United Nations University

WIDER

World Institute for Development Economics Research

Discussion Paper No. 2002/58

Measuring Vulnerability to Poverty

Gisele Kamanou[1]
and Jonathan Morduch[2]

June 2002

Abstract

Many argue that poverty is intimately linked with 'vulnerability'. Still, there is no consensus about how to define and measure 'vulnerability'. We review theory and describe strengths and limits of recently proposed measures. We then propose a definition of vulnerability and develop a general empirical framework that combines Monte Carlo and bootstrap statistical techniques. The approach estimates the expected distribution of future expenditures for each household and then calculates vulnerability measures as a function of those distributions. The approach addresses weaknesses in existing methods, and can be implemented with panel data. An application to Côte d'Ivoire in 1985–86 shows that by our definition there was considerable vulnerability in the cities outside of Abidjan, a finding obscured by existing methods.

Keywords: vulnerability, poverty, poverty measurement, Côte d'Ivoire

JEL classification: O1, D3, I3

[1]United Nations, Kamanou@un.org, [2]New York University, Jonathan.Morduch@nyu.edu

This study has been prepared within the UNU/WIDER project on Insurance Against Poverty, which is directed by Dr Stefan Dercon.

UNU/WIDER gratefully acknowledges the financial contribution to the project by the Ministry for Foreign Affairs of Finland.

Acknowledgements

We are grateful to the Direction de la Statistique in Côte d'Ivoire and to the Living Standards Measurement Surveys programme of the World Bank for access to the data used here. The data are available for download at http://www.worldbank.org/lsms. Kristin Mammen kindly shared Stata code that greatly facilitated assembly of the data. Stefan Dercon and James Foster provided helpful comments on the first draft. All views and errors are our own.

UNU World Institute for Development Economics Research (UNU/WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by Jaana Kallioinen at UNU/WIDER
Printed at UNU/WIDER, Helsinki

ISSN 1609-5774
ISBN 92-9190-240-3 (printed publication)
ISBN 92-9190-241-1 (internet publication)

1 Introduction

The notion of 'vulnerability to poverty' remains elusive. It is a condition that may be easy to recognize in oneself or one's neighbours, but there is no consensus about how to define the concept and measure it in a broad cross-section of people. In surveys, poor households often identify vulnerability as a condition that takes into account both exposure to serious risks and defenselessness against deprivation. Defenselessness in turn is often seen as a function of social marginalization that ultimately results in economic marginalization (see, e.g., Kanbur and Squire 2001).

While this is a start, for quantitative researchers, the problem is to isolate a simple measure (or set of measures) of vulnerability that is comparable across time and location, – i.e. something akin to poverty measures. Not only have the principles underlying such a measure been hard to pin down, but practical applications are demanding of data and turn up challenges posed by a variety of forms of measurement error. Given the complexities, we focus here only on vulnerability to low consumption in order to shed light as sharply as possible. To focus even more sharply, we restrict attention to measures that can be constructed using quantitative data drawn from large, representative household surveys. Our concern is with statistical properties, and in that pursuit we necessarily sacrifice the richness of narratives emerging from qualitative approaches like participatory poverty assessments (e.g. Narayan et al. 2000). This is thus a first step toward an approach that we hope would ultimately combine some of that richness as well.

We begin by discussing the economic context in Côte d'Ivoire in the late 1980s, to provide context for the illustration in Section 5. Section 3 then discusses data issues including measurement error, changing household compositions, and attrition bias. Section 4 offers critical perspectives on existing approaches to measuring vulnerability, and Section 5 suggests a new framework that combines Monte Carlo and bootstrap statistical techniques. An application to Côte d'Ivoire in 1985–86 shows that by our definition there was considerable vulnerability in the cities outside of Abidjan, a finding obscured by existing methods. The conclusion highlights possible extensions of the approach.

2 Côte d'Ivoire 1985–88

2.1 Economic context[1]

The 1980's were a period of sharp economic challenge. Faced with the pressure of financial crisis in the early 1980s, the government of Côte d'Ivoire turned to the World Bank and IMF for financial assistance. The first round of financing, coupled with a structural adjustment programme, covered 1981 to 1986 with the objectives of stabilization through fiscal and monetary restraint. The programme entailed raising taxes on petroleum products, alcoholic beverages and tobacco. Despite the programme

[1] This section is based on Husain and Faruqee (1994).

(and in part as a result of it), investment expenditures fell from 23 per cent of GDP in 1979 to just 5 per cent in 1986. Current expenditures were also cut, although to a lesser extent, and civil service wages were frozen between 1984 and 1986. Agricultural reforms under structural adjustment mainly aimed at raising export production. The consumer price of rice was raised, leading to price increases of all other food crops.

Improved terms of trade brought a slight economic recovery in 1984–86, but in 1987, the international prices of cocoa and coffee fell sharply. The country then plunged into another severe economic crisis. Agricultural policies in 1988 were design to retain part of cocoa crops. This policy disrupted the agricultural marketing cycle: exporters were not reimbursed by the Agricultural Price Stabilization Fund and were unable to obtain export permits. Consequently, they were unable to repay crop credit, which in turn aggravated the liquidity problems of the banking system. At the same time, the country withdrew from adjustment programmes and all bank lending to Côte d'Ivoire was interrupted between 1987 and 1989. The burden of adjustment fell on public investment expenditures, which declined from 18 per cent of GDP in 1978–83 to just 3 per cent 1988–91.

2.2 Data

We consider vulnerability using the 1985–88 rounds of the Côte d’Ivoire Living Standards Survey. Using the CILSS for illustration is particularly helpful as the data have been used by others to study poverty dynamics, allowing us to draw on related studies of consumption patterns, poverty, and household behavior (see, e.g., the studies cited in Deaton 1997).

The sample selection process for the CILSS household surveys yielded a nationally-representative cross-section of Ivorian households. Each year 1600 households were sampled, half of which were revisited the following year; the other half was replaced with new households. The net result is four cross-sectional datasets and three two-year panels (for 1985–86, 1986–87, and 1987–88). Consumption and income have been deflated to capture regional price variation and overall inflation (anchored by 1985 Abidjan prices).

The poverty line is set here at 128,600 CFA, following Deaton (1997, p.155) and Grootaert and Kanbur (1995). The line corresponds roughly to the dollar-a-day poverty line used often in international comparisons. Using this poverty line, we find that the headcount poverty rate rises from 31 per cent in 1985 to 48 per cent in 1988.[2]

The economic crisis in Côte d’Ivoire of the 1980s is echoed in the household survey data. Per capita consumption fell sharply between 1985 and 1988 and households faced considerable uncertainty. Table 1 shows that real per capita consumption fell from 223,226 CFA in 1985 to 175,327 in 1988 in our sample.

---

2 Here, as in much below, the sample is restricted to households with proportional consumption changes between 2 and –2 (i.e. those who report changes by less than 200 per cent up or down). The restriction removes outliers and restricts attention just to households in complete two-year panels. Results without the restriction are similar. In the unrestricted sample, poverty increases from 31 per cent to 44 per cent and per capita consumption drops from 231,971 CFA to 178,641 in 1985-88.

Table 1
Poverty measures and per capita consumption (CFA)
All households

	1985	1986	1987	1988
Headcount	0.31	0.34	0.35	0.48
Normalized poverty gap	0.094	0.094	0.099	0.157
Squared poverty gap	0.041	0.036	0.041	0.073
Expenditure per capita	223,226	206,872	216,179	175,327
Household size	11.8	10.8	9.8	9.7

Note: Sample is restricted to households with consumption changes between 2 and –2.

3 Data issues

Our main focus is on conceptual frameworks, but before turning there, we begin by reviewing three important data problems that can be particularly acute when studying consumption and income dynamics: attrition bias, changing household consumption, and measurement error.

3.1 Attrition

Even if data quality is high, biases emerge when panel data are incomplete. The most extreme form occurs when households leave the panel. Table 2 gives rates of attrition; they range between 11 and 15 per cent of the non-poor sample and between 5 and 13 per cent of the poor sample. The numbers are in the range of those for surveys of this kind. For the purposes here, attrition is particularly problematic when those who leave the survey are differently vulnerable than others. Those who leave may be among the most vulnerable, or, since migration can be a coping mechanism, they may be less vulnerable. Regression analyses show that rates of attrition are negatively associated with increased age of the household head and larger household size. On average, smaller, younger (and presumably more mobile) households are more likely to leave the sample, but without follow-up surveys it is impossible to discern the exact nature of biases.

Table 2
Attrition (% of base year sample)

	1985–86	1986–87	1987–88
Poor	8	13	5
Non-poor	11	14	15

3.2 Changing household composition

Table 1 shows that in the CILSS household size fell from 11.8 in 1985 to 9.7 in 1988. Partly this is due to a change in sampling frame procedures between the first and second panel in 1987 (Coulombe and Demery 1993). Table 3 shows this shift clearly. But even without a change in the sampling frame, households often change sharply in size and composition. For the poor population, for example, the household size fell on average by 7 per cent between 1985 and 1986. The changes are generally less sharp for the non-poor.

The changes are partly due to family splits and to migration, some of which was induced by the economic crisis. Counter-balancing the losses, many households also report births and the arrival of relatives and others. We show below that analyzing the data in terms of adult equivalence rather than per capita values goes some distance in addressing the problem, but the net effect on results is modest here.

The bigger issue concerns interpretation. Throughout the paper we take changes in per capita income and consumption (or adult equivalents) to signal 'shocks' (like price changes or low rainfall). But variance decompositions suggest that as much as a quarter of the variation in per capita consumption is due to the denominator – to changes in household size. These are typically not 'shocks' as we commonly think of them. Instead, they are often the product of deliberate choices made by households. Decreasing per capita consumption associated with the birth of children may reasonably be deemed desirable in many cases – not a 'negative shock' at all (Anand and Morduch 1999). In principle, the bootstrap Monte Carlo approach discussed below could be extended to allow prediction of shocks after conditioning on changing demographic structure, and this is left here as a caveat on interpreting our results.

3.3 Measurement error

Measurement error poses a serious challenge for analysts of vulnerability (Baulch and Hoddinott 2000). The error can come in several forms. First, errors in forming measures of consumption aggregates; second, inappropriate price deflations; third, inappropriate deflations for household size; fourth, errors in matching households in different waves of panel data.

Table 3

Mean changes in household size (%)

Panel households only

	Members			Adult equivalents		
	1985–86	1986–87	1987–88	1985–86	1986–87	1987–88
Poor	-7	-20	-4	-6	-17	-3
Non-poor	-3	-6	-3	-3	-6	-2

Table 4

Average proportional changes in income and consumption

	Income			Consumption		
	1985–86	1986–87	1987–88	1985–86	1986–87	1987–88
Poor						
Quartile 1	0.287	0.286	0.309	0.516	0.370	0.109
Quartile 2	0.439	0.136	0.301	0.332	0.159	0.041
Quartile 3	0.188	0.151	0.185	0.235	0.146	-0.067
Quartile 4	0.167	0.203	0.066	0.088	0.118	-0.113
Non-poor						
Quartile 1	0.143	-0.022	0.141	0.030	-0.022	-0.067
Quartile 2	0.027	-0.059	0.027	-0.156	-0.048	-0.133
Quartile 3	0.051	0.026	-0.010	-0.189	-0.076	-0.176
Quartile 4	-0.054	-0.076	-0.069	-0.315	-0.091	-0.250

Note: Sample is restricted to households with consumption and income changes between 2 and –2.

Grootaert et al. (1997, pp. 645–8) argue that these problems are apt to be minimal in the CILSS. Their assertion rests on several claims. First, the average consumption declines in the household data match quite closely with the declines in private consumption seen in the (independently derived) national accounts. Second, if there is misreporting in transitory income, it is apt to be under-reported, leading to an under-estimation of vulnerability not exaggeration. Third, first-differencing the data will eliminate all measurement error that is fixed from period to period. Fourth, changes in consumption appear to be systematically related to human capital variables and household composition, suggesting that consumption variation is not due mainly to error. And, fifth, the quality control of the CILSS fieldwork was excellent, with many built-in cross-checks on data quality (Ainsworth and Munoz 1986).

Unfortunately, we cannot assume that other data sets are of similar quality. And, even here, there are reasons for concern. Table 4, for example, shows mean reversion of a sort that is consistent with substantial measurement error. The table gives average proportional changes in per capita income and consumption in the three 2-year panels. (Table 5 shows that accounting for changing household composition through the use of adult equivalents and modest returns to scale does not change the picture very much – although in other contexts it might.) The data are disaggregated by poverty status and, within that, by quartile. The top half of Table 4 shows that average income improved for poor households in each panel, while consumption improved for all but the third and fourth quartiles of the poor in 1987–88. The early improvements are often large: on the order of 30 per cent per year for the poorest groups. But the improvements do not hold for the better-off households, where average consumption almost uniformly declined (by as much as 25 per cent in 1987–88 for the highest quartile). Figure 1 shows the results in a simple graph. For all years, average changes turn from positive to negative close to the poverty line (the vertical line at 128,600 CFA).

Table 5

Standard deviations

Proportional changes in income and consumption per capita

	Income			Consumption		
	1985–86	1986–87	1987–88	1985–86	1986–87	1987–88
Poor						
Quartile 1	0.609	0.559	0.580	0.629	0.363	0.523
Quartile 2	0.688	0.576	0.712	0.451	0.429	0.490
Quartile 3	0.678	0.593	0.589	0.612	0.386	0.406
Quartile 4	0.506	0.608	0.589	0.439	0.439	0.339
Non-poor						
Quartile 1	0.627	0.604	0.605	0.449	0.445	0.322
Quartile 2	0.652	0.546	0.566	0.339	0.438	0.397
Quartile 3	0.630	0.604	0.546	0.348	0.396	0.353
Quartile 4	0.590	0.572	0.695	0.259	0.384	0.358

Note: Sample is restricted to households with consumption and income changes between 2 and –2.

Figure 1

Top: Density of per capita expenditure (left) and per capita income

Bottom: Proportional changes in per capita expenditure and income versus base year levels

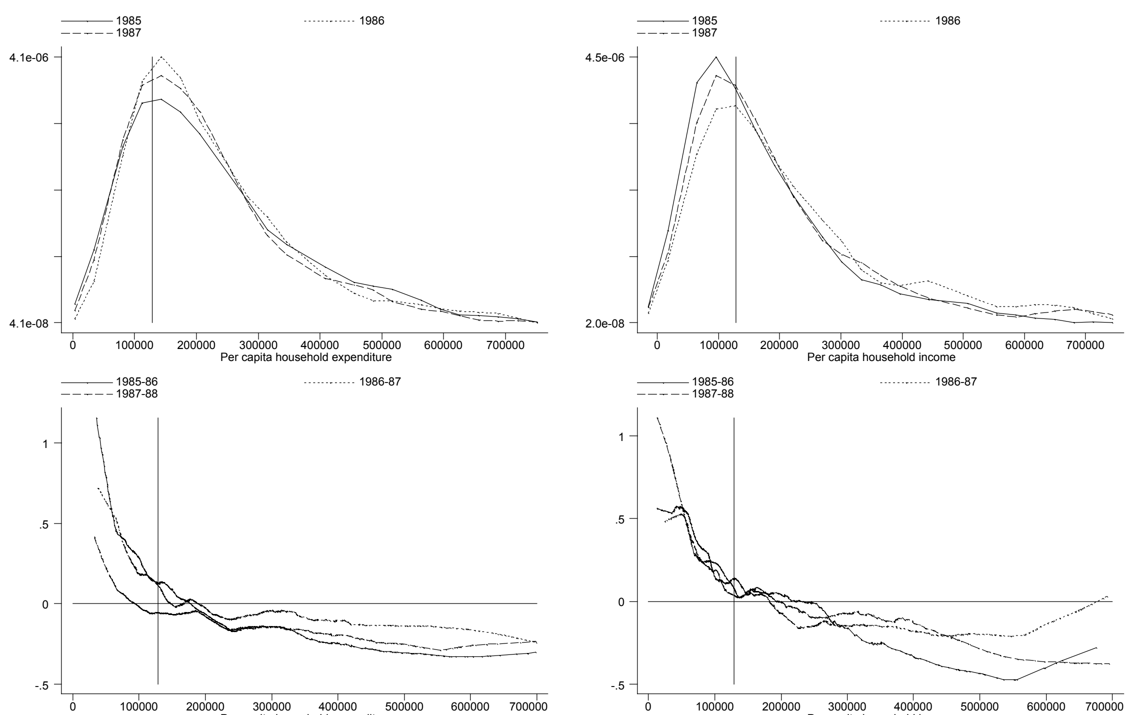

Note that the richest households have disproportionately large negative consumption changes, and the poorest have disproportionately large increases. This is consistent with measurement error in the base period, where error makes the 'rich' too rich and the 'poor' too poor. In the subsequent period, they tend to be closer to where they should have been. Some of the changes, of course, may be real, but Table 4 shows the changes to be too large to be fully credible as true variation. One of the ways that the problem could be approached is to use an alternative wealth measure to generate base-year categorizations; one possibility is the housing index of Kamanou (1999) which uses principal components methods to create alternative measures of well-being using CILSS data.

4 Conceptual frameworks

With those empirical issues in mind, we turn to conceptual frameworks.

4.1 Expected utility theory

To capture the idea of vulnerability, we could start with the microeconomic theory of risk and uncertainty. The theory of expected utility tells us that the expected utility of risk averse individuals falls as the variability of consumption rises, holding all else the same. If we knew the utility functions and expected consumption patterns of all individuals, we could then analyze poverty in terms of certainty-equivalent consumption (the level of consumption which, if unvarying, would yield an equivalent level of expected utility as a household's actual – higher mean but more variable – consumption levels). The theory allows the integration of consumption variability in a natural way, but the data requirements are too high to be of much practical use here. Not only are utility functions unseeable, but there are just a handful of longitudinal data sets from low-income countries with an adequate time dimension to yield precise measures of household-specific consumption variability.[3]

4.2 Mobility measurement

An alternative, and more promising, starting point is the literature on income mobility (e.g. Shorrocks 1978, Fields and Ok 1999). The literature avoids cardinally comparable utility functions and evaluates the evolution of income rankings, focusing on the degree of path dependence. The starkest beginning point is the question: if you are born poor, are you more likely than others to die poor? Concern has tended to focus on long-term transitions across ranks (e.g. from generation to generation), rather than the short- and medium-term fluctuations considered here, but it provides a jumping-off point.

---

[3] In addition, the tradition of poverty measurement has always had just a loose relationship with cardinal notions of utility, and introducing those notions via certainty-equivalent consumption introduces restrictions on measures that will create disjunctures with standard ordinal-based frameworks.

The mobility literature as a whole, though, has limitations with respect to our narrow interest. Most important, the literature is concerned primarily with changes in income rankings: it is relative movement that matters there. Here, though, the concern is with movements judged against an absolute standard given by official poverty lines. Given an absolute standard, simpler measures may suffice, where they would inadequately capture mobility. For example, Grootaert and Kanbur (1995) profitably employ Markov transition matrices in order to study changes in poverty status in the CILSS data.

Moreover, the mobility literature focuses on historical patterns, not prospective ones. This is a critical distinction. With just two observations on a set of households, for example, we can see which households became poor, which exited poverty, and which did not cross the line. But some among those that did not change status may well have been 'vulnerable' to a misfortune – however, good luck intervened to prevent a downturn. It is this condition of *possibility* that we are trying to capture. Measuring historical patterns is of course important, but for policy purposes the hope is to be able to generate measures that allow targeting of groups that are 'vulnerable' to loss, not just households that can be identified as actually having suffered in retrospect (a point stressed also by Chaudhuri 2000).

4.3. Vulnerability as variability

If the history of shocks is a predictor of future shocks, a simple starting point for measuring vulnerability would be to compare standard deviations of consumption and income changes. Households or groups are judged to be more vulnerable if standard deviations of past consumption changes are higher. The approach is simple and draws on the association of vulnerability with variability.

An immediate practical problem here is that few data sets have a long enough time dimension to yield a reliable standard deviation for each household over time. Most, like the Côte d'Ivoire Living Standards Survey, have two repeated observations for a household at best. In practice, the standard deviation then must be estimated from cross-sectional variation. Thus it picks up the dispersion of shocks not their average strength. A strong homogeneity assumption must also be made in order to interpret results of vulnerability, namely that all households observed in the cross-section receive draws from the same distribution of consumption changes. If this is so, measures of dispersion of the changes will then indicate the degree of exposure. One can do better by disaggregating further – by region, by income group, by education status, etc. But, at the end of the day, some kind of homogeneity assumption must be made, leaving concern with unobserved heterogeneity.

Strengths and weaknesses are illustrated in Tables 5 and 6. The right half of Table 5 and 6 show that, almost uniformly, the standard deviations of proportional consumption changes are lower than those of income, indicating considerable consumption smoothing in the limited sense above. In Table 6, for the top quartile of the non-poor in 1987–88, the difference is nearly half (0.695 for income versus 0.358 for consumption), but most differences are on the order of 20 per cent. Within consumption, variability is somewhat less for richer households than poorer, but the magnitudes are in the same range (roughly 0.35 to 0.45).

Table 6

Means and standard deviations

Proportional changes in consumption per adult equivalent

	Mean			Standard deviation		
	1985–86	1986–87	1987–88	1985–86	1986–87	1987–88
Poor						
Quartile 1	0.528	0.315	0.091	0.624	0.307	0.479
Quartile 2	0.287	0.128	0.016	0.489	0.396	0.439
Quartile 3	0.228	0.131	-0.075	0.565	0.369	0.375
Quartile 4	0.075	0.093	-0.140	0.407	0.446	0.264
Non-poor						
Quartile 1	0.015	-0.023	-0.077	0.415	0.424	0.322
Quartile 2	-0.161	-0.050	-0.129	0.332	0.405	0.372
Quartile 3	-0.138	-0.070	-0.157	0.523	0.378	0.335
Quartile 4	-0.289	-0.060	-0.254	0.239	0.403	0.347

Note: Sample is restricted to households with per capita consumption changes between 2 and –2.

Comparison of the tables with Table 4 illustrates one immediate concern with using the standard deviation as a measure of vulnerability: downside risk is weighed the same as upside risk.[4] A 10 per cent upward shock affects the standard deviation identically to a 10 per cent downward shock. The tables show that the standard deviations for richer households, while somewhat lower than for poorer households, are measured with respect to decreasing consumption. If standard deviations were all that we had to go by, it would be reasonable to conclude that variability is hurting the poor more than the non-poor. But vulnerability is not just about variability: most observers would argue that down-side risk is the chief concern and it may be that, with declining average consumption levels, down-side risk is greater for non-poor households. These statistics, while commonly employed, are not sufficient to discern the essence of vulnerability. Turning to coefficients of variation could help (standard deviation divided by the mean), but with some means at zero or very close to zero, the coefficients of variation blow up in some cells, obscuring comparisons.

A second concern is that the standard deviation can give odd results in times of persistent growth. Imagine that consumption is growing for a household so that over 8 periods their consumption is (1, 2, 3, 4, 5, 6, 7, 8). Take another household with a consumption pattern of (7,5, 2, 6, 3, 1, 4, 8). The standard deviations of both series are identical, but the second appears to be buffeted by shocks while the first is on a steady upward path. Labeling them both as identically vulnerable misses the key part of their stories.[5]

---

4 See World Development Report 2000/2001, Box 3, p.20, based on Dercon (2001).

5 A related tension emerges when defining transitory poverty as deviation from mean consumption in rapidly-growing regions like China (see, e.g., Jalan and Ravallion 1998).

A third, related concern is that the standard deviation offers no accounting for persistence of downturns – or negative serial correlation. Across eight periods, the pattern (1, 0, 1, 0, 1, 0, 1, 0) has the same standard deviation as (0, 0, 0, 0, 1, 1, 1, 1) but, if we interpret 1 as being poor and 0 being non-poor, the second pattern might suggest a more serious exposure to poverty. In terms of the standard deviation of the change in poverty status, the patterns above imply changes of the form (1, -1, 1, -1, 1, -1, 1) and (0, 0, 0, -1, 0, 0, 0). The former has a higher standard deviation, but again the latter flows from a situation which many would associate with greater vulnerability. With just two-years panels, this concern underlies our work below as well, since persistence can only be addressed when more data are available.

4.4 Risk of change in poverty status

Another starting point has been to identify the probability of becoming poor (or becoming poorer) in a given sample. For many, this implies that vulnerability is a condition of the *non-poor* only. The approach has appeal by highlighting the groups that could be helped by preventative measures before adverse events are realized.

Dercon and Krishnan (2000, pp.44–5) measure 'vulnerability' in rural Ethiopia, for example, by estimating determinants of consumption levels and then predicting the degree to which households would suffer severe consumption shortfalls given particularly poor rainfall (less than half the long-term mean). Their estimates suggest that the 'vulnerable' population (those that have a risk of falling below the poverty line) is 40 to 70 per cent higher than the observed poverty rate.

Pritchett et al. (2000) answer the question by estimating the standard deviation of consumption changes in the cross-section and then, given that variation, predicting the income level below which households are more than 50 per cent likely to be poor next period. The idea is clear and not particularly demanding of data (a two-year panel is sufficient). A limitation is that the problems with the standard deviation are unavoidable in this framework as well. In using the bootstrap Monte Carlo method in Section 5 below, it is possible to avoid relying on the standard deviation in this way (but, like Pritchett et al., we are forced to use cross-sectional variation to predict intertemporal variation).

A refinement of this basic approach is suggested in Morduch (1994) and a related idea is applied in Jalan and Ravallion's (1999) work on China. The idea is to disaggregate the population into groups that are 'structurally poor' (in the sense of having fundamentally low earning power) and those that are not. Then within each group, subgroups that become poor due to bad shocks or that exit poverty due to good shocks are identified separately. In this way, the 'non-poor' in a base year are disaggregated according to those who appear to be 'fundamentally' not poor and those who, in a given period, appear to be above the poverty line through good fortune only. As a practical matter, though, more than two years of data are required.

4.5 Ability to cope

Better data and new empirical methods have led to new studies on coping mechanisms and the efficacy of 'informal insurance' (see, e.g., Dercon 2001 and Morduch 2000).

Amin et al. (1999) use econometric methods for measuring the efficiency of informal insurance to form a measure of vulnerability. In their work, 'vulnerability' is associated with consumption fluctuations associated with imperfect risk sharing as set out by Townsend (1994). A similar approach is taken by Jalan and Ravallion (1999); see Morduch (2002) for a discussion of the approach and empirical evidence.

In the Amin et al. (1999) study of two villages in rural Bangladesh, a household is considered vulnerable in proportion to the extent to which income shocks translate into consumption shocks. Specifically, they estimate the equation:

$$\Delta \ln c_t^h = \alpha^h \Delta \ln y_t^h + \sum_t \beta_t X_t + \varepsilon_t^h, \qquad (1)$$

where the change in log consumption at time t for household h is regressed on a set of time-specific dummy variables for the village (X_t) and the change in log income for the household. Full insurance is assumed if the household-specific parameter α is 0. Complete autarky is implied by $\alpha = 1$, and values of α approaching 1 are taken to imply heightened vulnerability. For a quarter of the households α is negative, which is perplexing but nevertheless interpreted as implying a lack of constraints and thus full insurance. Most households are determined to be somewhat vulnerable, with female-headed households taking a coefficient on α that is 0.35 higher than that of male-headed households.

The approach captures an important, but selective, element of vulnerability. In general the vulnerability of a population is the product of three elements:

i) The pattern of possible 'shocks'. These may be losses due to, say, losing a job or experiencing a bad harvest. Included here are also increases in needs due, for example, to illness, child birth, or costly social occasions.

ii) The strength of coping mechanisms. This is the degree to which provisions are not in place to fully address shocks (and is captured by the Amin et al. 1999, method).

iii) Structural and behavioral ramifications of consumption declines. Are they apt to lead to temporary shortfalls or to lead to poverty traps?

The estimate of α only captures the second element, and the decomposition suggests that the focus on α is apt to be incomplete for many policy purposes. First, the estimates only get at idiosyncratic risk (since α captures the role of income fluctuations over and above those picked up by the time-sepecific dummy variable for the village). In Côte d'Ivoire, though, covariant risk is a critical element of vulnerability. Second, the measure of vulnerability here captures the transmission of both positive and negative income shocks. In principle, a household could fend off all downward income shocks adequately well (so they do not translate into consumption declines) but then consume out of positive income shocks (making hay while the sun shines). This would lead to an estimated α between 0 and 1, but these households would not be 'vulnerable' in the sense assumed here. It is unlikely that this is the case in the Bangladesh sample, but it is a concern as we look for general approaches.[6]

In the same light, there are concerns that surround policy implications. Consider two populations. Both could have the same average standard deviation of consumption changes. In one, though, there may be quite good coping mechanisms (a low estimated α) but many shocks. In the other, coping mechanisms may be far weaker (high α), but the shocks may be much weaker too. In the first case, policy might most effectively attempt to reduce income variability. In the latter, it might best strengthen coping mechanisms. Here, the Amin et al. (1999) yields very useful information. But consider instead a situation where both populations have the same estimated α such that both seem equally 'defenseless'. In the first, again, assume that there are many income shocks and in the second there are far fewer. Many would say that 'vulnerability' is greater in the first case, but the approach would view both cases identically.

One way to address this concern is to form an index from the product of α and a measure of income shocks. This, though, effectively leads us back to working with consumption fluctuations directly. All the same, estimating α provides a means to identify 'defenselessness' from 'vulnerability', and with a sufficiently long panel might profitably be used in conjunction with the ideas set out in Section 5 below. (See also Ligon (2002) which explores approaches to targeting vulnerable households in the relatively long ICRISAT panel.)

4.6 Asset holding

Another approach is to measure assets rather than consumption patterns (e.g. World Development Report 2000/2001*,* Box 3, p.20). The idea addresses frustration at identifying vulnerability measures (and at often lacking panel data). The idea flows from what we know about coping mechanisms: having more assets generally makes coping easier and households work hard to hold onto particular assets. An inventory of assets can show how much of a cushion households will have in time of crisis – and how costly addressing the crisis by selling assets is likely to be.

The approach has promise, and the first order of business is to better identify the exact relationship of assets and vulnerability. This, of course, requires having a clear measure of vulnerability to start with. As with the approach discussed above, where vulnerability is associated with the ability to smooth idiosyncratic shocks, an asset-based measure yields useful information about coping mechanisms conditional on shocks, but extra information is required to inform about the distribution of expected shocks.

5 A new framework

We set out a simple framework that combines elements of the approaches above and addresses some shortcomings. The framework developed here combines Monte Carlo simulations with the bootstrap, a nonparametric method for estimating the standard error

6 One way to address the tension is to distinguish between negative and positive shocks in the regression equation, following, e.g., Jacoby and Skoufias (1997).

of sample parameters (Efron and Tibshirani 1994). The idea is to generate a distribution of *possible* future outcomes for households, based on their observed characteristics and the observed consumption fluctuations of similar households. We are limited by having just two repeated observations on each household, but the approach could be refined if more data were available – in particular, a longer time series would allow consideration of persistent poverty versus shorter spells (see, e.g., Morduch 1991).

5.1 Preliminaries

To be concrete, we employ the commonly-used Foster-Greer-Thorbecke (1984) class of poverty measures,

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{G}\left(\frac{z-y_i}{z}\right)^\alpha \qquad (2)$$

where with $\alpha = 0$ we get the headcount index, with $\alpha = 1$ we get the (normalized) income gap, and with $\alpha = 2$ we get a measure that is sensitive to the distribution of consumption below the poverty line.[7] Here, z is the poverty line, y is household income, N is the total population size, and households are ordered from bottom to top: $y_1, y_2, \ldots y_G, z, y_{G+1}, \ldots y_N.$ Using the measure, we can gauge changes in poverty status across periods (generalizing earlier analysis of the CILSS by Grootaert and Kanbur 1995).

More important, we can use the framework to derive a measure of vulnerability that applies to all households, both currently poor and not. Before turning to the Monte Carlo bootstrap approach, we consider historical experience in the base data. To do this, we focus on a single measure, which with only two years of data is ($P_{\alpha t+1}$ - $P_{\alpha t}$).[8] If attention was just accorded to transitions into poverty, attention could be restricted just to households that begin in the basis year above the poverty line (so that the measure is just $P_{\alpha t+1}$). In the present application, we consider only households with two complete years of data, with base year summary statistics provided in Table 7. Following the earlier tables, Table 8 provides measures for quartiles of the poor and quartiles of the non-poor (defined by status in base years), and Figure 2 shows the relationship graphically. The first column of the table shows that the 45 per cent of the first quartile of those who were not poor in 1985 were poor by 1986. For the next richest quartile, 18 per cent became poor, and for the third quartile, 10 per cent became poor. Just 1 per cent of the richest group was poor in 1986.

If we take this evidence alone, there appears to be a great deal of 'vulnerability' (as measured by worsening poverty status). Yet poverty overall increased by just 8 per cent. This is because (as seen in Figure 1), many who were poor in 1985 exited poverty (from the poorest to least poor quartiles below the poverty line in 1985, the rate of exit was 15 per cent, 29 per cent, 25 per cent, and 41 per cent). The magnitude of changes should not be surprising once the frequency distributions of per capita consumption are seen in Figure 1– much of the population is very close to the poverty line so a small tip one way or the other can lead to a change in measured poverty.

7 Chaudhuri (2000) also suggests using this index in the context of vulnerability measurement.

8 Estimates of poverty changes are examined, for example, in Dercon and Krishnan's (2000) work on seasonality and poverty in Ethiopia.

Table 7
Per capita consumption (1985 Abidjan CFA)
First wave of households in 2-year panels

	1985	1986	1987
Poor			
Quartile 1	55,936	62,257	56,419
Quartile 2	82,552	90,400	85,615
Quartile 3	103,375	103,499	104,897
Quartile 4	120,724	119,682	120,152
Non-poor			
Quartile 1	153,921	149,990	153,154
Quartile 2	221,706	204,635	209,825
Quartile 3	326,669	295,264	293,815
Quartile 4	634,982	592,414	616,885
Total	222,727	225,583	210,394
Headcount Index	0.302	0.300	0.345

Table 8
Observed changes in measured poverty

	Change in headcount index			Change in squared poverty gap		
	1985–86	1986–87	1987–88	1985–86	1986–87	1987–88
Poor						
Quartile 1	-0.145	-0.107	-0.042	-0.150	-0.132	-0.011
Quartile 2	-0.287	-0.287	-0.161	-0.037	-0.004	0.060
Quartile 3	-0.246	-0.392	-0.118	-0.043	0.014	0.088
Quartile 4	-0.406	-0.522	-0.183	-.043	0.034	0.067
Non-poor						
Quartile 1	0.454	0.396	0.474	0.029	0.037	0.026
Quartile 2	0.177	0.194	0.249	0.012	0.016	0.012
Quartile 3	0.096	0.069	0.087	0.002	0.003	0.006
Quartile 4	0.013	0.010	0.007	0.002	0.000	0.000
Total	0.076	0.034	0.119	0.000	0.005	0.025

Note: Poverty rate in final year less the rate in the base year. Positive numbers reflect a worsening of conditions.

Figure 2

Probability of being poor in the subsequent period versus base year levels

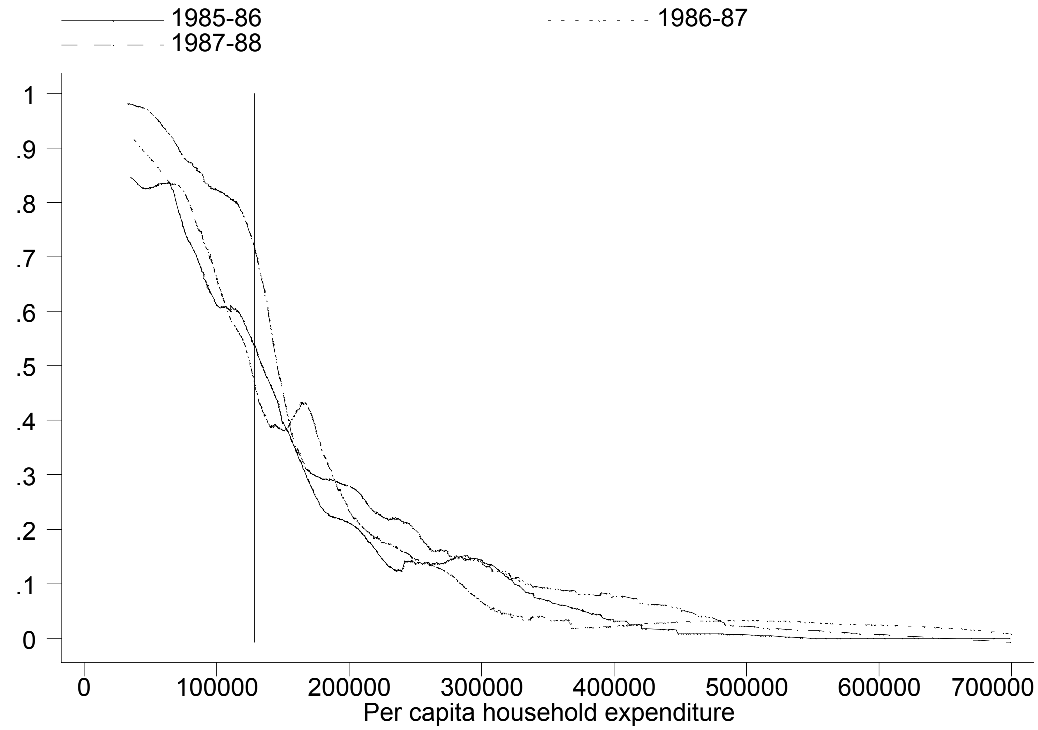

If these processes are independent, policymakers would want to encourage exit and discourage entry. But imagine that the processes stem from the same origin, that they are intrinsically linked. In that case, reducing the vulnerability of the non-poor may also reduce the chance for exiting poverty of the poor, an issue on which we can only speculate.

Considering the squared poverty gap in Table 8 offers a different story. Here, the top two quartiles of the poor also do worse from period to period. While some households are exiting poverty, others are becoming much worse off, and the downside outweighs the upside. The bottom lines are that (1) a large fraction of non-poor households face the risk of becoming poor in the next year. (2) By the same token, many poor households will exit poverty. A key question (which is unanswerable in these data) is whether those that are poor are also likely to soon exit. If so, notions of 'vulnerability' take a very different cast than if spells were persistent.

5.2 Monte Carlo bootstrap approach

In a natural extension of the framework above, we define vulnerability in a population as the difference between the *expected* value of a poverty measure in the future and its current value. This yields a set of measures of vulnerability that expand the framework illustrated in Table 8. The measures take the form:

$$E\,P_{\alpha t+1} - P_{\alpha t} \quad = \quad \frac{1}{N}\sum_{i=1}^{G_{t+1}} \sum_{S} \Pr(s, y_{it+1})\left(\frac{z - y_{it+1}}{z}\right)^{\alpha} - \frac{1}{N}\sum_{i=1}^{G_t}\left(\frac{z - y_{it}}{z}\right)^{\alpha} \qquad (3)$$

where E is the expectation operator, s is a given state of the world for which the joint probability distribution with Y_{t+1} is Pr(s, y), G_t and G_{t+1} are the number of poor households in the current and future periods respectively, and y_{it} and y_{it+1} denote the current and future consumption of household i. We assume that the true distribution of *possible* outcomes in the next period for households (y_{it+1}) could be known.

The practical implementation of this approach, however, is made difficult by the fact that the joint distribution of s and y_{it+1} is not known. In fact the 'states of the world' might be latent variables with an unknown distribution. The idea is to make up for the unknown joint distribution Pr(s,y) by generating a distribution of possible future outcomes for households, based on their observed characteristics and the observed consumption fluctuations of 'similar' households. In other words, the bootstrap technique enables us to construct several versions of possible future data by resampling the original data. The expected value is thus estimated by the mean of the bootstrap estimate of $P_{\alpha t+1}$.

The approach is implemented by starting with the base year of a panel and drawing a large number (say B=1000) of independent bootstrap samples. (A bootstrap sample is a random sample of size n drawn with replacement from the empirical distribution of some observed data of size n.) With each bootstrap sample, a regression equation is constructed to predict the change in consumption based on its correlation with a set of household covariates. The linear part of the predicted value is then augmented with a shock drawn at random from the empirical distribution of the regression residuals.

This yields a predicted per capita expenditure of the future period for each household in each of the bootstrap samples. Using these samples we compute $P_{\alpha t+1}^{b}$ for each b from 1 to 1000, and we then estimate $EP_{\alpha t+1}$ as the mean of the bootstrap estimates $P_{\alpha t+1}^{b}$.

The algorithm goes as follow for each region:

Step1: Draw 1000 bootstrap samples from the original data. Let $X = (x_1, x_2, x_3, ..., x_n)$ be the data on a given region (n is the number of households in the region) with $x_i = (y_{i1}, y_{i2}, h_1, h_2, ... h_p)^T$ and y_{i1} and y_{i2} are the first and second wave household consumption and $h_1, h_2, h_3, ..., h_p$ are p household covariates. Say $X^b = (x_1^b, x_2^b, x_3^b, ..., x_n^b)$, $b = 1,2,3,...,1000$ are the bootstrap samples drawn by resampling $X = (x_1, x_2, x_3, ..., x_n)$ with replacement.

Step2: With each new bootstrap sample, run a regression of $\delta_i = (y_{2i}^b - y_{1i}^b)/y_{1i}^b$ on covariates $h_1, h_2, h_3, ..., h_p$ and form the Monte Carlo estimates of the future period consumption by: $\hat{y}_i^{mcb} = y_{1i}^b(1 + \hat{\delta}_i + \varepsilon^{mc})$ where $\hat{\delta}_i$ is the fitted values from the regression for the household i, and ε^{mc} is a random draw for the empirical distribution of the residuals from the regression.[9] For the illustrative example shown below we constructed the prediction equation for second period consumption using the Generalized Linear Model to fit the proportional change in the per capita consumption (δ_i) on household covariates including first wave per capita consumption, household size, age of the head of the household, nationality, and indicators of whether the head of the household is literate and numerate. The independent variables also included an index of housing constructed using a combination of principal components analysis and the U-scores, a new statistical technique used for scaling discrete variables in the context of Principal Components (see Kamanou 1999).

***Step3*:** Form an estimate of the future period's poverty measure of the bootstrap sample as: $\hat{P}^{mcb}_{\alpha t_2} = \frac{1}{n}\sum_{i=1}^{G_b}\left(\frac{z-\hat{y}_i^{mcb}}{z}\right)^{\alpha}$. This is the Monte Carlo estimate of the future period's poverty measure obtained from the bootstrap sample.

***Step4*:** The Monte Carlo bootstrap estimate of vulnerability for the population for the period (t_1, t_2) is then defined by: $V^{mcb}_{\alpha} = \hat{P}^{mcb}_{\alpha t_2} - P_{\alpha t_1}$.

These data are used to calculate the change in the observed headcount in Column 4 and the vulnerability index in the final column. The most striking difference between the two occurs for the cities outside of Abidjan. There, the observed change in the headcount is small (-0.02) while the vulnerability index shows a larger, positive value (0.10). The index shows that many more households were 'vulnerable' to poverty than actually became poor, a condition that played out in later years as the economy worsened, pushing richer households below the poverty line.

The illustration shows the power of the basic approach. As described above, the application can be extended to incorporate distributionally-weighted poverty indices as well, yielding additional dimensions by which to view the changes in Côte d'Ivoire.

Table 9

Monte Carlo Bootstrap results

Vulnerability index (based on poverty headcount)

1985–86 panel (713 observations)

	1985 Observed headcount	1986 Observed headcount	1986 Bootstrap headcount	Change in observed headcount	Vulnerability index
Abidjan	0.05	0.16	0.17	0.11	0.12
Other Cities	0.25	0.23	0.35	-0.02	0.10
West Forest	0.11	0.22	0.28	0.11	0.17
East Forest	0.45	0.37	0.44	-0.07	-0.01
Savannah	0.46	0.48	0.43	0.02	-0.03

9 The method can be refined by allowing residuals to have a household-specific component (see Chaudhuri 2000).

To illustrate, we compared the five regions of the Cote d'Ivoire according to their vulnerability measures in the 1985–86 sample. Table 9 gives the observed poverty headcounts for the 1985 and 1986 waves for each region in the 1985–86 panel. Table 9 also gives the Monte Carlo bootstrap estimate of predicted poverty in 1986, $\hat{P}_{\alpha_2}^{mcb}$.[10]

6 Concluding comments

We have proposed a new definition of vulnerability, and we have developed an approach built around Monte Carlo bootstrap predictions of consumption changes. The approach both generalizes and extends previous approaches, and it avoids unappealing features of existing measures.

We have illustrated the methods using the two-year panels of the Côte d'Ivoire Living Standards Survey. Our method reveals substantial measured 'vulnerability' between 1985 and 1986 in cities outside of Abidjan, for example. The result is in contrast to impressions gleaned from viewing observed poverty outcomes in 1985 and 1986. Our method shows potential difficulties faced by households that are obscured when simply using the historical record as a guide to the extent of possible vulnerabilities. The method is a first step in targeting all households with a reasonable risk of worsening poverty status, not just those that in fact lose out.

To sharpen policy implications, the approach can be extended to test whether a specific condition common to a given group is likely to make the group more vulnerable than others. In particular, the effectiveness of a poverty alleviation programme can be measured by comparing the pre- and post-programme vulnerability along the same lines as those of the scheme presented above.

With panel data for just two consecutive periods available, we have not taken on longer-term issues. Those issues are apt to matter greatly when considering vulnerability more richly, though. Concerns that will immediately arise include how to judge the persistence of poverty. Second, how should variability be traded-off against changes in the mean? In the CILSS, poorer households on average saw increases in consumption, even while country-wide average consumption fell. While the approach can accommodate alternative poverty measures (and, hence, different social weights), it cannot resolve the deeper conceptual issues surrounding the trade-offs between

---

10 Since the histograms of the proportional change in the bootstrap estimate of poverty headcount in all three panels show that they are very close to normally-distributed, we could easily calculate the bootstrap confidence interval for $\hat{P}_{\alpha_2}^{mcb}$. The 90 per cent confidence interval can be computed along the same lines as those of the bootstrap confidence interval, given that P_{α_1} is a constant (and thus that the variability of $V_{\alpha_2 t_2}^{mcb}$ depends only on that of $\hat{P}_{\alpha_2}^{mcb}$) and that $\hat{P}_{\alpha_2}^{mcb}$ is close to normal. The bootstrap confidence interval is the standard confidence interval (i.e. based on an asymptotic normal theory): $p \in \hat{p} \pm z^{(1-\alpha)} * \hat{se}$ with probability (1-2α), where $\hat{p}$ is the estimated poverty measure obtained from the original data, $\hat{se}$ is the bootstrap estimate of the standard error of the poverty measures obtained form the bootstrap samples ($\hat{p}^{mcb}$), and $z^{(1-\alpha)}$ is the $100(1-\alpha)th$ percentile of a standard normal distribution.

‘winners’ and ‘losers’ involved. Real progress will require consideration of issues at the intersection of moral philosophy, policy analysis, and the broader social sciences.

References

Ainsworth, M. and J. Munoz (1986), ‘The Cote d’Ivoire Living Standards Surveys: Design and Implementation’, Living Standards Measurement Study Working Paper No.26, World Bank, Washington DC.

Amin, Sajeda, Ashok Rai, Georgio Topa (1999), ‘Does Microcreedit Reach the Poor and Vulnerable? Evidence from Northern Bangladesh’, CID Working Paper No.28, October, Harvard University.

Anand, Sudhir and Jonathan Morduch (1999), ‘Poverty and the ‘Population Problem’, in Massimo Livi-Bacci and Gustavo de Santis (eds), *Population and Poverty in Developing Countries*, Oxford: Oxford University (Clarendon) Press.

Baulch, Bob and John Hoddinott (2000), ‘Economic Mobility and Poverty Dynamics in Developing Countries’, *Journal of Development Studies*, Vol.36, No.6, August, pp.1–24.

Chaudhuri, Shubham (2000), ‘Empirical Methods for Assessing Household Vulnerability to Poverty’, preliminary draft, Department of Economics, Columbia University, April.

Coulombe H. and L. Demery (1993), ‘Household Size in Cote d'Ivoire: Sampling Bias in the CILSS’, Living Standards Measurement Study Working Paper No.97, World Bank, Washington DC.

Deaton, Angus (1997), *The Analysis of Household Surveys*, Baltimore: Johns Hopkins.

Dercon, Stefan, (2001), ‘Income Risk, Coping Strategies and Safety Nets’, Centre for the Study of African Economies, Department of Economics, Oxford University (draft). See also WIDER Discussion Paper No.2002/22, Helsinki: UNU/WIDER.

Dercon, Stefan and Pramila Krishnan (2000), ‘Vulnerability, Seasonality and Poverty in Ethiopia’, *Journal of Development Studies*, Vol.36, No.6, August, pp.25–53.

Efron, Bradley and Robert J. Tibshirani (1994), *An Introduction to the Bootstrap*, Chapman & Hall/CRC.

Fields, Gary and Efe Ok (1999), ‘The Measurement of Income Mobility: An Introduction to the Literature’, in J. Silber, (ed.) *Handbook on Income Inequality Measurement*, Boston: Kluwer Academic Press, pp.557–96.

Foster, James, Joel Greer, and Erik Thorbecke (1984), ‘A Class of Decomposable Poverty Measures’, Econometrica 52, 761–65.

Grootaert, Christiaan and Ravi Kanbur (1995), ‘The Lucky Few Amidst Economic Decline: Distributional Change in Côte d’Ivoire as Seen through Panel Data Sets, 1985–88,’ *Journal of Development Studies*, Vol.31, No.4, pp.603–19.

Grootaert, Christiaan, Ravi Kanbur, and Gi-Taik Oh (1997), 'The Dynamics of Welfare Gains and Losses: An African Case Study', *Journal of Development Studies*, Vol.31, No.4, pp.635–57.

Husain, Ishrat and Rashid Faruqee (eds) (1994), *Adjustment in Africa: Lessons from Country Case Studies*, World Bank, Washington DC.

Jacoby, Hanan and Emmanuel Skoufias (1997), 'Risk, Financial Markets, and Human Capital in a Developing Country', *Review of Economic Studies*, Vol.64, No.3, July, 311–35.

Jalan, Jyotsna and Martin Ravallion (1998), 'Transient Poverty in Postreform Rural China', *Journal of Comparative Economics*, Vol.26, No.2, 338–57.

Jalan, Jyotsna and Martin Ravallion (1999), 'Are the Poor Less Well Insured? Evidence on Vulnerability to Risk in Rural China', *Journal of Development Economics*, Vol.58, No.1, 61–81.

Jarvis, Sarah and Jenkins, Stephen, 'How Much Income Mobility is There in Britain?', *Economic Journal*, Vol.108, No.447, March 1998, pp.428–43.

Jarvis, Sarah and Jenkins, Stephen, 'Low Income Dynamics in 1990s Britain', *Fiscal Studies*, Vol.18, No.2, May 1997, pp.1–20. Reprinted in *IDS Bulletin*, Vol.29, No.1, January 1998, 32–41; and in David Rose (ed.) *Researching Social and Economic Change*, London and New York: Routledge.

Kamanou, Gisele (1999), 'An Index of Household Material Wealth Based on Principal Components of Discrete Indicators: An inquiry into Family Support and Human Capital within the Household Dynamics during Structural Adjustment in Côte D'Ivoire', Ph.D. dissertation, Department of Statistics, University of California, Berkley.

Kanbur, Ravi and Lyn Squire (2001), 'The Evolution of Thinking about Poverty: Exploring the Interactions', in Gerald Meier and Joseph Stiglitz (eds), *Frontiers of Development Economics: The Future in Perspective*, New York: World Bank and Oxford University Press, pp.183–226.

Ligon, Ethan (2002), 'Targeting and Informal Insurance', WIDER Discussion Paper No.2002/8, Helsinki: UNU/WIDER.

Morduch, Jonathan (1991), 'Risk and Welfare in Developing Countries', Unpublished Ph.D. dissertation, Department of Economics, Harvard University.

Morduch, Jonathan, (1994), 'Poverty and Vulnerability', *American Economic Review*, Vol.84, May: 221–5.

Morduch, Jonathan (2000), 'Between the State and the Market: Can Informal Insurance Patch the Safety Net?', *World Bank Research Observer*, Vol.14, No.2, August, 187–207.

Morduch, Jonathan (2002), 'Consumption Smoothing Across Space: Testing Theories of Consumption Smoothing in the ICRISAT Study Region of South India', WIDER Discussion Paper No.2002/55, Helsinki: UNU/WIDER.

Narayan, Deepa, with Raj Patel, Kai Schafft, Anne Rademacher, and Sarach Koch-Schulte (2000), 'Voices of the Poor: Can Anyone Hear Us?', New York: World Bank and Oxford University Press.

Oh Gi-Taik and M. Venkataraman (1992), 'Construction of Analytical Variables and Data Sets Using the Data from the Cote d' Ivoire Living Standards Survey 1985–1988: Concept, Methodology and Documentation', World Bank, Washington DC.

Pritchett et al. (2000), 'Quantifying Vulnerability to Poverty: a Proposed Measure with Application to Indonesia', SMERU working paper, May 2000. [www.smeru.or.id]

Ravallion, Martin and Shubham Chaudhuri (1996), 'Risk and Insurance in Village India: Comment', *Econometrica*, Vol.65, No.1, pp.171–84.

Shorrocks, Anthony (1978), 'The Measurement of Mobility', *Econometrica*, Vol.46, pp.1013–24.

Townsend, Robert (1994), 'Risk and Insurance in Village India', *Econometrica*, Vol.62, pp.539–92.

Townsend, Robert (1995), 'Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Countries', *Journal of Economic Perspectives*, Vol.9, pp.83–102.

World Bank, World Development Report (2000/2001), 'Attacking Poverty', Oxford: Oxford University Press.