Session Number: Parallel Session 6B
Time: Tuesday, August 26, PM

Paper Prepared for the 31st General Conference of
The International Association for Research in Income and Wealth

St. Gallen, Switzerland, August 22-28, 2010

Welfare Volatility or Measurement Error? Some Implications for Chronic Poverty Measurement

Catherine Porter

For additional information please contact:

Name: Catherine Porter
Affiliation: University of Oxford

Email Address: catherine.porter@economics.ox.ac.uk

This paper is posted on the following website: http://www.iariw.org

Welfare Volatility or Measurement Error? Some Implications for Chronic Poverty Measurement

Catherine Porter- DRAFT not for citation*

July 28, 2010

Abstract

A number of chronic poverty measures have recently been proposed, and are beginning to be used for the purposes of targeting social programs. As with cross-sectional poverty measures, it is important to minimise errors of inclusion and exclusion. Many measures have the property of being sensitive to volatility in welfare, often as measured by consumption. For any given average level of consumption below the poverty line, a higher variance would entail a higher degree of poverty. However, variance in consumption may well be caused by measurement error. No research to date has considered the implications of measurement error for chronic poverty measurement. This paper considers the relationship between the underlying assumptions about substitution of welfare between time periods in the proposed measures in the literature and the direction of bias in the case of measurement error, using empirical illustrations. In the case of perfect substitution (e.g. using a simple average of consumption over time), and classical measurement error, intertemporal poverty is (slightly understated). If there is a duration cutoff, then measurement error could have quite serious consequences around the discontinuity (analagous to the poverty line in static measures). Further, if the intertemporal substitution is different above and below the poverty line, measurement error will have different consequences. And, if measurement error is not random or classically distributed, the implications may very well change.

*University of Oxford. Preliminary work. Thanks to Natalie Quinn and Gaston Yalonetsky for useful discussions. Correspondence to catherine.porter@economics.ox.ac.uk

1 Introduction

Conceptualising and quantifying poverty when data on many periods are available has grown in both academic and policy interest over the past ten years, due mainly to improved data availability. Conceptually, most economists agree that a snapshot of poverty is not necessarily representative and that most people see their economic welfare (however measured) change over time. Problems of choosing a suitable welfare indicator are complicated when the time dimension is included (e.g. income vs consumption vs BMI) as they may have different movements over time, (Günther and Klasen, 2009). A number of chronic poverty measures have recently been proposed in the literature (Jalan and Ravallion, 2000; Bossert, Chakravarty, and D'Ambrosia, 2008; Porter and Quinn, 2008; Calvo and Dercon, 2009; Foster and Santos, 2006; Foster, 2009). The choice of welfare indicator may lead to different requirements in terms of the theoretical properties a chronic poverty measure should satisfy (Calvo and Dercon, 2009; Porter and Quinn, 2008). A key issue is the amount of intertemporal substitution in welfare an intertemporal measure should incorporate. For example, if using income data, we may wish to allow a higher degree of intertemporal substitution in the measure, whereby 'richer' spells can compensate for a period in poverty, than if we have consumption data (given that we assume households have already smoothed their consumption to the best of their ability (Deaton, 1992). Such issues are conceptual and also ethical in nature. However, measurement error in the data leads also to more practical considerations, and some of these may be in direct conflict with desirable conceptual properties.

This paper examines such considerations, in particular we focus (for now) on two of the most well-used 'chronic' poverty measures in the policy literature, those proposed by Jalan and Ravallion (2000) and Foster (2009). The two measures are relatively straightforward extensions of a widely used static poverty measure, but the intertemporal aggregation is quite different, leading to distinct differences in the degree of intertemporal substitution inherent in the measures, and therefore their properties. In this preliminary work, we begin to sketch the theoretical ideas around the consequences of measurement error for

both measures, and also provide some empirical examples showing sensitivity of the two measures to a randomly generated measurement error, and to outliers in the original distribution.

2 Theory (I): Chronic Poverty Measures

In a seminal paper, Sen (1976) criticised the headcount measure which was prevalent at the time (and noticably still is) for shortcomings that are now well understood. His paper introduced desirable axioms of (static) poverty measures that are now taken as given in the literature. Whilst this paper has set the standard for the approach used in most other papers on the subject of poverty measurement, Sen's measure in the end has however not been popularly adopted, not least due to the fact that his measure is not decomposable by sub-groups. Notable contributions in the 1980's built on Sen's analysis[1] including Clark *et al.* (1981), Chakravarty (1983) and Atkinson (1987). Whilst a number of measures have been proposed, Foster *et al.* (1984) made what has probably been the most frequently utilised contribution to static poverty measurement,[2] with their class of decomposable poverty measures, known in the literature as the FGT measures, or p-alpha measures:

$$P_\alpha(\mathbf{x}) = \frac{1}{N}\sum_i \left(\frac{z - x_i}{z}\right)^\alpha \mathbb{I}(x_i \leq z) \tag{2.1}$$

We base the initial analysis on two intertemporal poverty measures that have been published in the recent literature; for clarity we write these measures following the notational convention defined above.

Jalan and Ravallion (2000)'s 'total poverty' measure P_{JRT} is essentially FGT-2 (squared poverty gap) aggregated over time periods as well as individuals,

[1] Zheng (1993) provides an axiomatic derivation of an index by Watts (1968) that had actually preceded Sen's analysis by several years. Zheng postulated that the lack of interest in the Watts index was due to this, though it may be more related to Foster (1994)'s argument. The Watts index has had some revival in interest in recent years mainly due to its relationship to exit-time from poverty. See, for example, Morduch (1998) and the recent literature on so-called pro-poor growth (Ravallion and Chen, 2003).

[2] E.g. number of citations on Google scholar=2621, Chakravarty (1983) was cited by 162, on 11 June 2010. See also Foster, Greer, and Thorbecke (2010).

$$P_{JRT}(X;z) = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\left(1-\frac{x_{it}}{z}\right)^2 \mathbb{I}(x_{it}\leq z) = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\left(\max\left[0,1-\frac{x_{it}}{z}\right]\right)^2 \tag{2.2}$$

while their 'chronic poverty' measure P_{JRC} is FGT-2 applied to each individual's mean wellbeing,

$$P_{JRC}(X;z) = \frac{1}{N}\sum_{i=1}^{N}\left(1-\frac{\bar{x}_i}{z}\right)^2 \mathbb{I}(\bar{x}_i\leq z) = \frac{1}{N}\sum_{i=1}^{N}\left(\max\left[0,\frac{1}{T}\sum_{t=1}^{T}\left(1-\frac{x_{it}}{z}\right)\right]\right)^2. \tag{2.3}$$

Calvo and Dercon (2009) note that this measure can be generalised to the p-alpha class of measures by substituting varying values of α for the squared term in equation 2.3. This measure has been widely used in the empirical literature to quantify chronic poverty, as it has a very intuitive interpretation: on average over time, the household has had welfare levels below the poverty line. It also has the practical advantage of being able to distinguish a distinct group of households with low welfare, in order to analyse their characteristics and make poverty profiles etc. However, it has been noted by several authors Foster and Santos (2006); Calvo and Dercon (2009); Porter and Quinn (2008) that this measure implies perfect subtitutability of welfare between periods: a very poor spell could be perfectly compensated for by a high enough period of welfare afterwards, and thus fails to meet Intertemporal Transfer property as outlined by Porter and Quinn.

Foster (2009)'s proposed class of measures are also an extension of the FGT suite as expressed in equation 2.1 and include an aggregation to a composite measure for a society. The methodology involves two focus steps. Firstly compute the headcount for each time period as in the static case. The duration headcount is then the sum of the static headcount over time. Secondly, set a 'duration poverty line' $0 > \tau \leq 1$, expressed as periods under the poverty line as a proportion of time measured (e.g. number of panel data obser-

vations available). The focus step *again* by counting only those below the poverty line *and* the duration poverty line. Equation 2.4 below shows the transformation is identical to the static P^α measures, but for households with fewer (as a proportion of the total) periods in poverty than τ, the consumption is censored to z. but he incorporates a 'poverty line' τ in the time dimension so that a household's wellbeings only enter if that household is below the wellbeing poverty line z in a proportion of periods greater than τ. (He also permits a flexible power parameter α; for comparability, and because it yields attractive properties, we shall take $\alpha = 2$; in our empirical section we take $\tau = 0.6$ or three out of five periods.)

$$P_F(X;z) = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\left(1-\frac{x_{it}}{z}\right)^{\alpha}\mathbb{I}(x_{it}\leq z)\mathbb{I}\left(\sum_{t=1}^{T}\mathbb{I}(x_{it}\leq z)\geq \tau T\right) \tag{2.4}$$

The two measures have been increasingly adopted used in the policy context, (e.g. Cruces and Wodon (2003), OPHI references).

Calvo and Dercon (2009) also proposed to use a similar measure, but calculated without the use of the duration poverty line. In this way, the poverty measure is the 'stock' of accumulated lifetime poverty, and the authors state that there is then no substitution of welfare between periods. However, in practice whilst this measure is useful for ordering households or individuals, it does not define a group of chronic poor which is less useful for policy purposes (Dercon and Porter, 2010).

Porter and Quinn (2008) define the property of **strong focus**, that the poverty measure is not sensitive to changes in wellbeing in any period above the poverty line, for any individual (even if that individual's wellbeing lies below the poverty line in other periods). The property of **Non-Decreasing Compensation** states that the marginal rate of intertemporal compensation between an individual's welfare in two periods should not decrease, as the period wellbeings increase in proportion. Equivalently, the elasticity of compensation should not decrease as wellbeing increases. Porter and Quinn show that 1. Strong focus and increasing compensation are incompatible properties. 2. The

only trajectory ordering satisfying strong focus and non-decreasing compensation is the 'Rawlsian' ordering.

Jalan-Ravallion (2000) does not satisfy **strong focus**, since all welfare levels in all periods are averaged. Foster (2009) does not satisfy **Non-Decreasing Compensation**, since all periods above the poverty line have zero intertemporal compensation (and are by definition higher than those below the poverty line).

2.1 Exploring measurement error

In this section, we begin to explore the consequences of measurement error for the two types of chronic poverty measure discussed above. We begin with a very simple specification for the distribution of measurement error, and in future versions of the paper, we will explore the consequences of non-normally distributed measurement error. In the empirical section we present some simulations where we add measurement error to a dataset that we assume is correctly measured, in order to assess the differences. We note at this point though that in theory measurement error should be a less serious problem for an intertemporal poverty measure (i.e. one which aggregates poverty over time) than either a *static* poverty measure, or a measure of *intertemporal* mobility.

Results that were used in Glewwe (2007) can be expanded on to understand the consequences of the presence of measurement error for the Jalan and Ravallion (2000) measure of chronic poverty. We follow Glewwe (2007) in specifying actual income or consumption x^* as having a lognormal distribution:

$$log(x^*) \sim N(\zeta, \sigma^2_{x^2}) \tag{2.5}$$

and also assume that the observed value of consumption, x, is the product of x^* and a random measurement error that also follows a lognormal distribution:

$$x = x^* \varepsilon, \varepsilon > 0 \tag{2.6}$$

and

$$\epsilon \equiv log(\varepsilon) \sim N(0, 1) \tag{2.7}$$

Since the median of ϵ is zero, the median of ε is one. Because both x^* and ε follow a lognormal distribution, x does as well.

$$log(x^*) = log(x) + log(\varepsilon) \sim N(\zeta, \sigma^2_{x^*} + \sigma^2_\epsilon) \tag{2.8}$$

Glewwe also shows that the mean of x is in fact higher than that of x^*:

$$E[x^*] = exp(\zeta + \sigma^2_{x^*}/2) \tag{2.9}$$

$$E[x]) = exp(\zeta + (\sigma^2_{x^*} + \sigma^2_\epsilon)/2) \tag{2.10}$$

In this case then the Jalan-Ravallion estimator of chronic poverty will underestimate chronic poverty only very slightly, by a factor of only $exp(\sigma^2_\epsilon)/2)$. Given that ϵ is standard normal distribution, the factor is equal to a constant factor of 1.64.

In the case of Foster's (2009) measure, the consequences of measurement error may be more serious. There is no intertemporal substitution of welfare for periods above the poverty line. The measure is therefore indifferent to measurement error in periods above the poverty line *iff* measurement error does not push the welfare level below the poverty line. Porter and Quinn (2008) note that a non-continuous poverty measure would be excessively sensitive to measurement error at any point of discontinuity. Measurement error that pushes a 'classified poor' household (i.e. for $\tau - 1$ periods, where τ is the cutoff) above the poverty line in only one period would lead them to be misclassified

as non-chronic poor, and vice versa. The subsequent aggregation of chronic poverty to society includes a step where poverty shortfalls are censored to zero for the case of the subsample who are poor, but for fewer than τ periods.

Proof: sketch outline of discontinuity - possible asymmetric effect of symmetric distributed measurement error?

We see in the empirical results below that changing the duration poverty line in a fairly short panel can change the results quite dramatically (and increase by 2/3 the amount of poor people).

3 Measurement Error- How much of an issue?

Bound, Brown, and Mathiowetz (2001) outline the consequences of measurement error for various estimations and uses. To our knowledge, for developing countries there have been no validation studies of the type done using US data, that have re-surveyed the information using different methods (e.g. the Panel Study of Income Dynamics Validation Study (see Bound et al. (1990,1994) or the Current Population Survey-Social Security Earnings Records Exact Match File (Bound and Krueger (1991)). In the US data, research suggests that measurement error in earnings violates classical assumptions. However as Glewwe (2007) points out, there is no reason to assume that the form that measurement error in US earnings data takes would be in any way representative of that in other settings.

Studies on US data on poverty: Worts, Sacker, and McDonough (2010) find that that correcting for error the US and UK using panel data shows poverty as less temporary and risks as less widely dispersed than otherwise assumed, and that cross-national differences are more pronounced. McGarry (1995) also using US data that observation error causes poverty rate to be overestimated by around two percentage points on average. However, eliminating observation error increases the probability of transiting either into or out of poverty. The author finds that reductions imply that the amount of permanent poverty

is underestimated when measurement error is ignored.

Glewwe 2004(?) finds using the Vietnam Living Standards Survey at least one third of measured mobility in per capita expenditures is due to measurement error, and that about 13% of measured inequality is also due to measurement error.

There are several reasons to use consumption rather than incomes as a welfare indicator in developing countries. Deaton (1991), Deaton and Zaidi (2002) and Grosh, Glewwe, and Bank. (2000) discuss the issues involved. Less likely to be systematic, less likely to underreport for fear of tax collection or loss of benefits. Whilst still likely to be measured with error, it is less likely to be systematically underestimated e.g. for tax purposes, though there is of course a chance that it could be systematically underestimated if the households perceive the survey to be possibly scoping out beneficiaries for a support programme. Glewwe (2007) notes that in many developing countries earned income may be zero if non monetary income is not accounted for.

Glewwe (2007) finds that [when the (log) measurement errors have a mean of zero and are uncorrelated with income]: (i) measurement error leads to underestimation of the mean income of the poor at any point in time, (ii) increases (decreases) in measurement error over time, for a given level of inequality, lead to underestimation (overestimation) of income growth among the poor, and (iii) increases (decreases) in inequality over time, for a given level of measurement error, lead to overestimation (underestimation) of income growth among the poor.

We now turn to some empirical results in order to pursue the issue further. The next section describes the panel data used, and then we present some preliminary results. The final section outlines further work to be done in this area.

4 Data

Data are from the Ethiopian Rural Household Survey (ERHS) collected by the University of Addis Ababa and the Centre for the Study of African Economies (CSAE) at the Uni-

versity of Oxford, as well as the International Food Policy Research Institute (IFPRI). Data are available from fifteen districts[3] in several regions. Seven villages were originally included in IFPRI's survey of 1989, which were chosen primarily because they had suffered hardships in the period 1984–89 (the 1984–85 famine). For a detailed description see Webb, von Braun, and Yohannes (1992). In 1994, 360 of the households in six villages were retraced and the sample frame was expanded to 1477 households. The nine additional communities were selected to account for the diversity in the farming systems in the country. Within each village, random sampling was used. The households were resurveyed again in 1994 and 1995, and subsequently in 1997 and 1999. The sixth and latest round of the survey was completed in 2004. The attrition rate is low, less than one per cent per annum (annualised, or 12.1% in total between 1994 and 2004). Since the three surveys in 1994-1995 are within eighteen months of each other, we drop the second round of 1994, in order to use five rounds of the data for our subsequent analysis.

The dependent variable or welfare measure chosen is real household monthly consumption per adult-equivalent. This is comparable with other studies of consumption and poverty that have been conducted on the dataset, and other studies of poverty. In the ERHS, detailed information is also available on household income and assets. At the individual level, there is information on height and weight though not for all individuals and not for all rounds. Data on monthly consumption of food, purchased food and non-investment non-food items (excluding durables, as well as health and education expenditure) based on purchased items, gifts in cash and in kind, and a diary of consumption from own production from a two-week recall period was divided by adult equivalent units based on World Health Organisation (WHO) guidelines. This was deflated by a food price index constructed from data collected for each village at the same time as the household survey. For a detailed discussion on the construction of the consumption indicator, see Dercon and Krishnan (1998). Food represents around eighty per cent of the consumption basket. We use a consumption poverty line calculated by Dercon and Krishnan (1998) which is

[3] These communities are called Woredas, the equivalent of a county in the UK. They are further divided into Peasant Associations (PAs), the equivalent of a village, and consist of up to several villages (e.g. the ERHS comprises 15 Woredas, and 18 PAs). The administrative system of the PAs was created in 1974 after the revolution.

village specific (according to local prices) and averages 44.3 Birr per adult equivalent per month (all prices are specified in 1994 real terms). It is based on the monthly cost of the diet to achieve 2100 Kcal per day per adult, using the food consumed by the poorer half of the sample, and minimum 'essential non-food' expenditure.[4] For analytical convenience, we include only households in the analysis that have nonmissing values of consumption for each available round, restricting the sample to 1168 households. We also ignore changes in household composition, apart from in the calculation of consumption per adult equivalent, and therefore abstract somewhat from measuring individual welfare. In the ERHS, the only possibility to examine individual wellbeing would be through an anthropometric measure, which is a focus for future work.

5 Empirical results

Table 1: Standard poverty measures, by round

Year	Headcount F-G-T (0)	Poverty Gap F-G-T (1)	Sq-Poverty Gap F-G-T (2)	Mean cons. of poor (Birr)
1994	0.39	0.16	0.09	26.84
1995	0.44	0.19	0.10	25.72
1997	0.24	0.08	0.04	30.31
1999	0.28	0.09	0.04	31.13
2004	0.22	0.07	0.03	30.29

Notes: Source is ERHS data, own calculations. Poverty line is 44.3 Birr per adult, per month on average (though varies by village). Number of observations is 1168. Measures are weighted by household size.

Foster (2009) has a flexible duration poverty line, which alters the proportion of poor households considerably, as can be understood from table 2: just under ten percent of households are poor for four or five periods, rising to 21.5% if we include households that are poor three or more periods. Jalan and Ravallion (2000) set the chronic poverty identification step as those households whose arithmetic mean consumption is below the

[4]Dercon and Krishnan (1998) show the calculations in depth. Also note that the household was considered the same between rounds if the head of the household was unchanged, while if the head had died or left the household, the household was considered the same if the current household head acknowledged that the household (in the local meaning of the term) was the same as in the previous round. In 1994 the exchange rate of the Ethiopian Birr (ETB) with the USD, corrected for purchasing power parity was 0.21875. One PPP dollar is worth 1.75 its equivalent in USD [World Bank (2000)], and the exchange rate was 8 ETB to 1 USD. Therefore the Dollar a Day poverty line (1.08) would be calculated as approximately 148 Birr per adult per month, around three times that which we use here.

Table 2: Poverty Duration (Nr times hh classified as poor)

Counted poor	Nr households	Percent	Cum. Percent
Never	388	35.63	35.63
Once	277	25.44	61.07
Twice	190	17.45	78.51
Thrice	128	11.75	90.27
Four times	78	7.16	97.43
In every period	28	2.57	100
Total	1,089	100	

Notes: Cumulative percent works backwards: i.e. the number of households poor in *at least* n periods

Table 3: Poverty Duration (Nr times hh classified as poor - with measurement error)

Counted poor	Nr households	Percent	Cum. Percent
Never	163	14.97	14.97
Once	342	31.4	46.37
Twice	264	24.24	70.62
Thrice	198	18.18	88.8
Four times	95	8.72	97.52
In every period	27	2.48	100
Total	1,089	100	

Notes: Cumulative percent works backwards: i.e. the number of households poor in *at least* n periods

poverty line, which is 13.4% of households, a relatively low proportion, which is due to the high degree of substitution between periods.

If we consider in table 3 the number of periods in poverty when simulated measurement error is added, it appears that this changes most for the households who are not often poor: the proportion never poor halves from just over 30% to only 15%. At the opposite end, the proportion always poor changes only marginally.

Table 4 compares the two poverty measures, with $(alpha = 0, 1, 2)$ to see what the main differences are. The first column uses the original data, second column incorporates the simulated measurement error (with the ln measurement error having a standard normal distribution), and the final column censors outliers at the 5th and 95th percentile of the

distribution.[5]

Table 4: Poverty measures sensitivity to measurement error and outliers)

Jalan-Ravallion			
	Measure	Added m.e.	Outliers removed
Alpha=0	0.1341	0.1028	0.1331
Alpha=1	0.0298	0.0284	0.0294
Alpha=2	0.0100	0.0116	0.0098
Foster (d3)			
Alpha=0	0.1535	0.2037	0.1535
Alpha=1	0.0611	0.1048	0.0606
Alpha=2	0.0331	0.0684	0.0324

Notes: Measures defined in the text.

Adding the simulated measure increases the Foster measure, but decreases the Jalan-Ravallion measure. Proportionally, the Foster measure is more affected. This is as expected from our theoretical sketchings, given the discontinuity in the measure.

Outliers can change the results of any estimation if extreme enough (Wooldridge, 2009). Whilst it is true that if a person wins the lottery (or something less dramatic, but significant, such as gaining a job in the formal sector) in one time period, they are likely no longer poor, this should show in the subsequent data (e.g. consumption should increase significantly in all subsequent time periods).

The probability that an extremely high or low value is due to measurement error is probably higher than that of something around the mean? OR that measurement is likely worse (higher) for observations around the tails of the distribution.

The trimming of outliers does not appear to affect any of the measures significantly. We also recalculated the village poverty profile censoring the consumption data at the 5th and 95th percentiles. Jalan-Ravallion showed virtually no change, due to the high level of substitution inherent in this measure. Surprisingly, neither did Foster measure. Though, the upper tail is never included in the Foster measure. The lower tail may not be likely to push any household across the discontinuity [formalise this].

Data reliability should be a consideration when choosing the amount of intertemporal

[5]We censor rather than truncate the data as this would change our sample.

compensation that is required in the measure, and it is always important to conduct sensitivity analysis.

6 Conclusions and future research

Desirable normative properties of poverty measures in the intertemporal context may be in direct conflict with desirable practical properties. Choosing the welfare indicator, poverty line, and chronic poverty measure all include value judgements and it is important to be explicit what are the normative and practical consequences of decisions.

The estimation of poverty in the intertemporal context is already enhanced by the use of panel data, and either measure discussed performs better than a 'snapshot' or cross-section measure of poverty. However, it appears that the Foster measure is more sensitive to measurement error.

Issues to explore in future work:

- Formalising more proofs of the theoretical consequences
- Extreme values - if they are real then we do wish to weight them heavily - is there any way to distinguish
- easurement error that is correlated with the level of consumption/income?

References

Atkinson, A. B. (1987): "On the Measurement of Poverty," *Econometrica*, 55(4), 749–764.

Bossert, W., S. R. Chakravarty, and C. D'Ambrosia (2008): "Poverty and Time," Discussion Paper 87, ECINEQ, Society for the Study of Economic Inequality.

Bound, J., C. Brown, and N. Mathiowetz (2001): "Measurement error in survey data," in *Handbook of Econometrics*, ed. by J. Heckman, and E. Leamer, vol. 5 of *Handbook of Econometrics*, chap. 59, pp. 3705–3843. Elsevier.

Calvo, C., and S. Dercon (2009): "Chronic Poverty and All That: the Measurement of Poverty over Time," in *Poverty Dynamics: Interdisciplinary Perspectives*, ed. by T. Addison, D. Hulme, and R. Kanbur. Oxford University Press US, Washington, DC.

Chakravarty, S. R. (1983): "Ethically flexible measures of poverty," *Canadian Journal of Economics*, 16(1), 74.

Clark, S., R. Hemming, and D. Ulph (1981): "On Indices for the Measurement of Poverty," *The Economic Journal*, 91(362), 515–526.

Cruces, G., and Q. T. Wodon (2003): "Transient and chronic poverty in turbulent times: Argentina 1995-2002," *Economics Bulletin*, 9(2), 1–12.

Deaton, A. (1991): *Household saving in LDC's : credit markets, insurance and welfare*. Research Program in Development Studies Woodrow Wilson School Princeton University, Princeton, N.J.

——— (1992): *Understanding consumption*, Clarendon lectures in economics. Clarendon Press, Oxford.

Deaton, A., and S. Zaidi (2002): "Guidelines for constructing consumption aggregates for welfare analysis," Discussion Paper 0821349902, World Bank, Angus Deaton and Salman Zaidi. ill. ; 28 cm.

Dercon, S., and P. Krishnan (1998): "Changes in Poverty in Rural Ethiopia 1989-1995: Measurement, Robustness Tests and Decomposition," Discussion Paper WPS/98-7, Centre for the Study of African Economies.

Foster, J. (2009): "A Class of Chronic Poverty Measures," in *Poverty Dynamics: Interdisciplinary Perspectives*, ed. by T. Addison, D. Hulme, and R. Kanbur. Oxford University Press US, Washington, DC.

Foster, J., J. Greer, and E. Thorbecke (1984): "A Class of Decomposable Poverty Measures," *Econometrica*, 52(3), 761–766.

——— (2010): "The FosterGreerThorbecke (FGT) poverty measures: 25 years later," *Journal of Economic Inequality*, online, May.

Foster, J., and M. E. Santos (2006): "Measuring Chronic Poverty," .

Foster, J. E. (1994): "Normative measurement: Is theory relevant?," *American Economic Review*, 84(2), 365.

Glewwe, P. (2007): "Measurement Error Bias in Estimates of Income and Income Growth among the Poor: Analytical Results and a Correction Formula," *Economic Development and Cultural Change*, 56(1), 163–189.

Grosh, M. E., P. Glewwe, and W. Bank. (2000): *Designing household survey questionnaires for developing countries : lessons from 15 years of the living standards measurement study*. World Bank, Washington, D.C.

Günther, I., and S. Klasen (2009): "Measuring Chronic Non-Income Poverty," in *Poverty Dynamics: Interdisciplinary Perspectives*, ed. by T. Addison, D. Hulme, and R. Kanbur. Oxford University Press US, Washington, DC.

Jalan, J., and M. Ravallion (2000): "Is Transient Poverty Different? Evidence for Rural China," *Journal of Development Studies*, 36(6), 82.

McGarry, K. (1995): "Measurement Error and Poverty Rates of Widows," *The Journal of Human Resources*, 30(1), 113–134.

Morduch, J. (1998): "Poverty, economic growth, and average exit time," *Economics Letters*, 59(3), 385–390.

Porter, C., and N. Quinn (2008): "Intertemporal Poverty Measurement: Tradeoffs and Policy Options," .

Ravallion, M., and S. Chen (2003): "Measuring pro-poor growth," *Economics Letters*, 78(1), 93–99.

Sen, A. (1976): "Poverty: An Ordinal Approach to Measurement," *Econometrica*, 44(2), 219–231.

Watts, H. (1968): "An economic definition of poverty," in *On understanding poverty*, ed. by D. P. Moynihan. Basic Books, New York.

Webb, P., J. von Braun, and Y. Yohannes (1992): "Famine in Ethiopia: Policy Implications of Coping Failure at National and Household Levels," Discussion Paper 92, International Food Policy Research Institute (IFPRI).

Wooldridge, J. M. (2009): *Introductory econometrics : a modern approach*. South-Western, Australia ; United Kingdom, 4th edn.

Worts, D., A. Sacker, and P. McDonough (2010): "Re-Assessing Poverty Dynamics and State Protections in Britain and the US: The Role of Measurement Error," *Social Indicators Research*, 97(3), 419–438.

Zheng, B. (1993): "An axiomatic characterization of the Watts poverty index," *Economics Letters*, 42(1), 81–86.