Scand. J. of Economics 94(2), 253-273, 1992

Household Saving in LDCs: Credit Markets, Insurance and Welfare*

Angus Deaton

Princeton University, Princeton, NJ, USA

Abstract

Some ways in which farmers in LDCs can protect their living standards against fluctuations in income are discussed. After considering the theory of consumption under uncertainty when there is no or limited borrowing, the case where some borrowing is allowed is also examined. Empirical evidence from some LDCs is used to look at (i) household borrowing and lending, their importance and timing, and their role in smoothing consumption, and (ii) the life-cycle behavior of consumption and income. The results suggest that "hump" life-cycle saving is not likely to be a very important generator of wealth in LDCs and provide further evidence on the limited role of credit markets.

I. Introduction

Agricultural income is inherently uncertain. Weather, pests, disease and fires make yields uncertain, and the notorious variability of agricultural prices can generate fluctuations in farmers' incomes even when output is stable. In poor countries, most of the population earns its living from agriculture, either directly or indirectly. For many of these people, a poor harvest or a low harvest price can threaten disaster, even if, on average, agricultural incomes are sufficient to provide a sustainable standard of living. In such circumstances, the protection of living standards requires that resources be transferred across time, from good years to bad years. For this to work, someone, farmers, communities, or governments, have to be able to look ahead, and make adequate provision for the future. Until

*The original, longer version of this paper, Deaton (1991b) can be obtained from the Research Program in Development Studies at Princeton University. I should like to thank Kristin Butcher for splendid research assistance, and the Lynde and Harry Bradley Foundation for financial support. Some of the results reported here come from a research project on saving behavior supported by the World Bank; the views expressed here are those of the author and should not be attributed to the World Bank or its affiliated organizations.

recent years, much of the development literature, and in particular the literature on project evaluation in developing countries, UNIDO (1972), Little and Mirrlees (1974) and Squire and van der Tak (1975), supposed that individuals were unlikely to be able to make adequate intertemporal plans, so that the government should do so on their behalf. In practice, there are many government policies that play such a role, at least in part. Many governments, particularly in Africa, set agricultural procurement prices that do not vary with world prices, and argue that such schemes help stabilize farm incomes. Governments in many developing countries are involved in the provision of credit, which can also play a role in insuring consumption in the face of income variability. There are also a wide range of relief policies, from *ad hoc* famine relief to more regular food for work schemes, which can provide some minimum living standards in bad times.

This paper is concerned with non-governmental consumption insurance schemes. I discuss some of the ways in which farmers can protect their living standards against fluctuations in income. Where there are credit markets, individuals can borrow and lend, and there is no need for consumption to be tied to income in the short run. But even without credit markets, or with limited credit facilities, money or goods can be put by in good years to provide a buffer against the bad times that will sooner or later follow. Knowledge of the extent to which poor households use these mechanisms is necessary if we are to assess the appropriate role, if any, for programs provided by the state.

Section II of the paper considers the theory of consumption under uncertainty when there is no borrowing, or limited borrowing at penal rates. The starting point is the model of liquidity constraints in Deaton (1990a, 1991a). I consider the effectiveness of simple rules-of-thumb of the type that agents could easily implement, and evaluate the welfare consequences of these against fully optimal strategies. I also look at the case where some borrowing is allowed, albeit in limited amounts, and at high rates of interest. In this model with "money-lenders", agents can pay to avoid the worst outcomes in a pure buffering model, when a succession of low incomes, which has exhausted precautionary assets, is followed by a further bad year.

Section III considers some empirical evidence from Côte d'Ivoire, Thailand and Ghana. I look at two separate issues: (i) household borrowing and lending in Côte d'Ivoire and Ghana, their importance and their timing in an attempt to assess the part they play in smoothing consumption, and (ii) the life-cycle behavior of consumption and income, here using data from Thailand and Côte d'Ivoire. As argued by Carroll and Summers (1991), if credit markets are good enough to support long-term (low-frequency) consumption smoothing, and if the life-cycle hypothesis is true, then the ratio of young people's life-time resources to those of old people

is larger in fast growing countries (Thailand) than in slow-growing countries (Côte d'Ivoire), so that consumption profiles should be more tipped towards the young in the former than in the latter. In fact, the opposite is true, and I show that, in contrast, consumption tracks current income very closely. There results suggest that "hump" life-cycle saving is not likely to be a very important generator of wealth in these countries, and provides further evidence on the limited role of credit markets, at least for long-term consumption smoothing. Section IV provides some brief conclusions.

II. Theory: Saving with Borrowing Constraints

Individuals in poor countries borrow and lend money, and perhaps do so to prevent shortfalls in consumption. However, it remains implausible that agents can always borrow as much as they like for consumption purposes, so that it is important to examine behavior when borrowing is not permitted. The essential point to note is that the inability to *borrow* does not imply that the consumer cannot *save*. Indeed, the fact that borrowing may be unavailable when most needed is itself a reason to set aside something when times are good; liquidity constraints reinforce the precautionary demand for assets. Many consumers may never wish to borrow; people who are patient, or for whom the return on assets is sufficient to overcome their impatience, will tend to postpone consumption, building up assets as they go, so that temporary short-falls of income are unlikely to pose a problem, except perhaps early in life. But for those who are impatient, or who are poor enough to feel that future consumption is an inadequate reward for postponing current consumption, lack of borrowing facilities will be a real disadvantage. For such consumers, it is essential to hold some assets that can be used to buffer consumption when incomes are low.

A simple model of optimal buffering is constructed in Deaton (1991a), and is summarized here as a starting point. Preferences take the standard form,

$$u = E_t\left\{\sum_{\tau=t}^{\infty}(1+\delta)^{t-\tau}v(c_t)\right\} \tag{1}$$

where $\delta > 0$ is the rate of time preference, and $v(c_t)$ is the instantaneous (sub) utility function, assumed to be increasing, strictly concave, and differentiable. The evolution of assets is given by

$$A_{t+1} = (1+r)(A_t + y_t - c_t) \tag{2}$$

where y_t is labor income, A_t is real assets and r is the real interest rate. I

assume that $r < \delta$, so that the agent's impatience outweighs the incentive to accumulate. The real interest rate is treated as fixed and known, and all the uncertainty is focussed on labor income y_t. Labor is inelastically supplied, and y_t is a stationary random variable with support $[y_0, y_1]$, with $y_0 > 0$ and $y_0 \leq y_1 \leq \infty$, so that income cannot fall below the positive floor y_0. I start from the most obvious form of the borrowing restriction, *viz.*

$$A_t \geq 0. \tag{3}$$

The simplest way to solve the consumer's problem is to start from the modification of the usual Euler equations that is brought about by the presence of the borrowing constraint (3). Define the state variable x_t "cash on hand", by

$$x_t = A_t + y_t. \tag{4}$$

x_t is the maximum that can be spent on consumption in period t. Consumption in periods t and $t+1$ must satisfy

$$\lambda(c_t) = \max[\lambda(x_t), \beta E_t \lambda(c_{t+1})] \tag{5}$$

where $\lambda(c)$ is the marginal utility of c, i.e. $\lambda(c) \equiv v'(c)$, $\beta = (1+r)/1+\delta)$, and $\beta < 1$ since $r < \delta$. If the consumer is constrained, consumption can be no higher than x_t, and the marginal utility no lower than $\lambda(x_t)$. The constraint will bind if marginal utility at x_t is higher than the discounted expected marginal utility next period, otherwise the two marginal utilities are equated in the usual way. Note that the expectation itself takes account of the possibility of future constraints.

The solution to (5) depends on the stochastic structure of the income process y_t. Here, since my interest is in extending the model in other directions, I deal with only the simplest case where income is independently and identically distributed over time. Deaton (1991a) discusses cases where income is a first order autoregressive or moving average process. In the i.i.d. case, the optimal rule is to make consumption a function of cash on hand, i.e.

$$c_t = f(x_t) \tag{6}$$

where, by virtue of (4) and (2), x_t evolves according to

$$x_{t+1} = (1+r)\{x_t - f(x_t)\} + y_t \tag{7}$$

and y_t, labor income, is an i.i.d. stochastic process.

In general, it is not possible to write down an analytical form for the function $f(x)$. However, the theory can be used to infer its general properties, and in practice, given specifications for the utility function and the distribution function of incomes, as well as the interest rate and time-preference parameters, the function can be calculated numerically. Given

his cash on hand, the consumer should spend everything if the total is below some critical value x^*, say. Above x^*, something is put by for the future, and the marginal propensity to retain cash, although always less than unity, is an increasing function of cash on hand. If the consumer follows the optimal rule, consumption can be very much smoother than income; if the income process is normally distributed with mean 100 and standard deviation 10, if utility is isoelastic with a coefficient of relative risk aversion of 3, and if the interest and time-preference rates are 5 per cent and 10 per cent respectively, the standard deviation of consumption is 5, which is half that of income. This can be achieved with very low average levels of "buffer" assets, averaging only 7 and rarely more than 10. However, consumption fluctuations are necessarily asymmetric. It is always possible to prevent consumption being too high, since resources can be kept for the future. But if assets are exhausted, as must happen from time to time, there is always the *possibility* of low income immediately thereafter, and the unprotected consumer has no choice but to cut consumption to match. Of course, these outcomes are no worse than would be the case if consumption were equal to income, and the optimal buffering strategy much reduces their likelihood. When they occur, the shadow price of the borrowing constraint, the shadow price of loans, will rise to high levels; for the same parameter values, rates of 30 per cent occur every twenty years or so, and rates of well over 100 per cent are possible. There is a demand for loans at even very high rates of interest.

The calculation of the optimal policy (11) from the conditions (5) and (7) is not a trivial task. The function has to be approximated by points along a grid, or by some suitable polynomial, and then values or parameters chosen to satisfy the functional equations as closely as possible. The question then arises as to whether consumers could reasonably be expected to solve this problem for themselves, and if not, whether there might not exist simpler, more intuitive rules of thumb that might do nearly as well. In Figure 1, the broken line shows the optimal policy for a consumer with the preferences and income process described above, while the piecewise linear function shows the simple rule of thumb: "spend all cash on hand up to mean income, and 30 per cent of any excess". (The third function on the graph will be discussed below.) My choice of rule of thumb is not arbitrary, but was guided to some extent by knowledge of the optimal function. However, the critical point at which saving begins was taken to be mean income rather than any approximation to the optimal critical point. Note also that the rule-of-thumb function is never below the optimal function, so that the consumption is always too high, and I have made no attempt to achieve a better approximation. My concern is more that the rule should be simple, simple enough to have plausibly evolved from trial and error.

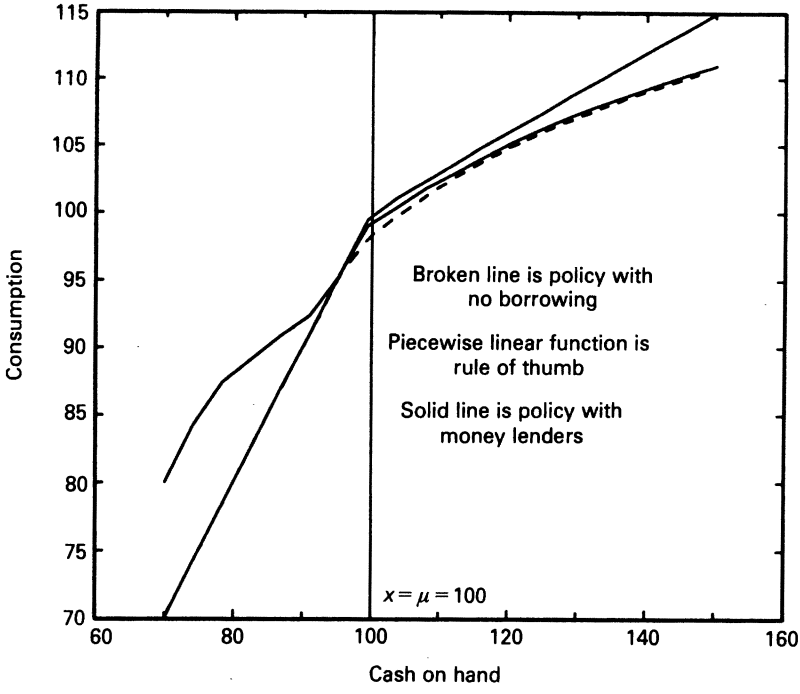

Fig. 1. Three consumption policies.

There are various ways of assessing the performance of such a rule. Figure 2 is one; it illustrates a 200-year realization of income, consumption, and assets under the rule of thumb. Income values are random independent draws from $N(100, 10)$, the agent begins with no assets, and the process is allowed to run, governed by (6) and (7) with $f(x_t)$ in (6) taken to be the piecewise linear function, i.e.

$$f(x_t)=x_t-0.7\,(x_t-\mu)I(x_t>\mu) \tag{8}$$

where $I(.)$ is the indicator function. The outcomes in the figure are very close to those generated by the optimal policy, indeed there seems to be no practical difference between the two. Not only does the rule-of-thumb reproduce all the characteristics of the stochastic equilibrium under the optimal policy, but there is no perceptible welfare difference. Figure 3 shows value functions for various policies, and that for the rule of thumb in Figures 1 and 2 is not visibly different from that for the optimal policy. (These value functions are calculated from the following procedure. Start from a grid on x, $\{xx_1, xx_2, \ldots, xx_N\}$ and corresponding grid on c, $\{c_1, c_2, \ldots, c_N\}$, with $c_i \equiv f(xx_i)$. Then the transition rule (7) is used to calculate a

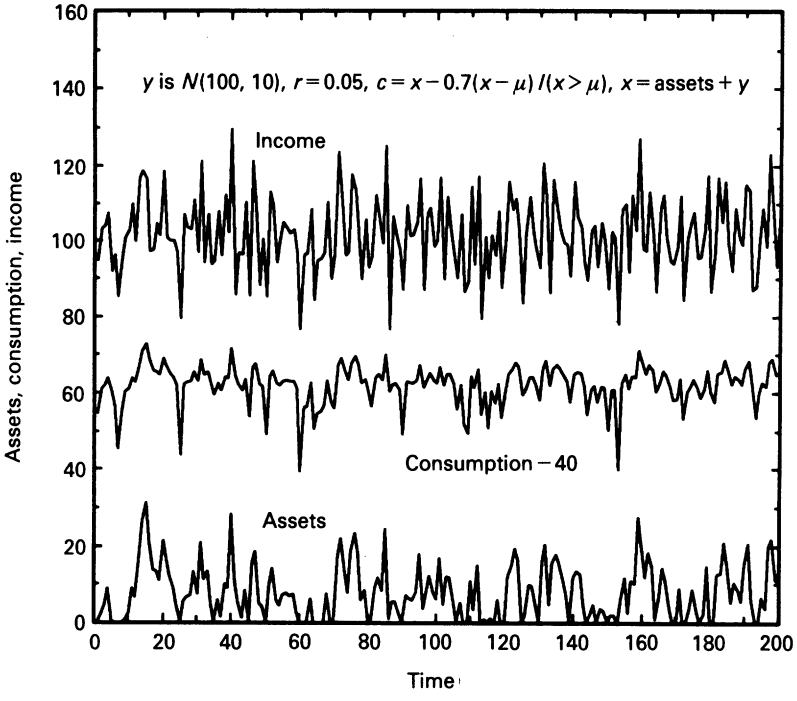

Fig. 2. Buffering without borrowing and rule of thumb behavior.

transition matrix T with element t_{ij}, the probability that, given x is xx_i in t, it is xx_j in $t+1$. We also have a grid of subutility values $v_i = v(c_i) = v\{f(xx_i)\}$, so that the expected utility associated with xx is the infinite sum of terms of the form $(1+\delta)^{-i}T^i v$ or $\{I-(1+\delta)^{-1}\ T\}^{-1}v$). Figure 3 also shows the value functions associated with a number of other possible consumption strategies. The next best after the optimal strategy and the rule-of-thumb is a second rule of thumb where, instead of saving 70 per cent of the excess above μ, only 50 per cent is saved, so that there is even more excess consumption than by the first rule. Quite some way below comes the value function that comes from the simplest policy of all, that associated with setting consumption equal to income and accumulating no assets. Finally, the worst policy shown is the ultra-conservative one of spending no more than mean income, and spending less if cash on hand is less. These rankings show that, at least for these parameters, the rules of thumb do well capturing most of the benefits of the more complicated optimal policy, and that any sensible rule of this form does a great deal better than the "obvious" policy of never saving anything. These results are not surprising given earlier results in the commodity storage literature,

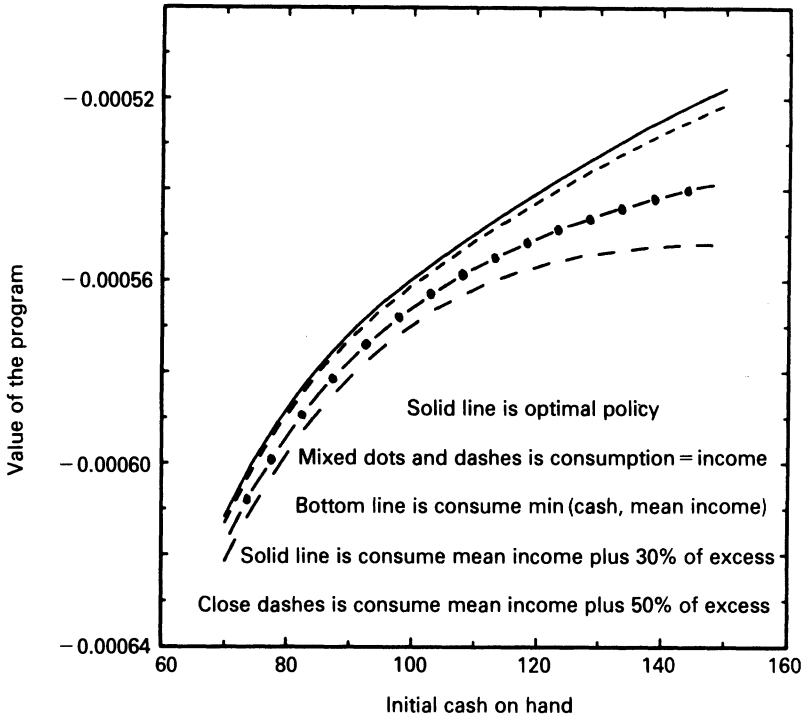

Fig. 3. Value functions for alternative policies.

where piecewise linear storage functions are known to perform well relative to fully optimal policies; see Newbery and Stiglitz (1981, Chapter 30).

One feature of these results that is perhaps surprising is the small values that are attached by consumers to any form of consumption smoothing. Figure 3 can be used to read off the amount that a "consumption equals income" consumer would pay to be instructed in the optimal buffering policy, and the answer is close to 10, which is one tenth of mean income, or one standard deviation of income. This amount is the *discounted present value* of the improvement, not a continuing per period gain, and seems a very small amount to pay for the consumption stream in Figure 2 rather than the income stream immediately above it. Of course, the agent who buffers consumption loses utility by holding assets, as well as gaining from the greater stability of his consumption and the net gain appears to be rather small. The low value attached to consumption smoothing is a standard result from other research, for example in the theory of commodity price stabilization; the continuing income equivalent of a reduction in variance is approximately the reduction in the coefficient of

variation multiplied by half the coefficient of relative risk aversion, an approximation that accords well with the results in Figure 3. But as has often been noted, these standard evaluations generate answers that reflect neither intuition nor the urgency with which policy-makers and agents in LDC's approach the stabilization issue.

It is also possible to calculate how much these consumers would pay to be rid of the uncertainty altogether, and how much to be rid of the liquidity constraints, In this context, it is the liquidity constraints that hurt, much more than the uncertainty. If there is no uncertainty, the optimal policy is derived from equation (5), but without the expectation. If we insert the isoelastic form $\lambda(c) = c^{-\rho}$, invert and take logarithms, we have

$$\ln c_t = \min[\ln x_t, \ln c_{t+1} - \rho^{-1}(r - \delta)]. \tag{9}$$

Since impatience dominates, $\delta > r$, so that, if initially, the consumer has assets, they will be run down along a declining consumption path until consumption equals income, at which value it will remain thereafter. The value of this policy is the middle line in Figure 4, higher than the value of the optimal policy with borrowing restrictions and uncertainty, but much lower than the function for the case where there are neither borrowing constraints nor uncertainty. If there are no borrowing constraints, then these impatient consumers will plan a falling consumption trajectory, with initial consumption set to satisfy the long-run budget constraint. In the case illustrated, which has the same parameters as before, initial consumption is 84 per cent higher than mean income, so that very substantial use is made of the borrowing facility, consumption is brought forward, and there are large utility gains. (I have been unable to calculate the fourth value function, for the case where there is uncertainty, but no restrictions on borrowing, but I conjecture that it lies not far below the top curve in Figure 4.)

In practice, while agents in LDC's do borrow and lend, it is unlikely that they would be able to borrow the very large sums that their impatience might dictate, and certainly not at rates that are the same as those at which they can lend. To consider a more realistic situation, I computed one further policy function, with the same parameters as before, but now allowing a limited amount of borrowing, up to 10, which is 10 per cent of income and one standard deviation, but at the "usurious" rate of 25 per cent. The Euler equation (5) is modified fairly straightforwardly to include four "regimes" instead of two. In the worst possible states, when cash on hand is very low, the maximum of 10 is borrowed from the moneylender, consumption is cash on hand plus 10, and consumption moves one for one with resources. This is the bottom segment of the function in Figure 1 above. As resources increase, less is borrowed from the moneylender until a segment is reached on which consumption is equal to cash on hand, and

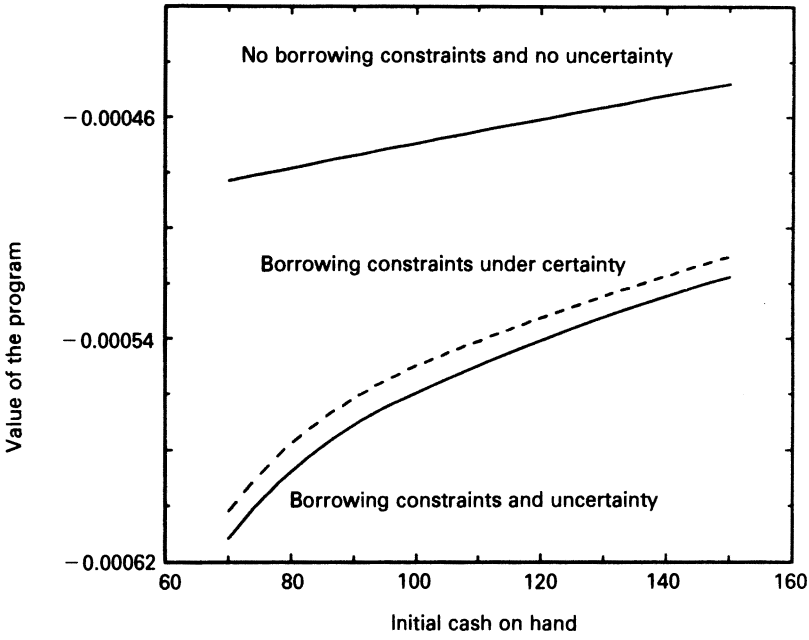

Fig. 4. The costs of uncertainty and borrowing constraints.

for this segment the policy function runs along the 45-degree line. At yet better positions, we get the same qualitative behavior as in the original model, with something being saved, and the original and "moneylender" policy functions asymptote to one another as the level of cash on hand becomes very large, and the probability of ever having to resort to the moneylender becomes correspondingly small.

Figures 5 and 6 illustrate the differences in behavior and value functions generated by the presence of the moneylender. Given the penal borrowing rate, it is not surprising that the moneylender is rarely used. However, there is a noticeable effect on the downward peaks in consumption, which are much less severe in the presence of the moneylender than without. Note too that the presence of the borrowing facility means that the consumer holds fewer assets; the precautionary need for assets is less when borrowing is available, even in limited amounts and at unattractive rates. The value functions hold no great surprises and once again, the evaluation of the gain in Figure 6 is small given what appears to be the marked usefulness of borrowing in Figure 5.

It is no simple matter to turn the theory of this section into fully articulated econometric specifications that could be confronted with the data. The lack of explicit functional forms for the optimal policy functions

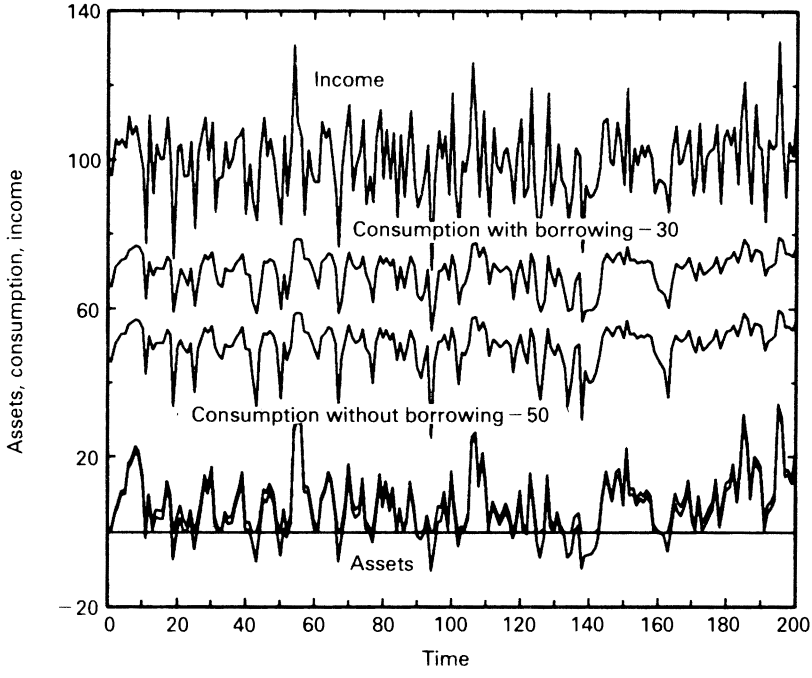

Fig. 5. Consumer behavior with and without a moneylender.

presents computational difficulties, which would be made worse if the income process were allowed to be temporally dependent, a generalization that would certainly be required by the data. It is also unclear whether the quality of the microeconomic data on income and saving is sufficient to support the estimation of complex non-linear models. At this stage, it seems better to use these models as a guide for data exploration, and that is the purpose of the empirical section below. Note, however, the broad implications of the sort of models discussed here:

(i) The dynamics of consumption and income are such as to detach consumption from income in the short-term, but not in the long-term. There is "high-frequency" consumption smoothing, but no "low-frequency" smoothing, over the life-cycle, nor over longer secular periods. The model is consistent with some lending and borrowing for buffering purposes, but not with the existence of long-term loans for consumption purposes. Agents look ahead in an entirely rational manner, but their horizons are naturally truncated at periods when they run out of assets, beyond which there is no point in planning. Even so, rational agents would save in anticipation of short-term falls in income, and vice versa.

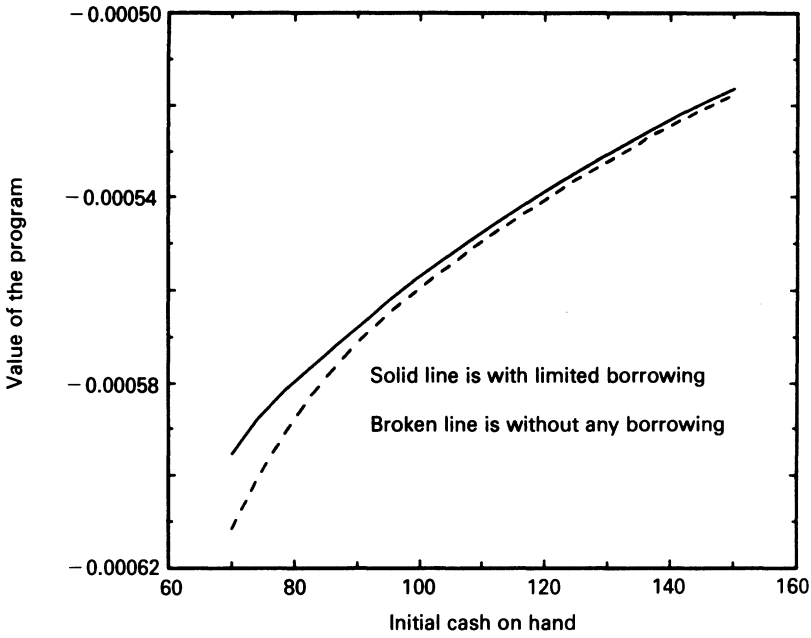

Fig. 6. The value of a moneylender.

(ii) In the cross-section, dissaving is likely to be common. If income is stationary, there will be as much dissaving as saving in the long run, and depending on the point in the agricultural cycle, large fractions of households may spend more than their income. The relationship between consumption and cash in hand depicted, for example, in Figure 1, will not be observed directly in the data, because income processes will differ from agent to agent. Abstracting from short-term buffering behavior, consumption will match income across agents, so that the implications of the theory for cross-sectional behavior are essentially identical to those of the standard permanent income theory, and consumption rises less rapidly than income because of the increasing proportion of transitory income as income rises.

(iii) Unlike the life-cycle model, buffering models do not provide an explanation for wealth holding in the society as a whole. There is no "hump" retirement saving and alternative explanations must be sought for the wealth that is actually observed. A crucial assumption in the buffering model is that consumers are *impatient*, so that assets are a necessary evil, lost consumption opportunities that must be held as insurance against an uncertain and dangerous future. In fact, it is reasonable to suppose that impatience is a characteristic that varies from individual to individual in the population, and that while most are too impatient to accumulate, there

are a minority of patient accumulators, who save to finance higher consumption later, either for themselves or for their heirs. In such a model, we have the opposite of Irving Fisher's contention that poverty generates high rates of time preference. It is high rates of time preference that make it optimal for agents to remain poor, and low rates of time preference that generate individual fortunes and long-lived dynasties. But patience, unlike wealth, is not easily bequeathable from parent to child, so that fortunes do not last for ever. To adapt the saying, families go from liquidity constraints to liquidity constraints in three generations. In this world, most individuals have little or no wealth, they smooth consumption in the short-run, but there are no long-standing deviations of consumption from income, and no room for accumulation. Wealth is owned by a small minority of individuals who are abnormally patient, who had access to abnormally large rates of return, or who were fortunate enough to inherit wealth from patient parents and grandparents who were in such a position. Whether or not this picture describes the United States is still being hotly debated; see Kotlikoff and Summers (1981), Modigliani (1988) and Kotlikoff (1988). That wealth transmission is governed by the sort of process discussed here has previously been argued by Bevan and Stiglitz (1979).

III. Empirical Evidence from LDCs

Patterns of Lending and Borrowing

I begin with evidence from West Africa. The World Bank, in collaboration with the Governments of Côte d'Ivoire and Ghana, has been collecting household data on a wide range of variables since 1985 in Côte d'Ivoire and 1987/8 in Ghana. The Living Standard Surveys (LSS) have relatively small samples by the standards of many household surveys in developing countries, but are distinguished by their range and depth, enquiring into the whole range of economic and social activities of the household members. Both data sets contain panel elements, with a half of the households surveyed in each year retained for one subsquent appearance in the next; the results here are based on the three years 1985, 1986, and 1987 from Côte d'Ivoire, and the single year 1987/8 from Ghana.

Table 1 summarizes the information on the extent of debts and credits in all three years of the Ivorian survey and the first year of the Ghanaian one. A little more than a third of the survey households on Côte d'Ivoire owe loans, while closer to a half have made loans that are still outstanding. Note that the same loans may be counted in more than one year, although two-thirds of loans are for durations of twelve months or less. Since there may be more than one lender or borrower in each household, and since there are lenders who are not private individuals (but 89 per cent of rural

Table 1. *Fractions of households with outstanding debts and credits*

	Rural		Urban	
	Creditors	Debtors	Creditors	Debtors
Côte d'Ivoire				
1985	0.474	0.399	0.487	0.396
1986	0.502	0.364	0.525	0.373
1987	0.441	0.386	0.484	0.333
Ghana				
1987/88	0.272	0.286	0.315	0.352

Notes: Creditors are those households who have at least one member who is a creditor for an outstanding loan, and debtors those households who have at least one member who have a loan outstanding. Even in the absence of measurement error, fractions of creditors and debtors need not be the same because there are lenders other than households, and because there can be more than one creditor and debtor in each household.

and 68 per cent of urban loans are from private individuals) the number of households who are creditors need not be the same as those who are debtors. Even so, it is surprising that there are so many more creditor households than debtor households in Côte d'Ivoire, and it may well be that respondents are more willing to report assets than liabilities. The Ghanaian figures are a little lower than the Ivorian ones, and there is a closer match between the numbers of creditors and debtors. Again, informal arrangements dominate the picture. Only about 7 per cent of loans in both urban and rural come from private banks, government banks, or cooperatives, about the same fraction from private moneylenders, and three quarters from private individuals, a third of whom are relatives among rural borrowers, and a fifth among urban borrowers. Very few loans carry interest payments or require collateral, and in most cases do not involve regular, pre-specified payments. There is clearly an active market in informal credit.

The survey data cannot be used directly to assess the role played by these loans in smoothing consumption. The size of the loans is not very large, especially in the rural sector, where the need for consumption smoothing would seem to be largest. For urban Ivorian households who have any loans at all, the average amount owed is around ten per cent of average consumption. For the forest zones, the ratio is closer to 5 per cent and for the Savannah it is only 1.5 per cent. For Ghana, the proportions are even smaller, between one and two per cent of average consumption. Moreover, not all loans are for consumption purposes. The survey distinguishes loans for farm, for business, for school, and for other purposes and the last category, which presumably includes borrowing for consumption, comprises little more than half of all loans in Côte d'Ivoire

Table 2. *Distribution of loans by season of contract and termination (percentage)*

	Season in which loan began		Season in which loan terminated	
	Rural	Urban	Rural	Urban
Côte d'Ivoire				
Off season	16.8	22.8	44.3	20.3
Planting season	13.0	13.8	3.2	9.9
Growing season	30.5	24.7	5.7	22.5
Harvest season	15.2	21.7	6.3	17.7
Missing	25.6	17.1	40.6	30.0
Ghana				
Off season	26.6	25.2	19.5	14.7
Planting season	18.3	23.1	11.5	9.4
Growing season	37.2	33.6	24.3	22.6
Harvest season	17.9	18.1	19.0	16.8
Missing	—	—	25.8	36.5

Notes: The Ivorian data are pooled over the three years 1985–87; those from Ghana are from the single year 1987/88. Since there are many different crops in Côte d'Ivoire and Ghana, the seasons correspond only loosely to actual agricultural seasons which vary from place to place and crop to crop. Here the off season is December, January, and February, the planting season is March and April, the growing season is May, June, July, and August, and the harvest season is September, October, and November. The figures are percentages of all loans in each sector by season. Missings include some genuinely missing observations, Côte d'Ivoire in the first two columns, but also loans that are still outstanding in the last two columns.

and 60–70 per cent in Ghana. These figures do not suggest that loans play a major role in consumption smoothing, although even such small amounts may be helpful, as for example in Figures 5 and 6 above.

Another way of looking at the role played by loans is to examine their seasonal pattern, and the way in which loans are linked to the agricultural calendar. Table 2 shows the distribution of loans contracted and terminated by seasons, defined in terms of calendar months as detailed in the notes to the table. In the urban sectors of both countries, loan contracts and terminations are relatively evenly distributed throughout the year, but this is much less true in the rural areas, particularly for loan repayments. The χ^2 tests that the two distributions are the same are, for Côte d'Ivoire, 49.1 for contracts and 370.3 for terminations, and for Ghana, 6.0 for contracts, and 22.6 for terminations, all with three degrees of freedom. In Côte d'Ivoire, a large fraction (40 per cent or so) of loans in the rural sectors are taken out in the growing season, in the months before the harvest, and are repaid in the off-season, after the harvest (75 per cent). The pattern is much less clear in Ghana, but in general is quite consistent with loans being used to smooth consumption, even though the amounts

are quite small. However, if such was the case we should expect changes in income between one year and the next to be positively associated with the net credit position (loans owed to less loans owed by) of the household. This can be tested in Côte d'Ivoire but not on the single year of data in Ghana. The predicted correlation exists in the 1985–86 panel, but is replaced by a *negative* correlation in the subsequent 1986–97 panel perhaps because of a different pattern of common shocks in the two years. In spite of the difficulties, the data from these two countries are perhaps consistent with a modest role for credit markets in smoothing consumption and as such are consistent with Udry's (1990) evidence from northern Nigeria.

Note finally that if borrowing restrictions are a problem for many households, assets, particularly liquid assets, should play a role in guaranteeing future consumption. A simple way of testing the role of these cash balances is to regress consumption in year t on assets in year $t-1$, including consumption and income in $t-1$ in an attempt to control for individual fixed effects. Lagged income has no predictive power for consumption conditional on lagged consumption, so the role of assets can be assessed by regressing consumption on its lag and on the amount of liquid assets in hand at the end of the previous period, nearly always cash, since few households report holdings of other financial assets. In the 1985–86 panel, this experiment generates a significant positive role for lagged assets in predicting consumption in the rural zones, but not in Abidjan or in other urban areas, but the 1986–87 panel shows no significant results for any of the five regions. Again we have the uncomfortable fact that the two panel data sets do not have the same structure.

Life-Cycle Patterns of Consumption and Income

If a large fraction of individuals are impatient and liquidity constrained, their consumption will track their incomes fairly closely over time. In cross-sectional data, there typically is such a close relationship, and Figures 7 and 8 show plots of average household consumption and household income by age of household head using household survey data for Côte d'Ivoire for 1985, 1986, and 1987, and for Thailand in 1986, distinguished by level of urbanization (sanitary districts are intermediate between urban (municipal areas) and rural (village) Thailand). The data have been smoothed over age groups, with each point showing averages for the age and the two (Thailand) or three (Côte d'Ivoire) ages on either side, with declining triangular weights. The Thai sample is larger, so it needs less smoothing, but five points are still required because people tend to report rounded ages (ending in 5's and 0's) and because those who do so have lower average consumption and income levels.

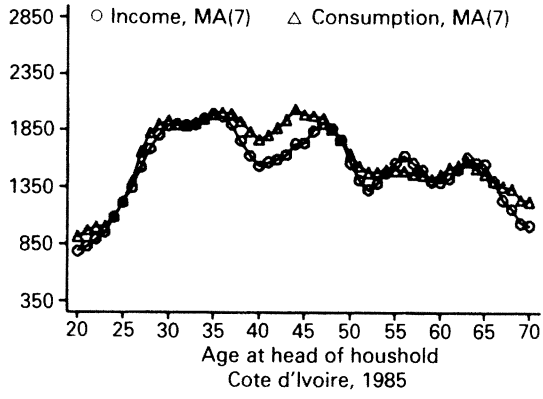

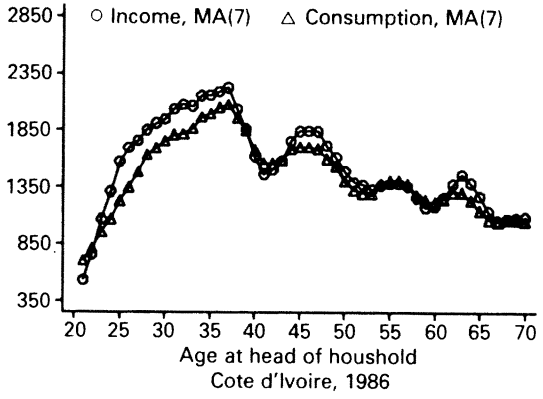

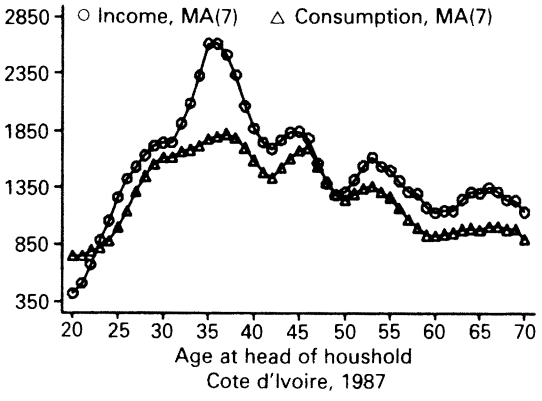

Fig. 7. Consumption, income and age: Côte d'Ivoire.

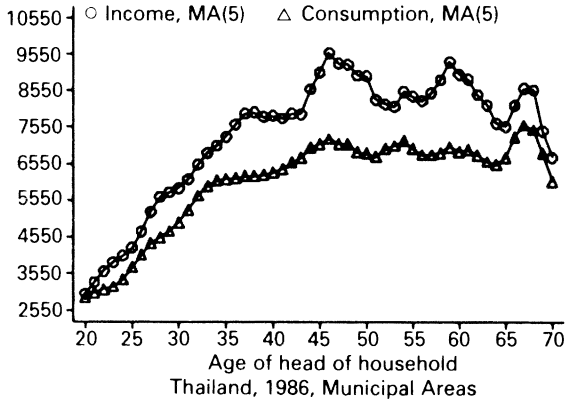

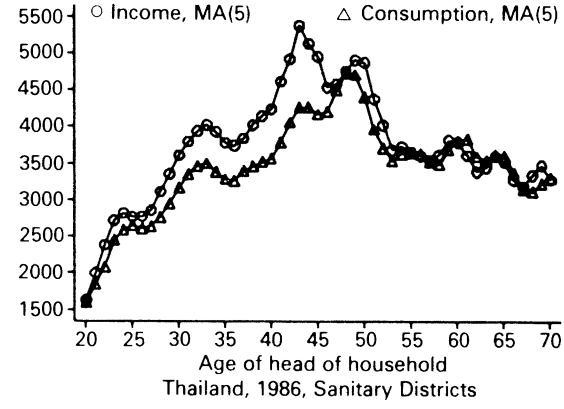

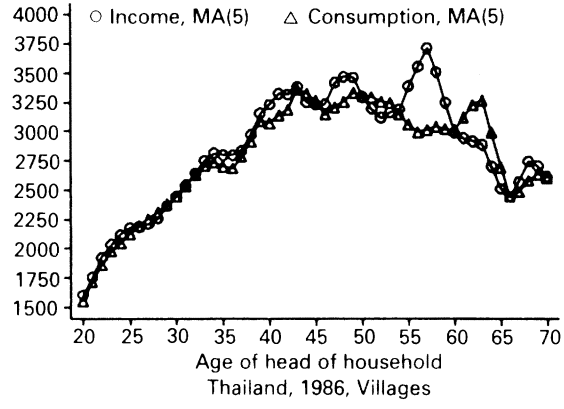

Fig. 8. Consumption, income and age: Thailand.

Although there is some evidence of "hump" saving in urban Thailand, and to a lesser extent in the sanitary districts and in Côte d'Ivoire in 1987, there is generally a close relationship between the evolution with age of consumption and income. Indeed, the relationship is even closer when purchases of durable goods are included in consumption. While such evidence casts doubt on the importance of life-cycle saving in the generation of national wealth, cross-sectional evidence is hardly conclusive. These households may have saved at some time in the past, or will do so at some time in the future, and we cannot tell how that saving may evolve with age. Nor can this evidence, by itself, by used to cast doubt on the central life-cycle proposition that consumption is determined by life-time wealth and the pattern of preferences, not by the life-time evolution of income. Tastes may vary systematically with age in a manner that is correlated with income. Indeed, as has long been recognized, the explanation is a plausible one, because the evolution of family size and responsibilities is similar to that of income.

However, the combination of the Thai and Ivorian evidence does provide evidence against the life-cycle hypothesis, and in favor of the proposition that consumption tracks income. The argument is due to Carroll and Summers (1991). According to Summers and Heston (1988), the Thai economy has been growing at about 4 per cent per annum over the last 25 years, while the Ivorian economy has grown at a little less than 1 per cent per annum over the same period. If such growth rates are reflected in individual incomes, and if they are expected to continue, then a 25 year old in Côte d'Ivoire has expected total lifetime resources that are two-thirds larger than those of his or her 75 year-old grandparents. For Thailand, the 25 year old is 7.1 times better off than his or her grandparent. Hence, if consumption is determined by life-time resources, consumption by age patterns should be more tipped in favour of young households in Thailand and more tipped in favor of the old in Côte d'Ivoire. Precisely the opposite is true. Consumption peaks at a much earlier age in the slow-growing Côte d'Ivoire than it does in the rapidly growing Thailand. Of course, tastes and preferences may be different in the two countries, but there is a much simpler explanation, which is that the patterns of income are different in the two countries, and consumption tracks income. Indeed, if the taste explanation is to be relied upon, then family size is not the relevant variable, because family size peaks much later in Côte d'Ivoire, rising steadily from head's age 25 to 50 or later, than in Thailand, where the peak is reached by age 40.

In Figure 8, for Thailand, note that the vertical scales for the three regions are not the same, and that the range of variation of life-time consumption in the municipal areas, from 2550 to 7550 baht, is much larger than the range in the villages, from 1500 to 3250 baht. Again, the

difference is more plausibly linked to the similar paths in income than to any difference in taste variation between urban and rural regions. It is also reasonable to suppose that incomes in urban areas have been growing more rapidly over time than those in the villages. Income differences between the two areas are very large, Bangkok is a rapidly growing city that accounts for much of Thailand's urban population, and there is a limit to the rate of growth of agricultural productivity. Yet there is no evidence in Figure 8 that consumption patterns are relatively tipped towards the young in the urban areas. Once again, the simplest hypothesis, that consumption tracks income, is the most plausible. Even if this is not the case, and the life-cycle hypothesis works in each country individually, with the tracking of income by consumption largely coincidental, the comparison of the two countries must reject any model that relies on common tastes across countries, as does the standard life-cycle explanation for the positive international correlation between saving and growth.

IV. Conclusions

The results of this study should be read in conjunction with the evolving literature on saving in developing countries, and are placed in more extensive context in the full version of this paper, Deaton (1991b). My own earlier work for the Côte d'Ivoire, Deaton (1990b), and Paxson's (1991), (1992) work for Thailand, confirm some but not all aspects of life-cycle theory. In particular, these papers provide evidence that farmers (and others) look ahead at least some way when deciding how much to consume and to save, but they do not support the cross-equation restrictions generated either by the full permanent income hypothesis, nor by any simple modification based on the way liquidity constraints might work. More general reviews of the undoubted progress in the area can be found in Gersovitz (1988) and Deaton (1990a). However, it will also be clear that much remains to be done. Although the data and the models generate what appear to be real insights, we are far from a really satisfactory understanding of all the evidence. The data themselves are difficult to use. Income is not a concept that is easily measured, particularly for people who are self-employed (in agriculture or elsewhere), and it is in the analysis of saving that data deficiencies are probably at their most acute. Except for the fact that there are more self-employed workers in LDC's, I do not believe that saving data are any worse in developing than in developed economies; indeed, high response rates and the willingness of "fresh" interviewees to spend a great deal of time with survey staff suggests the reverse. Even so, there is a great deal of measurement error in both income and saving, and the measurement error in each is almost certainly correlated, facts that must be recognized in any credible econometric analysis. The theories

examined in the paper are useful, but none is strongly supported by the evidence. The model without borrowing is a useful benchmark and shows that it is possible for consumption to be very much smoother than income even without the intermediation of credit markets. However, there are important elements of reality that are not captured by such a formulation, and it is difficult to use the model to derive empirical predictions that can be readily tested on the data.

References

Bevan, D. & Stiglitz, J. E.: Intergenerational transfers and inequality. *Greek Economic Review 1,* 8-26, 1979.

Carrol, C. & Summers, L. H.: Consumption growth parallels income growth: Some new evidence. In B. D. Bernheim & J. B. Shoven (eds.), *National Saving and Economic Performance.* Chicago University Press for NBER, Chicago and London. 305-47, 1991.

Deaton, A. S.: Saving in developing countries: Theory and review. *Proceedings of the World Bank Annual Conference on Development Economics 1989; The World Bank Economic Review,* Washington, DC, 1990a.

Deaton, A. S.: Saving and income smoothing in Côte d'Ivoire. Research Program in Development Studies, Princeton University (processed), 1990b.

Deaton, A. S.: Saving and liquidity constraints. *Econometrica 59,* 1221-48, 1991a.

Deaton, A. S.: Household saving in LDC's: Credit markets, insurance, and welfare. Research Program in Development Studies WP No. 153, Princeton University, 1991b.

Gersovitz, M.: Saving and development. In Hollis Chenery & T. N. Srinivasan (ed.), *Handbook of Development Economics 1,* North-Holland, Amsterdam, 381-424, 1988.

Kotlikoff, L. J.: Intergenerational transfer of saving. *Journal of Economic Perspectives 2,* 41-58, 1988.

Kotlikoff, L. J. & Summers, L. H.: The role of intergenerational transfers in aggregate capital accumulation. *Journal of Political Economy 89,* 706-732, 1981.

Little, I. M. D. & Mirrlees, J. A.: Project Appraisal and Planning. Heinemann, London, 1974.

Modigliani, F.: The role of intergenerational transfers and life-cycle saving in the accumulation of wealth. *Journal of Economic Perspectives 2,* 15-40, 1988.

Newbery, D. M. G. & Stiglitz, J. E.: *The Theory of Commodity Price Stabilization.* Oxford University Press, Oxford, 1981.

Paxson, C. H.: Consumption seasonality and consumption smoothing. Research Program in Development Studies WP No. 150, Princeton University, 1991.

Paxson, C. H.: Using weather variability to estimate the response of savings to transitory income in Thailand. *American Economic Review 82,* 15-33, 1992.

Squire, L. & van der Tak, H. G.: *Economic Analysis of Projects.* Johns Hopkins University Press, Baltimore, 1975.

Summers, R. & Heston, A.: A new set of international comparisons of real product and prices for 130 countries. *Review of Income and Wealth 34,* 1-26, 1988.

Udry, C.: Credit markets in Northern Nigeria: Credit as insurance in a rural economy. *World Bank Economic Review 4,* 251-69, 1990.

UNIDO (United Nations Industrial Development Organization) *Guidelines for Project Evaluation.* Project Formulation and Evaluation Series, No. 2, United Nations, New York, 1972.