On risk, insurance, and intra-village consumption smoothing

Angus Deaton

Research Program in Development Studies
Princeton University

PRELIMINARY DRAFT
November 1990

I am grateful to Kristin Butcher for assistance in the preparation of this paper. I am grateful to the World Bank and to the Lynde and Harry Bradley Foundation for support.

INTRODUCTION

For the many people in developing countries whose incomes are derived from agriculture, fluctuations in income present a constant threat to living standards, even if on average, incomes are sufficient to guarantee at least a minimum level of consumption. There are many ways in which consumption can be protected against variations in income. When they exist, credit markets provide a means to smooth out incomes, and when they do not, people can build up stocks of money and commodities during good times, and run them down in bad. In addition to these private responses, there are a range of public actions that can protect individual living standards. Commodity price stabilization is sometimes justified on these grounds. Rationing, "fair-price" shops, and food- or cash-for-work schemes are other examples. However, there is also the possibility that some consumption stabilization takes place at the *local* level. Individuals or families within villages may share the risk involved in their activities through formal or informal schemes that diversify some proportion of their individual income risk. In many villages, the inhabitants are well-known to one another, and there is a great deal of common information about endowments and about the random events that affect agricultural activities and incomes. Such detailed local information is perhaps sufficient to allow the village as a whole to insure individual living standards without the problems of moral hazard that would beset any outside insurer. While the remaining informational and incentive effects make it unlikely that the complete diversification of individual risk would be either possible or desirable, the existence of some village-level risk spreading is worth investigating. Indeed, it has recently been argued by Townsend (1989, 1990) that even the extreme case of perfect diversification is given empirical support by the evidence from the ICRISAT villages in southern India.

This paper uses panel data from the 1985, 1986, and 1987 waves of the Living Standard Survey to look for evidence of risk pooling within villages in Côte d'Ivoire. The tests are extremely simple; I look at consumption changes from one year to the next and ask whether there are common village-level patterns in the changes for the individual households. Such effects are predicted by either total or partial risk sharing, since households who do better than the village average in any year will give up some (or all) of their incomes to support the living standards of their unluckier neighbors. I also test for complete diversification of risk, which implies, among other things, that individual consumption is unresponsive to individual income changes once the village effects are controlled for. I argue that the simplest way of addressing these questions is the straightforward procedure of calculating F-tests for village effects in regressions of consumption change on income change.

As noted by Altonji, Hayashi, and Kotlikoff (1989) in their related work on altruism, there is a formal similarity between standard permanent income models, in which income is spread over time by a single individual, and perfect diversification, where income is spread over individuals at a single moment of time. The object of the exercise is the same in both cases, to protect living standards against fluctuations in income. In the permanent income model, consumption is determined by the typically unobservable average of individual incomes over time, while in the risk-pooling model, consumption is determined by the (again) typically unobservable average of incomes over the members of the village. A standard test for the permanent income model is whether consumption responds to individual income by more or less than is warranted by its effect on permanent income, whereas perfect risk sharing can be tested by asking whether individual consumption responds to individual income changes other than through their effect on group income. In the permanent income literature, it is often found that consumption responds to income by more than it should, given the pure permanent income hypothesis, see for example Flavin (1981), the survey by Hayashi (1987), as well as Deaton (1990) for results from Côte d'Ivoire. Using data for American households, Altonji, Hayashi and Kotlikoff (1989) find that consumption is responsive to individual family income even when they control for the consumption levels of the extended family or dynasty. Testing for this sort of excess sensitivity at the village level is one of the objects of this paper.

1. MODELS AND EMPIRICAL STRATEGY

1.1 A simple methodology

A convenient starting point is an equation for individual consumption under perfect risk-pooling. Let i indicate the individual or household, v the village, and t the year, so that c_{ivt} is consumption by household i in village v in year t. Under perfect risk-sharing, individual consumption is a function only of aggregate or average village consumption, for example,

$$c_{ivt} = \alpha_{iv} + \beta_{iv}\mu_{vt} + e_{ivt} \quad (1)$$

where α_{iv} and β_{iv} are sharing parameters that are specific to the household, μ_{vt} is the mean consumption in the village, and e_{ivt} is an "error" term, the interpretation of which will be elucidated below. The functional form of the sharing rule, as well as the values of the parameters, depend on individual preferences and on the way in which the village aggregates preferences. Townsend (1989) derives (1) explicitly, but there is nothing sacred about the linear case. Since (1) is a *sharing* rule, the α parameters must sum to zero, and the β-parameters average to one; taking expectations of (1) conditional on the village v we have at once:

$$E(\alpha_{iv}) = 0 \quad E(\beta_{iv}) = 1. \quad (2)$$

Strictly interpreted, perfect risk pooling would imply that (1) holds *without error;* consumption levels within each cluster are perfectly synchronized. Even within this interpretation, the presence of the error terms could be justified as measurement error in consumption totals. However, without further restrictions on either the parameters or on the properties of the error terms, equation (1) has no empirical content. If the error terms were zero, the equation would be testable on a panel of three or more periods, since the ratios of consumption changes for two households in the same village should be invariant over time. As it is, each household in the Ivorian data is observed for only two years before being replaced in the panel, so that this test is not available to me. More useful in the current context is the implication that individual household incomes do not appear in (1) except in so far as they contribute to μ_{vt}. Without explicitly cataloging the properties of the error terms, it is natural to turn this implication into a testable hypothesis by asserting that, under complete diversification, individual incomes are orthogonal to e_{ivt}.

The first difference of (1) can be written

$$\Delta c_{ivt} = \beta_{iv}\Delta\mu_{vt} + \Delta e_{ivt}, \quad (3)$$

which, since the mean of the β_{iv}'s is unity, is equivalent to

$$(\Delta c_{ivt} - \Delta\mu_{vt}) = \varepsilon_{ivt} - \Delta e_{ivt} + (\beta_i - 1)\Delta\mu_{vt}. \tag{4}$$

The village means are typically not observable from survey data. In the Ivorian survey, households are grouped by village but there are no more than 16 sample households in each, so that μ_{vt} cannot be treated as known. Instead, I estimate the following model:

$$\Delta c_{iv} = \sum_{j=1}^{V} \gamma_j \delta_{vj} + b\,\Delta y_{iv} + \varepsilon_{iv}, \tag{5}$$

where δ_{vj} is a dummy variable equal to 1 if $v=j$ and zero otherwise, V is the total number of villages, and I have dropped the time subscript since there is only one set of changes in my data. The parameter b should be zero if there is complete diversification of risk; it is an "excess sensitivity" parameter comparable to that in the permanent income literature. Note from (4) that, in order to obtain a consistent estimate of b, the sharing rule, and in particular the parameters β_i, should be independent of *changes* in income. This may or may not be true, but some such assumption must be made, and it is much more plausible than the assumption that the sharing rules are independent of *levels* of income.

My procedure is therefore as follows. I estimate equation (5) and test the hypothesis that b is zero. If this fails, complete risk sharing within the village is rejected. The alternative hypothesis, that b is non-zero and presumably positive, has various possible interpretations. There could be partial insurance within the village, as would be the case if a linear tax system were operating. Consumption would be partly determined by own income, and partly by the average income of the village. In such a case, one might expect the villages dummies to be important, especially if mean income changes vary significantly from village to village. Indeed, if there are significant village level effects in income changes, and if risk is being totally or partially diversified within each village, the regression of consumption change on income change will generate an upwardly biased coefficient, which will be reduced when the dummies are introduced. I shall also test whether this is the case, but note the requirement that there be village effects in income changes. Without this, the insignificance of village dummies in (5) cannot be interpreted as evidence against risk pooling. Note also that village level dummies could be important for reasons that are unrelated to risk sharing. For example, households may be following a purely individualistic permanent income rule for consumption, but there may be news about future incomes that effects all members of the village, for example that a road is to be built, or an irrigation project begun. The only strict test here is for b=0, but much can also be learned from the patterns revealed by the other coefficients.

1.2 Alternative econometric procedures

This section, which deals with other possible estimation and testing strategies, is self-contained, and can be skipped without loss of continuity. Return to equation (1), reproduced here for convenience:

$$c_{ivt} = \alpha_{iv} + \beta_{iv}\mu_{vt} + e_{ivt} \tag{6}$$

then, given an estimate of μ_{vt} and time-series data on individual households, estimates of α_i and β_i could be obtained by running separate regressions for each household. This is the situation with the ICRISAT data used by Townsend who replaces μ_{vt} by the sample mean for the village computed excluding the consumption level of the ith household. Hence, the regression is

$$c_{ivt} = \alpha_{iv} + \beta_{iv}\mu_{vt}^{(i)} + e_{ivt}, \tag{7}$$

where the "leave out" mean $\mu_{vt}^{(i)}$ is calculated according to

$$\mu_{vt}^{(i)} = \frac{1}{n_v - 1}\sum_{j \neq i} c_{jvt} = \frac{n_v}{n_v - 1}\bar{c}_{vt} - \frac{1}{n_v - 1}c_{ivt} \tag{8}$$

where $\bar{c}_{vt}$ is the sample mean of consumption in the village. Equation (8) can be rewritten as

$$\mu_{vt}^{(i)} = \mu_{vt} + \frac{n_v}{n_v - 1}(\bar{c}_{vt} - \mu_{vt}) - \frac{1}{n_v - 1}(c_{ivt} - \mu_{vt}) \tag{9}$$

which shows that the leave-out mean can be regarded as the true mean plus measurement error. As the number of sample points in the village becomes large, $\mu_{vt}^{(i)}$ converges to μ_{vt} and the parameter estimates will be consistent. Since the ICRISAT village samples contain 30-40 households each, this is perhaps a practical estimator, although for any finite number of households in the village, the leave out mean is an error-ridden estimate of the true mean, and there will be attenuation bias according to (9), as in the standard errors in variables model.

Suppose for moment that the village sample is large enough so that these issues are unimportant. Then the estimated β-parameters in (1) or (6) are given by:

$$\tilde{\beta}_{iv} = cov(c_{ivt}\,\mu_{vt}) / var(\mu_{vt}), \tag{10}$$

where the variances and covariances are taken over time. But the village average of c_{ivt} is μ_{vt}, so that averaging (10) over the village gives, corresponding to (2),

$$n_v^{-1}\sum_{i \in v}\tilde{\beta}_{iv} = 1. \tag{11}$$

The important thing to note about (11) is that the result is purely *mechanical*; it is always true, irrespective of the nature of the data, the sample being used, or whether there is in fact any risk sharing in the village. As a consequence, there is no evidence for or against risk sharing to be gleaned from the fact that the estimates of β_i average to unity. In practice, the attenuation bias is likely to make the sum less than unity, and must do so, if as is to be expected, the individual $\tilde{\beta}_{iv}$'s are positive.

For the Ivorian data, with only two observations per household, these techniques are not feasible, but there are immediate analogs. Write the first difference equation (4) in the form:

$$\Delta c_{ivt} = \beta\,\Delta\mu_{vt} + \varepsilon_{it}, \tag{12}$$

where β is the mean of the β_i's, theoretically equal to unity. We might be tempted to run (12) as a regression and use the estimated β to measure the extent to which individual consumption changes move with the village means. Note first that if (12) is estimated with sample village means in place of $\Delta\mu_{vt}$, the OLS estimates are identically equal to unity, as simple algebra quickly shows. Once again, it is possible to use the leave out mean given by (8) applied to changes instead of levels. Now the calculation is a little more tedious, but it is straightforward to show that

$$\tilde{\beta} = 1 - \frac{\sum_v (n_v - 1)^{-1} \hat{\sigma}_v^2}{\sum_v n_v (\Delta \bar{c}_v - \Delta \bar{c})^2 + \sum_v (n_v - 1)^{-1} \hat{\sigma}_v^2} \tag{13}$$

where $\Delta \bar{c}$ is the mean consumption change over the whole sample, and $\hat{\sigma}_v^2$ is the estimate of the within village variance in the consumption change for village v, *i.e.*

$$\hat{\sigma}_v^2 = \sum_i (n_v - 1)^{-1} (\Delta c_{iv} - \Delta \bar{c}_v)^2 \tag{14}$$

Once again, the estimate of β converges to unity as the village sample sizes go to infinity, *irrespective* of whether there is risk pooling, or indeed of whatever is actually generating the data. However, for the Ivorian case, the village sample size is 16 or so, and the within village variances are large relative to the between variances, so that estimates of β are much closer to zero than to unity. Unless we take risk pooling to imply that all consumption changes within each village should be the same, estimates of β near zero are no evidence against the hypothesis, just as estimates averaging unity in the time series data are no evidence in its favor. However, it is clear from (13) that for finite cluster size, the estimate of β is based on much the same information as is the F-statistic, *i.e.* on the relationship between the within and between village variances in consumption changes. The F-test for village effects is given by

$$F(V-1, n-V+1) = \frac{\sum_v n_v (\Delta \bar{c}_v - \Delta \bar{c})^2 / (V-1)}{\{\sum_v (n_v - 1) \hat{\sigma}_v^2\} / (n-V+1)} \tag{15}$$

so that, in the case of a balanced sample of n_v households from each of V villages, we have the monotone relationship

$$\tilde{\beta} = \frac{(n_v - 1)^2 (V-1) F}{(n_v - 1)^2 (V-1) F + (n - V + 1)}. \tag{16}$$

Since the estimate of β and the F-statistic convey the same information, and since the latter is familiar and well-understood, it seems advisable to use it.

3. EMPIRICAL RESULTS

The data used here come from the 1985, 1986, and 1987 Living Standards Surveys of Côte d'Ivoire, described in detail in Ainsworth and Muñoz, (1986). In principle, there are 1600 sample households in each year, 800 of whom are resampled in the next year, but who are then dropped in favor of 800 fresh draws in the year after that. There are 100 sample clusters in all, 43 of which are in urban areas, so that there are 57 clusters corresponding to rural villages, each of which is designed to contain 16 households. There are 31 village clusters that are common to the 1985 and 1986 surveys, and which provide the first panel. In 1986 and 1987 there are 26 common panel clusters; these are not the same villages as made up the first panel. The structure of the sample is shown in Table 1. These figures are for the households actually used in the analysis below; inevitably, some households have been excluded. The majority are households that produced no useable data but a few others were excluded because the data on income or consumption were implausible either in levels or in changes. However, consumption and income totals are the sums of a large number of components, more than a hundred for consumption, and several hundred for income, and for all but a very few cases it was possible to repair the estimates by identifying and replacing a component. There are three regions of the country shown in the table. The West and East Forest comprise the tropical south of the country. Both areas derive much of their incomes from cocoa and coffee production, while the northern Savannah depends on rainfed agriculture, with cotton and rice the main cash crops.

Table 2 shows estimates of income and consumption for the panel households in each the years covered by the surveys. There are two sets of estimates for 1986 because it appears in both of the panels, but the two estimates come from different households in different clusters. The consumption estimates are almost certainly a good deal more reliable than those for income. The former are derived from the answers to a large number of detailed questions about individual non-durable goods and services, together with imputations at local market prices of home produced food. While there are almost certainly errors in these estimates, the concepts involved are clear enough. This is not true for income, where very few rural households earn wages, almost the only category of income that is conceptually straightforward. For the others, the calculation of income involves the construction of a set of household accounts for farm

Table I: Structure of the sample, Côte d'Ivoire 1985-86 and 1986-87
(Numbers of households or clusters)

	1985-1986 panel		1986-1987 panel	
	clusters	households	clusters	households
West Forest	11*	150	5	77
East Forest	12	181	11	165
Savannah	8	123	10	150
All rural	31	454	26	392

* One cluster contains only a single household

production and other self-employment activity. Allowances have to be made for wear and tear on tools and equipment, as well as for changes in the value of livestock and stocks of commodities. All of this leaves an uncomfortably large margin for judgment, and it is always unclear whether the estimated magnitudes correspond to any decision variable that would be recognized by the household. Even so, for the purposes of this

Table II: Income and consumption estimates, 1985, 1986, and 1987
CFA '000, annual rates, household averages

	1985		1986(i)		1986(ii)		1987	
	c	*y*	*c*	*y*	*c*	*y*	*c*	*y*
W Forest	1188	1098	903	790	1026	1225	893	1151
E Forest	1005	896	1042	1056	1147	1283	1048	1220
Savannah	671	630	828	788	770	742	659	716
All	975	891	938	896	977	1061	867	1011

paper, the precise measurement of income is perhaps of limited importance, provided only that it captures some of the idiosyncratic resource flows to the household.

The Savannah is the poorest of the three areas; the relatively well-off coffee and cocoa farmers reside in the Forest. The first panel shows little change in average consumption or average income, although there is a very large drop in both in the West Forest. Other information from Côte d'Ivoire provides neither explanation nor independent confirmation for this regional phenomenon, and the households in the second panel on the right of the table record higher incomes in the same year. However, the second panel contains only 77 households in 5 villages in the West Forest, and the discrepancies may reflect little more than the combination of small samples and large measurement error.

Changes in incomes and consumption are presented in Table 3, together with F-tests for the significance of cluster effects in these changes. The statistics for income changes are generally low; common village components explain very little of the variation in individual income changes. In the 1985-6 panel, the F for the West Forest has an associated p-value of 0.03, and that for the rural areas as a whole one of 0.0002, albeit from a much larger sample size. Clearly, a few of the West Forest villages did worse than others in 1986, and it is those same income falls in one region that generate the F-statistic for all regions. But these effects, as well as those for 1986-87, are much smaller than I had expected before looking at the data. None of the F-statistics are close to the log(sample size) critical value that asymptotically balances Type I and Type II errors, Schwartz (1978), a value that is frequently exceeded by tests for various sorts of cluster effects in household survey data. Nor is it easy to explain these results by appealing to measurement error. If, for example, the observed within village variation in income changes were composed of equal parts of measurement error and "real" change, the observed F-statistics would be a little more than half what they ought to be. But doubling the figures in the top half of the table still does not generate very impressive values.

The literature on agricultural risk in developing countries places great emphasis on the covariation between income risks for different farmers in the same location, and the consequent desirability of risk-pooling schemes that link the village to different

Table III: Income and consumption changes 1985-6 and 1986-7
CFA '000, household averages

	1985-86			1986-87		
Income	Δy	*F*	*df*	Δy	*F*	*df*
W Forest	-308	2.04	10,139	-74	1.93	4,72
E Forest	160	1.28	11,169	-64	2.05	10,154
Savannah	159	1.52	7,115	-26	1.17	9,144
All	5	2.30	30,423	-51	1.65	25,370
Consumption	Δc	*F*	*df*	Δc	*F*	*df*
W Forest	-285	4.10	10,139	-133	4.04	4,72
E Forest	37	4.15	11,169	-99	2.16	10,154
Savannah	157	1.35	7,115	-111	0.70	9,114
All	-37	4.62	30,423	-110	2.14	25,370

agricultural zones or to urban areas. There is little evidence of such covariation in these data, even though they span at least two quite distinct agricultural zones, the forest and the savannah. Côte d'Ivoire may be an exception to the general pattern, or these two years may be atypical in some ways. One possibility is that there is little covariance in normal years, but that there are infrequent but important events, such as fires or droughts, that affect everyone together, and that it is these rare events that shape risk bearing strategy. Even so, crops in Côte d'Ivoire are certainly affected by weather conditions, and fires in the cocoa and coffee growing areas are frequent occurrences, so that the absence of significant village level effects remains something of a puzzle. It is also possible that village effects do indeed exist, but are swamped by intra-village variation from one household to another.

From the point of view of this paper, the absence of strong village effects in income changes is unfortunate since it makes it more difficult to test for the presence of risk sharing in consumption. In the presence of village effects, the response of consumption to income changes should be larger in the absence of village dummies than in their presence, and this prediction cannot be tested with any power with the current data. Even so, inspection of the lower half of Table 3 reveals that the tests for village effects in consumption changes are typically much larger than those for income changes. Although there is little evidence of village effects in the Savannah, there are strong effects in the East and West Forest regions in 1985-86, and in the West Forest in 1986-87. Some factor is causing consumption levels within each village to move more closely together than do income levels. Risk pooling is certainly one possibility. Another is that consumption is measured more accurately than income, and that for correctly measured magnitudes, the F-statistics would be much closer. Nor is it difficult to think of economic explanations for intra-village comovement in consumption that do not involve risk sharing. What is clear is that the comovement does indeed exist.

Note that Table 3 shows a close relationship across regions and over time between consumption change and income change. Table 4 explores these connections more closely. The top panel of the Table gives estimates for 1985-86, and the bottom panel for 1986-87. The first row in each case shows that for each of the regions the ordinary least squares regression of the change on consumption on the change in income generates a significant positive coefficient. These values differ a good deal from region to region, and from the first panel to the second, but the coefficient is always positive and the t-values are large, ranging from 3.2 to 8.8. The second rows for each of the panels show that the introduction of village dummies makes little difference to these coefficients. The dummies themselves are often significant, see the third rows, but these F-statistics are much the same as in the previous Table. As far as these regressions are concerned, income changes and the village effects are essentially orthogonal. There is certainly no suggestion that the introduction of the village effects reduces the income coefficient at all, let alone to zero as required by complete risk sharing.

For many of these rural households a significant share of income and consumption is accounted for by food that is produced and consumed by the household and neither sold nor bought in the market. A value is imputed to this home-produced food

Table IV: OLS and IV estimates of marginal propensities to consume
(Absolute t-values or asymptotic t-values in brackets)

	West Forest		East Forest		Savannah		All Rural	
1985-86								
OLS: mpc without village effects, with village effects, and F-test for effects:-								
no dummies	0.290	(6.2)	0.153	(3.2)	0.368	(5.8)	0.259	(8.8)
dummies	0.265	(3.5)	0.155	(3.5)	0.373	(5.7)	0.223	(7.7)
F-test	3.79		4.80		1.42		4.15	
IVE: mpc without village effects, with village effects, and F-test for effects:-								
no dummies	0.192	(3.9)	-.003	(0.1)	0.271	(4.0)	0.126	(4.0)
dummies	0.171	(3.5)	0.029	(0.6)	0.270	(3.8)	0.107	(3.4)
F-test	3.94		4.19		1.38		4.37	
1986-87								
OLS: mpc without village effects, with village effects, and F-test for effects:-								
no dummies	0.458	(8.8)	0.162	(5.3)	0.168	(4.0)	0.239	(10.)
dummies	0.424	(8.1)	0.173	(5.6)	0.164	(3.8)	0.235	(10.)
F-test	2.90		2.61		0.56		2.12	
IVE: mpc without village effects, with village effects, and F-test for effects:-								
no dummies	0.418	(7.8)	0.090	(2.8)	0.088	(2.0)	0.177	(7.4)
dummies	0.388	(7.3)	0.105	(3.2)	0.087	(1.9)	0.177	(7.3)
F-test	3.10		2.43		0.61		2.14	

and the figure is included in both income and consumption. In consequence, any errors in this imputation, and finding the correct price is an obvious source of difficulty, will add a *common* error component to whatever other errors of measurement are present in the income and consumption totals. Positively correlated measurement errors do not necessarily generate an upward bias in the regressions in Table 4. For positive coefficients, the upward bias works in the opposite direction to the standard downward attenuation bias produced by the income errors alone, and the net effect cannot be signed in advance. To offset these problems, I have calculated instrumental variables estimators using as an instrument income changes excluding changes in the imputed value of home-produced food and other income in kind. Note that I am not challenging the appropriateness of including these imputations in both income and consumption. Rather the instrumental variable technique is an attempt to correct for the consequences of measurement error given those imputation procedures.

The instrumental variable estimates show that these concerns have some foundation, especially for the East Forest and the Savannah, where the estimated marginal propensities to consume are a good deal lower. Indeed for the East Forest in 1985-86 and for the Savannah in 1986-87, there is no longer any significant relationship between consumption and income changes. However, the instrumentation makes no difference to the relationship between the estimates with and without the village effects. As was the case for OLS, the introduction of the village dummies does not shift the estimates, and for the West Forest in both panels, the East Forest in the second panel, and the Savannah in the first, there remains a strong positive relationship between consumption change and income change that is robust to the inclusion of the village dummies.

These results seem to provide fairly firm evidence against the most extreme hypothesis, that there is *complete* consumption insurance within each of these Ivorian villages. In terms of the analogy with the permanent income hypothesis, there is clear evidence of excess sensitivity of individual consumption to individual income once group income has been controlled for. But the possibility remains that there is *some* insurance between households in the same village. Certainly, there is greater intra-village comovement in consumption than there is in incomes. However, there are many possible explanations for such a finding, including the all too likely effects of measurement error. The Ivorian data contain a great deal of within-village variance in income movement, variance that is large relative to whatever between village variance exists. Such an environment is not a good one with which to construct tests of partial consumption insurance.

LIST OF WORKS CITED:-

Ainsworth, M. and J. Muñoz, (1986), *The Côte d'Ivoire living standards survey*, Living Standards Measurement Study Working Paper No. 26, Washington, D.C. The World Bank.

Altonji, J.G., Hayashi, F., and L.J. Kotlikoff, (1989), "Is the extended family altruistically linked? Direct tests using micro data," Departments of Economics, Northwestern University, University of Pennsylvania, and Boston University, processed. (May)

Deaton, A., (1990), "Saving and income smoothing in the Côte d'Ivoire," Research Program in Development Studies, Princeton University, processed. (October)

Flavin, M., (1981), "The adjustment of consumption to changing expectations about future income," *Journal of Political Economy*, **89**, 974-1009.

Hayashi, F., (1987), "Tests for liquidity constraints: a survey and some new observations," in Bewley, T., editor, *Advances in Econometrics II: Fifth World Congress*, New York. Cambridge University Press.

Schwartz, G., (1978), "Estimating the dimension of a model," *Annals of Statistics*, **6**, 461-4.

Townsend, R., (1989), "Risk and insurance in village India," Department of Economics, University of Chicago, processed. (February)

Townsend, R., (1990), "Addendum on 'Risk and insurance in village India,'" Department of Economics, University of Chicago, processed. (May)