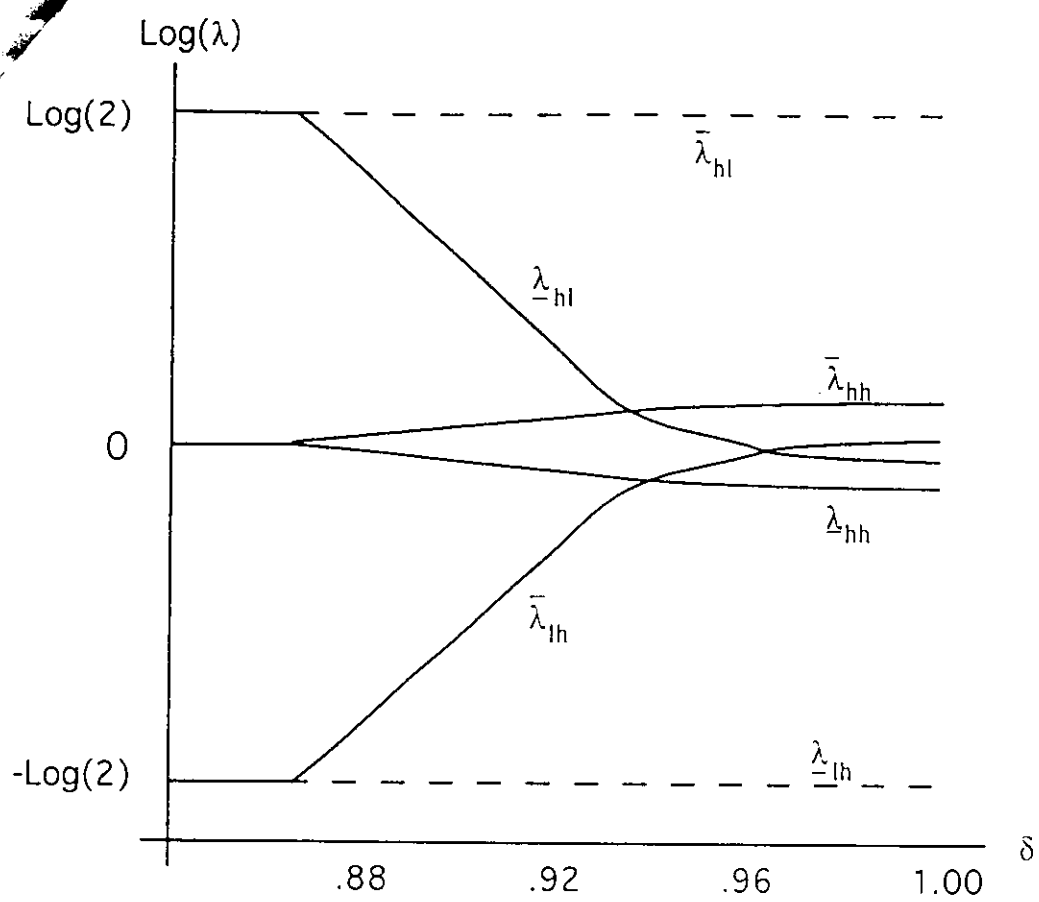


FIGURE 2. Dependence of interval endpoints on δ .

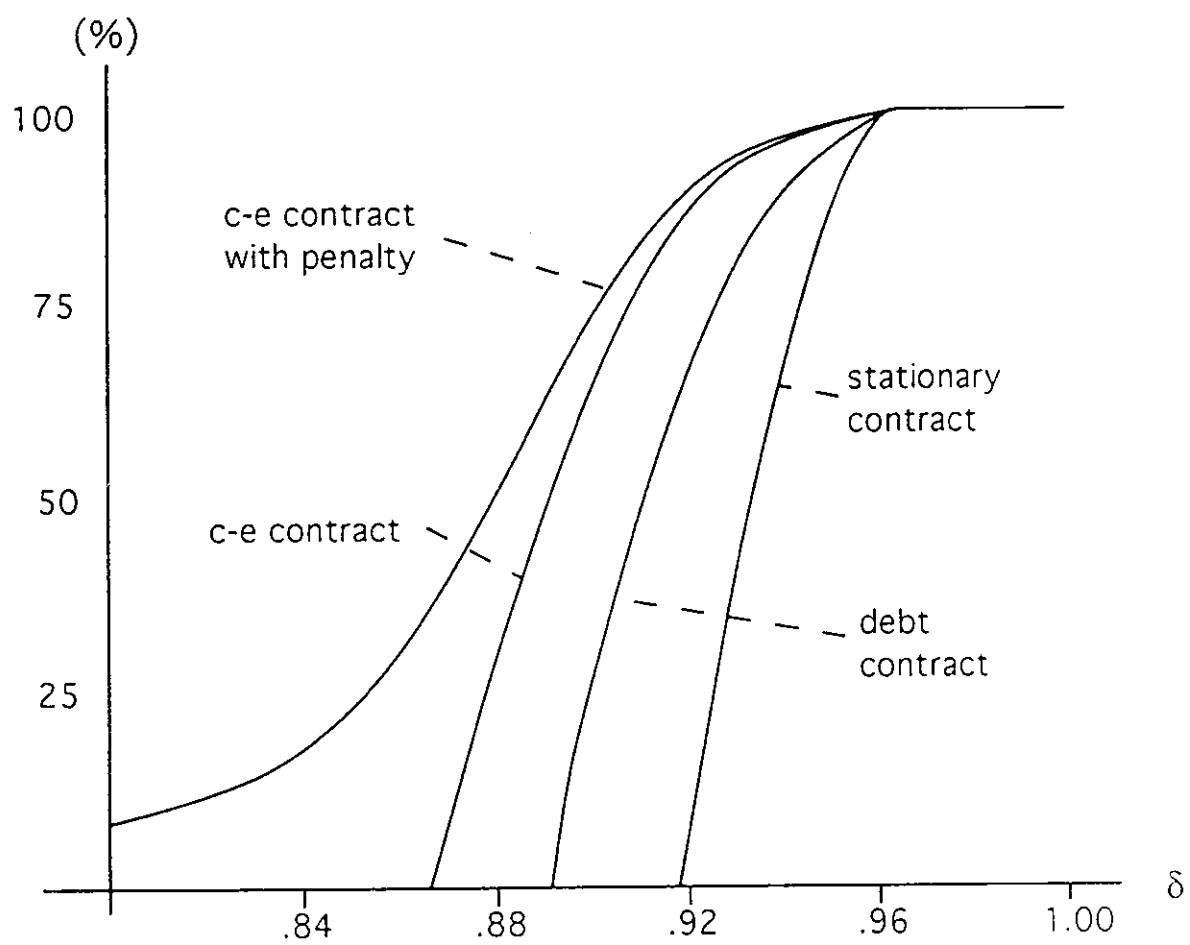


FIGURE 3. Contract surplus as a percentage of the first-best surplus